Online, Time-Varying and Multi-Period Optimization with Applications in Electric Power Systems

by

Julie Mulvaney-Kemp

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Javad Lavaei, Chair
Professor Shmuel Oren
Professor Duncan Callaway

Fall 2022

Online, Time-Varying and Multi-Period Optimization with Applications in Electric Power Systems

Abstract

Online, Time-Varying and Multi-Period Optimization with Applications in Electric Power Systems

by

Julie Mulvaney-Kemp

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Professor Javad Lavaei, Chair

Decision-makers often face environments which vary over time and deal with uncertainty as a result. Problems with temporal variation differ based on how frequently they are solved and what information about the present or future is available to the decision-maker; each class of problem poses distinct challenges and requires tailored solutions. In this dissertation, we design and analyze methods to support decision-makers in these settings, with an emphasis on applications in electric power systems.

Online optimization, in which a player makes a sequence of decisions aimed at minimizing the damage inflicted by their adversary, is the first class of problems considered. Here, we shed light on online nonconvex optimization problems in which algorithms are evaluated against the optimal decision at each time using the notion of dynamic regret. The adversary's loss functions are arbitrarily nonconvex but have global solutions that are slowly time-varying. To address this problem we first analyze the region around the global solutions to define time-varying target sets. Then, we introduce two algorithms and prove that the dynamic regret for each algorithm is bounded by a function of the temporal variation in the optimal decision. The first algorithm assumes the decision-maker has some prior knowledge about the initial loss function. The second algorithm makes no assumption about prior knowledge, and instead it relies on random sampling and memory to find and then track the target sets over time. In this case, the landscape of the loss functions determines the likelihood that the dynamic regret will be small. Numerical experiments validate these theoretical results and highlight the impact of a single low-complexity problem early in the sequence.

Time-varying optimization, in which a series of linked problems are solved sequentially using information about past decisions, but not considering the future, is the second class of problems considered. Specifically, we analyze the optimality behavior of solution trajectories for optimal power flow (OPF) with time-varying load. Despite its nonlinearity, time-varying OPF is commonly solved every 5-15 minutes using local-search algorithms. Failing to obtain

the globally optimal solution of power optimization problems jeopardizes the grid's reliability and causes poor financial and environmental outcomes. An empirical study on California data shows that, with enough variation in the data, local search methods can solve OPF to global optimality, even if the problem has many local minima. To explain this phenomenon, we introduce a backward mapping that relates the time-varying OPF's global solution at a given time to a set of desirable initial points. We show that this mapping could act as a stochastic gradient ascent algorithm on an implicitly convexified formulation of OPF, justifying the escape of poor solutions over time. This work is the first to mathematically explain how temporal data variation affects the complexity of solving power operational problems.

Multi-period optimization under uncertainty, in which decisions for the entire time horizon are made at once, is the third and final class of problems considered. Within this class, we focus on the decision-making problem facing hybrid power plant operators when participating in wholesale electricity markets with the goal of enabling researchers to accurately incorporate these resources into simulations of future electricity systems. To this end, a stochastic optimization model is developed to maximize the revenue of the plant, which consists of a renewable generator and an energy storage resource that appear as a single, combined unit to the market, while limiting risk due to uncertainty in market prices and generation levels, through price-quantity bid curves. The uncertainty is represented by scenarios, and a detailed methodology is provided for creating scenarios which both reflect the type of information a hybrid operator would have and only require the limited data which is available in simulation settings. This work advances existing models by allowing greater cooperation between the generator and storage components while also enforcing market limits on bid curves' complexity. Further, the approach ensures economically-sound behavior even when faced with unforeseen events.

To my family, for ensuring every educational opportunity was available to me.

To Bea, Kelsey and Lucy, for always being in my corner.

Contents

Contents **ii**

List of Figures **iv**

List of Tables **v**

1 Introduction **1**
1.1 Related Publications 3

2 Dynamic Regret Bounds for Constrained Online Nonconvex Optimization Based on Polyak-Łojasiewicz Regions **4**
2.1 Introduction 4
2.2 Theoretical Results for a Fixed Time Step 7
2.3 Online Projected-Gradient Descent with Random Exploration 13
2.4 Conclusion 23

Appendices **26**
2.A Proof of Theorem 2.2 26
2.B Proof of Lemma 2.1 28
2.C Capture property 28

3 Smoothing Property of Load Variation Promotes Finding Global Solutions of Time-Varying Optimal Power Flow **30**
3.1 Introduction 30
3.2 Empirical Study of Time-Varying OPF 32
3.3 Mathematical Analysis of Time-Varying OPF 36
3.4 Connection to Stochastic Gradient Ascent 41
3.5 Conclusion 49

Appendices **51**
3.A Proof of Theorem 3.1 51
3.B Proof of Proposition 3.1 55

4 Hybrid Power Plant Bidding in Models of Future Electricity Systems 57
4.1 Introduction 57
4.2 Hybrid Bidding Problem Formulation 60
4.3 Simulating Hybrid Bids Within a Future Electricity Market 73
4.4 Conclusion 85

Bibliography 88

List of Figures

2.1 Examples of functions which satisfy the PL inequality on all or part of their domain 8
2.2 Visualization of the proximal-PL region and the target set for an example optimization problem 14
2.3 Contour plots of objective functions used in the empirical study of Algorithm 2.2 23
2.4 Empirical validation of Theorem 2.3 probability bound 24
2.5 Empirical regret resulting from Algorithm 2.2 25

3.2.1 Average daily net load for California during August 2019 33
3.2.2 Solution convergence for points on discrete local trajectories 34
3.2.3 Cost for points on discrete local trajectories and point-wise local solutions, relative to the cost of the best trajectory 35
3.2.4 Real and reactive power output of select generators: points on discrete local trajectories and point-wise local solutions 36
3.2.5 Data and results for an empirical study on a 9-bus system 37
3.3.1 Three scenarios highlighting the role of load variation 38
3.4.1 Diagram of the 2-bus system. 45
3.4.2 The effect of load variation on the landscape of a 2-bus system 47
3.4.3 Numerical studies on a 39-bus system 50

4.1.1 Diagrams of participation models for hybrid resources 59
4.2.1 Bidding curve illustration for a single hour highlighting the role of price points in the grouping of scenarios 70
4.3.1 Forecast error correlation for a wind plant in New York for 2019 75
4.3.2 Example of generation scenario development 77
4.3.3 Example of day-ahead price scenario development 79
4.3.4 Example of analysis used to determine the number of generation scenarios . . . 82
4.3.5 Illustration of solution to initial state of charge estimation challenge 82
4.3.6 Sensitivity analysis of the initial state-of-charge parameter 83
4.3.7 Illustration of the high-price event heuristic in Algorithm 4.4 85
4.3.8 Illustration of the low-price event heuristic in Algorithm 4.5 85

List of Tables

2.1 Parameter and constant values for the empirical study of Algorithm 2.2 22

4.2.1 Overview of recent works on stand-alone storage bidding or scheduling 62
4.2.2 Overview of recent works on hybrid bidding or scheduling 63
4.3.1 Structure of the generation data available to the hybrid operator 74
4.3.2 Values used to define and select state-of-charge dependent bidding curve sets . . 84

Acknowledgments

I have many people to thank for their guidance and support of this research and my academic career.

First, I would like to acknowledge my advisor, collaborators, and committee members. I want to thank Javad Lavaei, my advisor, for introducing me to interesting and challenging problems in optimization and providing direction as I worked to solve them. Further, I am grateful to him for assembling a research group of collaborative, bright, and dedicated scholars whom I was fortunate to observe and work alongside. To my collaborators, Salar Fattahi, SangWoo Park, Ming Jin, Andrew Mills, Will Gorman, Miguel Heleno, and Dev Millstein, thank you for your willingness to share and to teach and for trusting me. As members of my dissertation committee, Shmuel Oren and Duncan Callaway provided advice and encouragement at pivotal times, and I am thankful for their support.

Second, I would like to thank many more individuals who have had a positive impact on my graduate school experience: IEOR colleagues Baturalp, Caleb, Erik, Han, Igor, Ilgin, Mahan, Pelagie, Ruijie, Ruojie, Tomas, Yann, Yoon, and Yuhao for sharing part of your journey with me, brightening my day countless times, and being teachers in addition to peers. IEOR staff Heather, Keith, and Rebecca for your timely help and care. The Electricity Markets & Policy department at Berkeley Lab for welcoming and supporting me fully. My fellow student workers in UAW 2865 and Student Researchers United for your solidarity.

Finally, I would like to acknowledge the funding sources which supported this research: UC Berkeley Chancellor's Fellowship, Department of Energy's Office of Energy Efficiency & Renewable Energy and Office of Electricity, and grants from ARO, ONR, AFOSR, and NSF.

Chapter 1

Introduction

Decision-makers regularly face environments which evolve over time and deal with uncertainty as a result. Problems with temporal variation differ based on how frequently they are solved, what information about the present or future is available to the decision-maker, how actions at one time influence the landscape of the future, and whether that landscape is designed by an adversary, among other dimensions. Each class of problem poses distinct challenges and requires tailored solutions. In this dissertation, we design and analyze methods to support decision-makers in these time-varying settings, with an emphasis on applications in electric power systems.

The electricity sector is a compelling and timely focus area because it is a vast complex network of infrastructure, policies and actors that is undergoing a transformation. As generation profiles shift away from a few large conventional power plants toward a multitude of smaller distributed energy resources, the size of problems facing system operators increases. Further, uncertainty within these problems is growing as variable renewable energy sources comprise increasing shares of generation capacity and climate change makes weather, and therefore electricity generation from these sources, less predictable [107]. To this field, this dissertation contributes analysis which may help to find better system operating points more quickly and a tool which enables researchers to incorporate a new type of participant in models of future electricity markets.

The decision-making problems studied in this dissertation can be posed as static optimization problems in the form of

$$\begin{aligned} \mathcal{P}(\theta) \; := \; & \min_{x \in \mathbb{R}^n} f(x; \theta) \\ & \text{subject to } x \in \mathcal{X}(\theta) \end{aligned} \tag{1.1}$$

or sequences of such problems, $\{\mathcal{P}(\theta_t)\}_{t=1}^{T}$. Here, x is the n-dimensional decision variable which the decision-maker selects from the feasible space $\mathcal{X}(\theta)$ with the objective of minimizing the function $f(\,\cdot\,;\theta) : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$. The vector $\theta \in \mathbb{R}^m$ represents the problem parameters which may or may not be known to the decision-maker and may directly or

indirectly influence the problem formulation. The following chapters each address a class of problems within this framework:

- **Chapter 2, "Dynamic Regret Bounds for Constrained Online Nonconvex Optimization Based on Polyak-Łojasiewicz Regions,"** focuses on online nonconvex optimization problems. In this class of problem sequences, the feasible space is known to the decision-maker and independent of θ, that is, $\mathcal{X}(\theta_1) = \mathcal{X}(\theta_2) = \cdots = \mathcal{X}(\theta_T) = \mathcal{X}$. When selecting x_t, the decision-maker is oblivious to θ_t and therefore $f(\,\cdot\,;\theta_t)$, as well, but has information about $\{\theta_k\}_{k=1}^{t-1}$. In fact, θ_t may be chosen by an adversary after x_t is selected. However, the adversary's choice of θ_t is limited in the following way: the optimal solution to $\mathcal{P}(\theta_t)$ must be in a neighborhood of the optimal solution to $\mathcal{P}(\theta_{t-1})$. This chapter analyzes the nonconvex landscapes of $\{\mathcal{P}(\theta_t)\}_{t=1}^{T}$, proposes algorithms for selecting x_t based on knowledge of θ_{t-1}, and proves theoretical bounds on the algorithms' performance.

- **Chapter 3, "Smoothing Property of Load Variation Promotes Finding Global Solutions of Time-Varying Optimal Power Flow,"** focuses on time-varying optimal power flow problems, an important problem sequence for power systems operations in which electricity supply and demand are matched throughout the network, ideally in the lowest-cost way, while respecting various physical, operational, and security constraints. In this setting, θ_t, and therefore $f(\,\cdot\,;\theta_t)$ and $\mathcal{X}(\theta_t)$, depends on the choice of x_{t-1} and is known to the decision-maker when selecting x_t. Each problem, $\mathcal{P}(\theta_t)$, is difficult and decisions must be made expeditiously, so the decision-maker may be resigned to selecting a locally-optimal decision, instead of a globally optimal decision. Since each local solution to $\mathcal{P}(\theta_1)$ spawns a different instance of $\mathcal{P}(\theta_2)$, there are multiple possible time-varying problem sequences and corresponding locally-optimal decision sequences (also called trajectories). This chapter explores local solution trajectories for time-varying optimal power flow problems both empirically and analytically and explains how temporal variation in $\{\theta_t\}_{t=1}^{T}$ enables these trajectories to avoid poor solutions over time.

- **Chapter 4, "Hybrid Power Plant Bidding in Models of Future Electricity Systems,"** focuses on a static multi-period problem of making bidding decisions in an electricity market. The market participant operates a "hybrid" power plant composed of a renewable generator and a battery. In this setting, decisions for the entire time horizon are made in advance and the single objective function $f(\,\cdot\,;\theta)$ is the summation of costs at each time step. Here, θ is stochastic due to uncertainty in future weather conditions and market prices. This chapter provides a methodology for designing a discrete set of scenarios which approximate the distribution of θ and formulates $\mathcal{P}(\theta)$ as a stochastic optimization problem. Further, the approach manages risk and ensures economically-sound bidding behavior even when θ falls outside of the developed scenario set, as may occur during a low-probability extreme event.

Each chapter is designed to be self-contained – they each contain background, motivation, and key definitions for the topic at hand – and can be read independently, if desired.

1.1 Related Publications

This dissertation includes material that has been previously published in the following locations:

- **Chapter 2**
  Main paper:

  1. Julie Mulvaney-Kemp, Salar Fattahi and Javad Lavaei, "Smoothing Property of Load Variation Promotes Finding Global Solutions of Time-Varying Optimal Power Flow," *IEEE Transactions on Control of Network Systems*, 2021

  Related paper:

  1. Julie Mulvaney-Kemp, Salar Fattahi and Javad Lavaei, "Load Variation Enables Escaping Poor Solutions of Time-Varying Optimal Power Flow," *IEEE Power & Energy Society General Meeting (PESGM)*, 2020
     - PESGM Best-of-the-Best Paper Award - Power System Operations, Planning, and Economics, 2020
     - INFORMS Energy, Natural Resources and the Environment (ENRE) Best Student Paper Award, 2020

- **Chapter 3**
  Main paper:

  1. Julie Mulvaney-Kemp, SangWoo Park, Ming Jin and Javad Lavaei, "Dynamic Regret Bounds for Constrained Online Nonconvex Optimization Based on Polyak-Lojasiewicz Regions," *IEEE Transactions on Control of Network Systems*, 2022

  Related paper:

  1. SangWoo Park, Julie Mulvaney-Kemp, Ming Jin and Javad Lavaei, "Diminishing Regret for Online Nonconvex Optimization," *American Control Conference (ACC)*, 2021

Chapter 2

Dynamic Regret Bounds for Constrained Online Nonconvex Optimization Based on Polyak-Łojasiewicz Regions

2.1 Introduction

Nonconvex optimization is ubiquitous in real-world applications, such as the training of deep neural nets [76], matrix sensing/completion [17, 16], state estimation of dynamic systems [101], and the optimal power flow problem [132]. Moreover, most of these practical problems are solved sequentially over time with time-varying input data, leading to online (real-time) versions of the aforementioned examples [6, 101, 112].

In this paper, we study an online nonconvex optimization (ONO) problem whose loss (objective) function changes over discrete time periods, namely,

$$\underset{\mathbf{x}\in\mathbb{S}}{\text{minimize}} \quad f_t(\mathbf{x}) \tag{2.1}$$

where $t \in \mathbb{Z}_+$ denotes the time and $\mathbb{S} \subseteq \mathbb{R}^n$ is the time-invariant convex feasible region. At each time $t = 1, \ldots, T$ in this ONO framework, the decision maker first chooses an action $\mathbf{x}_t \in \mathbb{S}$ while oblivious to the loss function $f_t : \mathbb{S} \to \mathbb{R}$. Once the action is played, it is evaluated against f_t, which may be chosen by an adversary in response to the action. Then, the decision maker is granted access to the loss function and its gradient.

The performance of a decision maker, or equivalently an algorithm, in online settings is typically evaluated by a metric called *regret* [50]. In this paper, we exclusively focus on the strictest version of regret, *dynamic regret*, which is defined as

$$\mathbf{Reg}_T^d(\mathbf{x}_1, ..., \mathbf{x}_T) := \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t^*, \tag{2.2}$$

where f_t^* denotes the global optimal objective value of (2.1). Dynamic regret (also called *non-stationary regret*) compares the decision maker's actions to an optimal action at each time t. In comparison, *static regret* (also called *stationary regret* or simply *regret*) compares the decision maker's actions to the best fixed action in hindsight:

$$\mathbf{Reg}_T^s(\mathbf{x}_1, ..., \mathbf{x}_T) = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{y}\in\mathbb{S}} \sum_{t=1}^{T} f_t(\mathbf{y}). \tag{2.3}$$

In general, nonconvex optimization problems are NP-hard. Therefore, commonly used local search algorithms, such as first-order and second-order descent algorithms, may converge to a spurious local minimum (i.e., a local minimum that is not globally optimal). As a result, dynamic regret can be arbitrarily high in a general setting due to the inability to efficiently find a near-optimal point $\mathbf{x}_t$. Existing works in online optimization literature have derived regret bounds in terms of various quantities, such as the temporal variation in the loss functions [10, 57, 19], the temporal variation in the gradients of the loss functions [57, 79], and the temporal variation in a decision sequence (also called path length or path variation) [131, 48, 57, 123, 79, 126]. Details on many of these variation measures used to bound dynamic regret can be found in [90], where the online convex optimization problem is analyzed. The existing regret bounds for ONO typically either focus on static regret [122, 83, 51, 127] or require the loss functions to be weakly pseudo-convex which is a restrictive condition that excludes spurious local minima [39]. Dynamic regret for unconstrained ONO is studied in [75], under the assumptions that an initial point near the optimal solution is known, the loss function is strongly convex around the optimal solution, and the decision maker has access to second-order information. In [96], the authors of this paper established probabilistic nonconvexity regret bounds for a variation of the ONO problem in which f_t is known to the decision maker at time t, future loss functions are unknown but fixed, the global minima are "sufficiently superior" to all local minima, and limit points of a continuous-time projected gradient algorithm can be found precisely. Finally, [30] and [35] explored how variability in the input data can help ONO solution trajectories escape non-global local solutions over time, but they did not study dynamic regret and focused on asymptotic regret.

The main goal of this paper is to analyze how the quality of the obtained solutions evolves in ONO settings where the global solution changes slowly over time. To this end, we first develop mathematical tools for characterizing the landscape of constrained nonconvex optimization problems and analyze the behavior of the projected gradient descent algorithm on such problems. There are many conditions in the literature that guarantee linear convergence of local search algorithms. In the unconstrained case when $\mathbb{S} = \mathbb{R}^n$, the Polyak-Łojasiewicz (PL) condition has been proven to be weaker than other common assumptions (such as strong convexity, essential strong convexity, weak convexity, and restricted secant inequality) that guarantee linear convergence [62]. Despite its favorable characteristics, requiring that a function satisfy the PL condition still significantly restricts the type of nonconvex functions that one can study. For instance, functions satisfying the PL condition cannot have local minima which are not globally optimal.

We leverage the generalization of the PL condition for constrained optimization, called the proximal-PL condition (originally proposed in [62]), to study dynamic regret minimization in a nonconvex setting. The first contribution of this paper is to establish a target set for each time instance with the property that once the algorithm finds a point in the corresponding target set at a given time, the global minimizers of future problems can be found efficiently. These time-varying target sets are defined with respect to the proximal-PL condition and the global solution. We show several important properties of these sets, including linear convergence to the global minimizer and quadratic growth.

The design and regret analysis of two online algorithms constitute the second contribution of this paper. Specifically, we equip local search algorithms with memory, random exploration, and multiple gradient queries and establish dynamic regret bounds for each algorithm in terms of the path length and squared path length of the optimal decision sequence, when the difference between consecutive points in this sequence is bounded appropriately. The first algorithm assumes that the decision maker has some prior knowledge about the initial function and can start at a point that is within its target set. This algorithm ensures bounded dynamic regret by producing decisions which track the time-varying target sets. The second algorithm obviates this initial condition assumption by using random exploration. In this case, dynamic regret depends on when the decision maker first finds a point within the corresponding target set, as after that time all decisions will track the time-varying target sets. Therefore, the relative volume of the time-varying target sets with respect to the entire feasible domain–a measure of how favorable the loss function landscape is–influences the likelihood that the dynamic regret will be small. In particular, a single low-complexity problem in the sequence can have a large influence on the outcomes.

The remainder of this paper is organized as follows. In Section 2.2, we analyze the optimization problem for each fixed time step, focusing on a neighborhood of the global solution. In Section 2.3, we introduce ONO algorithms, derive bounds on their dynamic regret, and support the analysis with empirical results. Finally, we conclude the paper in Section 2.4.

2.1.1 Notations

Let $||\cdot||$ indicate the ℓ_2-norm of a vector and $|\cdot|$ represent the cardinality of a set. The symbols $\mathbb{R}^n$ and $\mathbb{Z}_+$ denote the space of n-dimensional real vectors and the set of positive integers, respectively. The globally optimal objective value of the optimization problem at time t is denoted by f_t^*. If there is a unique global optimum at time t, it will be denoted as $\mathbf{x}_t^*$, in which case $f_t(\mathbf{x}_t^*) = f_t^*$. The indicator function $\mathbb{I}_{\mathbb{S}}(\mathbf{x})$ returns zero if $\mathbf{x}$ belongs to the set $\mathbb{S}$ and infinity otherwise. We define the projection operator as follows:

$$\Pi_{\mathbb{S}}(\mathbf{x}) := \underset{\mathbf{y}\in\mathbb{S}}{\operatorname{argmin}} \|\mathbf{x}-\mathbf{y}\|. \tag{2.4}$$

The tangent cone of a convex set $\mathbb{S}$ at $\mathbf{x}$ is denoted as $\mathbb{T}_{\mathbb{S}}(\mathbf{x})$. The sublevel set $\mathcal{L}_t$ is defined as $\mathcal{L}_t(\alpha) := \{\mathbf{x}\in\mathbb{S} | f_t(\mathbf{x}) < \alpha\}$. Finally, $\mathbb{P}[\cdot]$ denotes the probability of the argument.

2.2 Theoretical Results for a Fixed Time Step

2.2.1 Properties of the Problem Structure

Throughout this paper, we make the following assumptions on the problem structure:

1) The time-invariant feasible region $\mathbb{S} \subset \mathbb{R}^n$ is a compact, convex set known to the decision maker.

2) f_t is continuously differentiable on $\mathbb{S}$, but potentially nonconvex in $\mathbf{x}$ with many local minima, for all $t \in \{1, 2, \ldots, T\}$.

   – This assumption ensures that the magnitude of the gradient is bounded above by a positive constant M_1 for all $t \in \{1, 2, \ldots, T\}$. That is, $\sup_{\mathbf{x}\in\mathbb{S}, 1\leq t\leq T} ||\nabla f_t(\mathbf{x})|| \leq M_1$.

3) f_t has a unique global minimum $\mathbf{x}_t^*$ over $\mathbb{S}$ for all $t \in \{1, 2, \ldots, T\}$.

4) The first derivative of f_t is L-Lipschitz continuous on $\mathbb{S}$ for all $t \in \{1, 2, \ldots, T\}$, implying the following inequality for some constant L:

$$f_t(\mathbf{y}) - f_t(\mathbf{x}) \leq \langle \nabla f_t(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2}||\mathbf{y} - \mathbf{x}||^2, \ \forall \mathbf{x}, \mathbf{y} \in \mathbb{S}. \tag{2.5}$$

2.2.2 Proximal Polyak-Łojasiewicz Regions

In the context of unconstrained optimization problems, a differentiable function f_t satisfies the Polyak-Łojasiewicz (PL) condition [97] if the following condition holds for some parameter $\mu > 0$:

$$\underbrace{\frac{1}{2}||\nabla f_t(\mathbf{x})||^2 \geq \mu(f_t(\mathbf{x}) - f_t^*)}_{\text{PL inequality}}, \quad \forall \mathbf{x} \in \mathbb{R}^n. \tag{2.6}$$

If a function satisfies the PL condition and the magnitude of its gradient is small at some $\mathbf{x}$, then the function value at $\mathbf{x}$ will be close to the global minimum. This is the reason why the PL condition is also referred to as the gradient domination condition [36]. For a general (unconstrained) nonconvex optimization problem, first-order methods such as gradient descent may not converge to a global minimizer. However, if a function f_t satisfies the Polyak-Łojasiewicz condition, then every stationary point is a global minimizer. Moreover, PL is one of the most general conditions under which gradient descent offers linear convergence to a global minimizer [62]. Note that, in general, functions satisfying the PL condition may not have a unique global minimizer.

The top plot in Figure 2.1 shows an example of a nonconvex function that satisfies the PL condition. On the other hand, the function in the bottom plot of Figure 2.1 manifests spurious local minima and therefore cannot satisfy the PL inequality for all $\mathbf{x}$ for any $\mu > 0$. However, for a given μ, we can identify a subset of $\mathbb{R}$ that satisfies the PL inequality. The idea of focusing on regions where the PL inequality is satisfied, rather than only considering

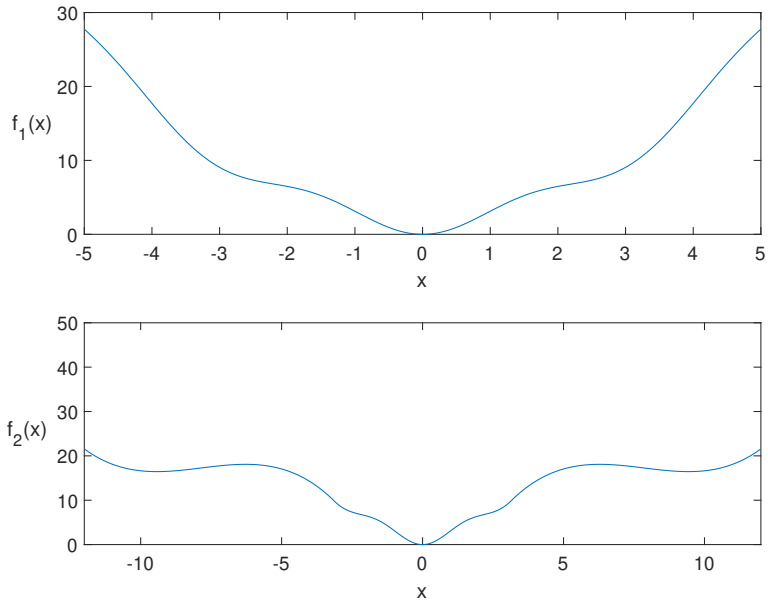

Figure 2.1: The top figure shows the nonconvex function $f_1(x) = x^2 + 3\sin^2(x)$, which satisfies the PL inequality with the parameter $\mu = 1/32$ for all $x \in \mathbb{R}$ [62]. The bottom figure shows an example of a nonconvex function that satisfies the PL inequality with the parameter $\mu = 1/32$ only for $x \in \{[-55.9, -10.2] \cup [-5.4, 5.4] \cup [10.2, 55.9]\}$. The function for the bottom figure is given below:

$$f_2(x) = \begin{cases} -\frac{1}{3\pi}x^3 - \frac{5}{2}x^2 - 6\pi x - \frac{17}{6}\pi^2, & \text{if } x < -\pi \\ x^2 + 3\sin^2(x), & \text{if } -\pi \leq x \leq \pi \\ \frac{1}{3\pi}x^3 - \frac{5}{2}x^2 + 6\pi x - \frac{17}{6}\pi^2, & \text{if } \pi < x \end{cases}$$

functions which satisfy the PL inequality over $\mathbb{R}^n$, leads to our definition of time-varying target sets in Section 2.2.3.

Next, we return to considering constrained optimization problems. Constrained optimization can be cast in the framework of unconstrained optimization by appending the objective function with $\mathbb{I}_{\mathbb{S}}$, an indicator function of a convex set $\mathbb{S}$. This indicator function is non-smooth and convex. Subsequently, a natural generalization of gradient descent to the constrained case is the proximal gradient method, whose iteration is described by

$$\mathbf{x}_t^{k+1} = \underset{\mathbf{y}}{\operatorname{argmin}} \left[\langle \nabla f_t(\mathbf{x}_t^k), \mathbf{y} - \mathbf{x}_t^k \rangle + \frac{\|\mathbf{y} - \mathbf{x}_t^k\|^2}{2s} + \mathbb{I}_{\mathbb{S}}(\mathbf{y}) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^k) \right] \tag{2.7}$$

for every $k \in \mathbb{Z}_+$, where s is a positive constant. It can be shown that the above algorithm

is equivalent to the projected gradient descent algorithm:

$$\mathbf{x}_t^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}_t^k - s\nabla f_t(\mathbf{x}_t^k)). \tag{2.8}$$

A matching generalization of the PL inequality, namely the proximal-PL inequality, was first proposed in [62].

Definition 2.1. *(Proximal-PL inequality) For a function f_t, define the proximal-gradient norm with parameter $\beta > 0$ as*

$$\mathcal{D}_t(\mathbf{x}, \beta) = -2\beta \min_{\mathbf{y}} \left[\langle \nabla f_t(\mathbf{x}), \mathbf{y} - \mathbf{x}\rangle + \frac{\beta\|\mathbf{y} - \mathbf{x}\|^2}{2} + \mathbb{I}_{\mathbb{S}}(\mathbf{y}) - \mathbb{I}_{\mathbb{S}}(\mathbf{x})\right]$$

We say that a point $\mathbf{x} \in \mathbb{S}$ satisfies the proximal-*PL inequality with the parameters $\mu > 0$ and $\beta > 0$ if*

$$\frac{1}{2}\mathcal{D}_t(\mathbf{x}, \beta) \geq \mu(f_t(\mathbf{x}) - f_t^*). \tag{2.9}$$

Note that by virtue of the equivalence between equations (2.7) and (2.8), the proximal-gradient norm can also be expressed as follows:

$$\mathcal{D}_t(\mathbf{x}, \beta) = -2\beta\left[\langle \nabla f_t(\mathbf{x}), \Pi_{\mathbb{S}}(\mathbf{x} - \frac{1}{\beta}\nabla f_t(\mathbf{x})) - \mathbf{x}\rangle + \frac{\beta}{2}\|\Pi_{\mathbb{S}}(\mathbf{x} - \frac{1}{\beta}\nabla f_t(\mathbf{x})) - \mathbf{x}\|^2\right]. \tag{2.10}$$

While [62] considers functions that satisfy the proximal-PL inequality at all points in $\mathbb{S}$, in this work we instead identify a subset of the entire space that satisfies the inequality. Hereby, we define the time-varying *proximal-PL region*, denoted $\mathcal{P}_t(\mu, \beta)$, as the set of all $\mathbf{x} \in \mathbb{S}$ satisfying the proximal-PL inequality with the parameters $\mu > 0$ and $\beta > 0$. That is,

$$\mathcal{P}_t(\mu, \beta) := \left\{\mathbf{x} \in \mathbb{S} \;\middle|\; \frac{1}{2}\mathcal{D}_t(\mathbf{x}, \beta) \geq \mu(f_t(\mathbf{x}) - f_t^*)\right\}. \tag{2.11}$$

In the remainder of the paper, we assume that the parameters μ and β are chosen to guarantee existence of an open set containing $\mathbf{x}_t^*$ for each $t \in \{1, \ldots, T\}$ such that the intersection of this set with $\mathbb{S}$ lies in the proximal-PL region. That is, there exists an open set $\mathcal{S}_t^P(\mu, \beta)$ such that $\mathbf{x}_t^* \in (\mathcal{S}_t^P(\mu, \beta) \cap \mathbb{S}) \subset \mathcal{P}_t(\mu, \beta)$.

2.2.3 Time-Varying Regions of Attraction and Target Sets

A proximal-PL region can span over multiple regions of attraction associated with different local minima. Also, note that a region of attraction (RoA) is algorithm dependent. In this paper, we define RoAs (with respect to the global minimizer) that respect the proximal-PL inequality under the projected gradient descent method and also under the projected gradient flow system, a continuous version of the former.

Definition 2.2. *Let $\mathcal{RP}_t^D$ and $\mathcal{RP}_t^C$ denote the subsets of the discrete and continuous RoAs that are contained within the proximal-PL region at time t:*

$$\begin{aligned}\mathcal{RP}_t^D(\mu,\beta,s) := \big\{\mathbf{x} \mid \mathbf{x}_t^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}_t^k - s\nabla f_t(\mathbf{x}_t^k)),\\ \mathbf{x}_t^0 = \mathbf{x}, \lim_{k\to\infty} \mathbf{x}_t^k = \mathbf{x}_t^* \text{ and } \{\mathbf{x}_t^k\}_{k=0}^{\infty} \subset \mathcal{P}_t(\mu,\beta)\big\} \qquad (2.12)\\ \mathcal{RP}_t^C(\mu,\beta) := \big\{\mathbf{x} \mid \dot{\mathbf{x}}_t = \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}_t)}(-\nabla f_t(\mathbf{x}_t)), \mathbf{x}_t(0) = \mathbf{x},\\ \lim_{\ell\to\infty} \mathbf{x}_t(\ell) = \mathbf{x}_t^* \text{ and } \mathbf{x}_t(\ell) \in \mathcal{P}_t(\mu,\beta)\ \forall \ell \geq 0\big\} \qquad (2.13)\end{aligned}$$

One can view the continuous-time RoA as the limit of the discrete-time RoA as we take the step size towards zero. The continuous-time RoA does not depend on any step-size, and therefore is directly related to the properties of the function. The discrete-time RoA is a function of the step size and therefore is algorithm-dependent.

Next, we define a region that we call the target set. In subsequent sections, we will show that if our proposed algorithm enters the target set at any time t, then it is possible to approach the global minimizer and track it at all subsequent times.

Definition 2.3. *(Target set) We define our target set for time t to be a subset of a sublevel set around the global minimizer that belongs to both $\mathcal{RP}_t^D(\mu,\beta,s)$ and $\mathcal{RP}_t^C(\mu,\beta,s)$:*

$$\mathcal{T}_t(\mu,\beta,s) := \mathcal{L}_t(\alpha_t) \cap \mathcal{RP}_t^D(\mu,\beta,s) \qquad (2.14)$$

where α_t is the largest α satisfying the following condition:

$$\mathcal{L}_t(\alpha) \cap \mathcal{RP}_t^D(\mu,\beta,s) \subseteq \mathcal{RP}_t^C(\mu,\beta) \qquad (2.15)$$

All points in each target set are feasible, satisfy the proximal-PL inequality, and lead to the global solution under the continuous and discrete projected gradient descent methods initialized at those points. Theorem 2.1 will show that these target sets are invariant, and Lemma 2.1 will show that their sizes are nonnegligible. One useful way to measure the size of a target set is with respect to the global solution.

Definition 2.4. *(Reach) Define the reach of a target set as the maximum distance between the global minimum and any point in the target set:*

$$\rho_t(\mu,\beta,s) := \max_{\mathbf{x}\in\mathcal{T}_t(\mu,\beta,s)} \|\mathbf{x}_t^* - \mathbf{x}\|. \qquad (2.16)$$

2.2.4 Properties of Target Sets

In [62], the authors showed the linear convergence of the proximal-gradient algorithm when applied to functions satisfying the proximal-PL condition. In this paper, we show that initializing the proximal-gradient algorithm in the corresponding target set ensures linear convergence to the global minimum, regardless of whether the proximal-PL inequality is satisfied for all feasible points. Additionally, there is an open ball around the global solution whose intersection with the feasible set $\mathbb{S}$ is also contained in the corresponding target set.

Theorem 2.1. *Given $\mu > 0, \beta \geq L$ and a fixed instance of t, consider the problem of minimizing f_t over $\mathbb{S}$ (Problem (2.1)) via the projected gradient descent method (2.8) with the step size s. If $\mathbf{x}_t^0 \in \mathcal{T}_t(\mu, \beta, s)$, then the projected gradient descent method with $0 < s < \min(\frac{1}{\mu}, \frac{1}{\beta})$ converges linearly to the optimal value f_t^*, i.e.,*

$$f_t(\mathbf{x}_t^N) - f_t^* \leq (1 - \mu s)^N \left[f_t(\mathbf{x}_t^0) - f_t^*\right], \tag{2.17}$$

and $\mathbf{x}_t^N \in \mathcal{T}_t(\mu, \beta, s)$, where $N \in \{0, 1, 2, \dots\}$ indicates the number of iterations.

Proof. The proof is similar to that of Theorem 5 in [62]. Let $F_t(\mathbf{x}) := f_t(\mathbf{x}) + \mathbb{I}_{\mathbb{S}}(\mathbf{x})$. By using the Lipschitz continuity of the gradient of f_t, one can write:

$$\begin{aligned} F_t(\mathbf{x}_t^1) &= f_t(\mathbf{x}_t^1) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) \\ &\leq f_t(\mathbf{x}_t^0) + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) + \langle \nabla f_t(\mathbf{x}_t^0), \mathbf{x}_t^1 - \mathbf{x}_t^0 \rangle + \frac{L}{2}||\mathbf{x}_t^1 - \mathbf{x}_t^0||^2 + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) \end{aligned}$$

Then, noting that $\mathbf{x}_t^0 \in \mathcal{T}_t(\mu, \beta, s) \subset \mathbb{S}$ and $L \leq 1/s$, we obtain an upper bound of the form:

$$\begin{aligned} F_t(\mathbf{x}_t^1) &\leq f_t(\mathbf{x}_t^0) + \langle \nabla f_t(\mathbf{x}_t^0), \mathbf{x}_t^1 - \mathbf{x}_t^0 \rangle + \frac{1}{2s}||\mathbf{x}_t^1 - \mathbf{x}_t^0||^2 + \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^1) - \mathbb{I}_{\mathbb{S}}(\mathbf{x}_t^0) \\ &= f_t(\mathbf{x}_t^0) - \frac{s}{2}\mathcal{D}_t(\mathbf{x}_t^0, 1/s) \end{aligned}$$

where the equality follows from the definition of $\mathbf{x}_t^{k+1}$ and the proximal-gradient norm. Finally, we upper bound the equation above by using the facts that $\mathbf{x}_t^0$ satisfies the *proximal*-PL inequality with parameters μ and β and that $\mathcal{D}_t(\mathbf{x}_t^0, 1/s) \geq \mathcal{D}_t(\mathbf{x}_t^0, \beta)$ since $\frac{1}{s} \geq \beta$ [62]:

$$F_t(\mathbf{x}_t^1) \leq f_t(\mathbf{x}_t^0) - \mu s[f_t(\mathbf{x}_t^0) - f_t^*].$$

Since $\mathbf{x}_t^1$ is feasible by the definition of projection, we have

$$f_t(\mathbf{x}_t^1) \leq f_t(\mathbf{x}_t^0) - \mu s[f_t(\mathbf{x}_t^0) - f_t^*],$$

which subsequently implies

$$f_t(\mathbf{x}_t^1) - f_t^* \leq (1 - \mu s)\left[f_t(\mathbf{x}_t^0) - f_t^*\right]. \tag{2.18}$$

Repeating the process for N steps, we have the final result:

$$f_t(\mathbf{x}_t^N) - f_t^* \leq (1 - \mu s)^N \left[f_t(\mathbf{x}_t^0) - f_t^*\right].$$

Further, by showing a non-increasing objective value in (2.18), it holds that $\mathbf{x}_t^1 \in \mathcal{L}_t(\alpha_t)$. The definition of $\mathcal{RP}_t^D(\mu, \beta, s)$ in (2.12), paired with knowledge that $\mathbf{x}_t^0 \in \mathcal{RP}_t^D(\mu, \beta, s)$, guarantees that $\mathbf{x}_t^1 \in \mathcal{RP}_t^D(\mu, \beta, s)$. Therefore, $\mathbf{x}_t^1 \in \mathcal{L}_t(\alpha_t) \cap \mathcal{RP}_t^D(\mu, \beta, s) = \mathcal{T}_t(\mu, \beta, s)$,

proving that the target set is invariant under the projected gradient descent method. □
Theorem 2.1 also gives a lower bound on f_t^* after N iterations:

$$f_t^* \geq \frac{f_t(\mathbf{x}_t^N) - \big(1-\mu s\big)^N f_t(\mathbf{x}_t^0)}{1-\big(1-\mu s\big)^N}, \quad \forall N \in \mathbb{Z}_+. \tag{2.19}$$

For unconstrained problems, a function satisfying the PL condition implies that it also satisfies the quadratic growth condition [62]. Next, we prove a similar relationship between the proximal-PL inequality and quadratic growth.

Theorem 2.2. *(Quadratic growth) The following inequality holds:*

$$\sqrt{f_t(\mathbf{x}) - f_t^*} \geq \sqrt{\frac{\mu}{2}}\|\mathbf{x} - \mathbf{x}_t^*\|, \quad \forall \mathbf{x} \in \mathcal{RP}_t^C(\mu, \beta). \tag{2.20}$$

Proof: See Appendix 2.A.

While the proof of Theorem 2.2 relies on the continuous version of the projected gradient flow algorithm, this paper does not require implementing or solving this continuous dynamical system. The algorithms in Section 2.3 use the discrete-time projected gradient descent algorithm. The next lemma establishes what we will refer to as the *robustness property* of a target set.

Lemma 2.1. *(Robustness of a target set) The target set $\mathcal{T}_t(\mu, \beta, s)$ includes a feasible ball of radius at least r around the global solution for some $r > 0$. That is, $\exists \mu, \beta, s, r > 0$: $\mathcal{T}_t(\mu, \beta, s) \supseteq (\mathcal{B}(\mathbf{x}_t^*, r) \cap \mathbb{S})$ for all $t \in \{1, \ldots, T\}$, where $\mathcal{B}(\mathbf{x}_t^*, r) := \{\mathbf{y} \in \mathbb{R}^n \,|\, \|\mathbf{x}_t^* - \mathbf{y}\|^2 \leq r\}$.*

Proof: See Appendix 2.B and Appendix 2.C.

Assumption 2.1. *There exists a constant M_2 such that*

$$f_t(\mathbf{x}) - f_t^* \leq M_2\|\mathbf{x} - \mathbf{x}_t^*\|^2, \ \forall \mathbf{x} \in \mathcal{T}_t(\mu, \beta, s), \ t = 1, \ldots, T. \tag{2.21}$$

Note that if $\mathbf{x}_t^*$ is in the interior of $\mathbb{S}$ at every $t = 1, \ldots, T$, then this assumption is automatically satisfied due to the assumption that ∇f_t is L-Lipschitz. Specifically, it holds with $M_2 = L/2$, which can be derived by substituting $\mathbf{x} = \mathbf{x}^*$ and $\nabla f_t(\mathbf{x}_t^*) = 0$ into (2.5). In the general case, this assumption is similar to f_t being 2-order calm at $\mathbf{x}_t^*$ relative to the set $\mathcal{T}_t(\mu, \beta, s)$ [32, 70, 72].

2.2.5 Visualization of a Proximal-PL Region and Target Set

To develop intuition about proximal-PL regions and target sets, it is beneficial to visualize these sets in an example. Consider the optimization problem

$$\begin{aligned} \min \ & f(x_1, x_2) = x_1^4 - 4x_1^3 + x_1^2 + 2x_1 + \frac{3}{2}\sin(2\pi x_1) \\ & \quad + x_2^4 - 4x_2^3 + x_2^2 + 2x_2 + \frac{3}{2}\sin(2\pi x_2) + 28.87 \\ \text{s.t.} \quad & -1 \leq x_1 \leq 3, \quad -1 \leq x_2 \leq 3 \end{aligned} \tag{2.22}$$

which is depicted in Figure 2.2a. This problem has the optimal value of 0 at $\mathbf{x}^* = (2.75, 2.75)$ and includes many spurious local solutions.

The proximal-PL region and target set for this problem with the parameters $\mu = 0.5$, $\beta = 250 > L$ and $s = \frac{1}{2\beta}$ are depicted in Figure 2.2b and Figure 2.2c, respectively. The proximal-PL region includes a neighborhood of the global solution, as well as points far from the global solution. However, many points in the feasible set do not satisfy the proximal-PL inequality, in particular those near stationary points. Observe that the target set is a subset of $\mathcal{RP}^D, \mathcal{RP}^C$ and the proximal-PL region. The symmetry in Figure 2.2b and Figure 2.2c is a result of the symmetry in the loss function f.

2.3 Online Projected-Gradient Descent with Random Exploration

In this section, we leverage the results developed in Section 2.2 to study the ONO problem (2.1). We introduce and analyze two algorithms for different scenarios:

1. Scenario 1: An initial point in the target region around the global solution $\mathbf{x}_1^*$ is known.
2. Scenario 2: No information about the loss functions or their minimizers is known in advance.

2.3.1 Scenario 1 - Known Desirable Initial Point

Algorithm 2.1 provides a natural approach to solving the ONO problem (2.1) in the setting where a suitable initial point is known. At each time t, the decision maker performs S_t iterations of projected gradient descent on f_t, with the final iteration becoming the decision maker's action at $t + 1$. The assumption is that the decision maker has enough knowledge about the problem at $t = 1$ to select an initial point in the corresponding target set and that the change in the global optimum between time steps is upper-bounded based on parameters reflecting the functions' landscapes. The latter assumption restricts the adversary's choice of loss function and can be regarded as requiring the global solution sequence to have *steadiness*. This assumption is formalized next.

Assumption 2.2. *(Steadiness of global solution) The change in the global optimum between consecutive time steps is upper-bounded by $\bar{r} < r$, where r is as defined in Lemma 2.1. That is, for $t = 1, \ldots, T-1$,*

$$||\mathbf{x}_{t+1}^* - \mathbf{x}_t^*|| \leq \bar{r}, \tag{2.23}$$

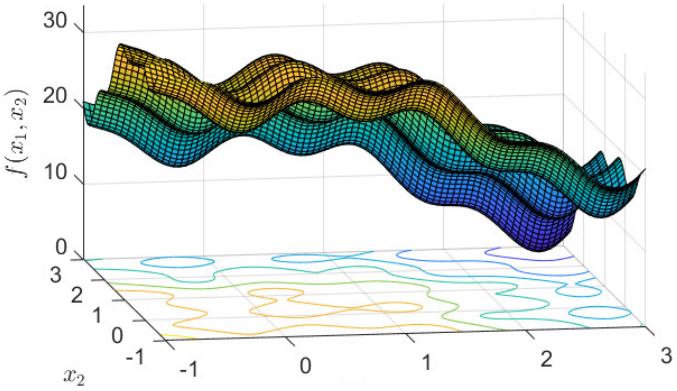

(a) Topology of the objective function f over the feasible set. Observe that this problem has many local minima.

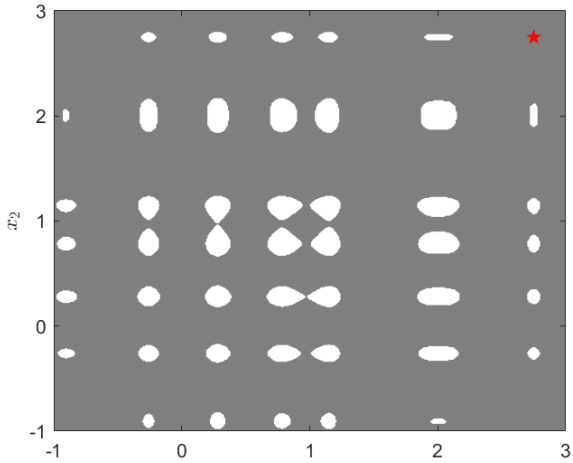

(b) Points in the grey region satisfy the proximal-PL inequality for the function f over the set $\mathbb{S} = [-1, 3] \times [-1, 3]$ with parameters $\mu = 0.5$ and $\beta = 250$, while points in the white regions do not. The unique optimal solution $\mathbf{x}^* = (2.75, 2.75)$ is identified by a red star.

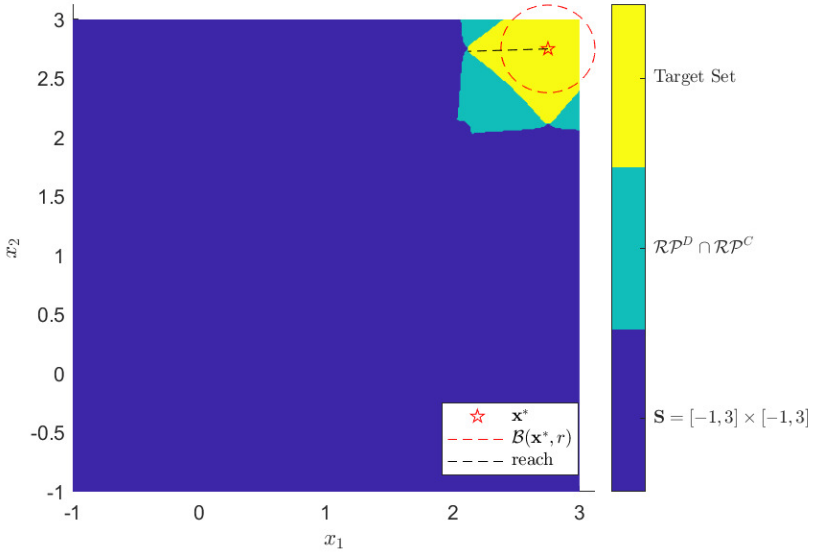

(c) Illustration of the target set and other sets critical to its definition, where points are colored based on the most restrictive set to which they belong. (Recall that Target Set $\subset (\mathcal{RP}^D \cap \mathcal{RP}^C) \subset \mathbb{S}$.) The red dashed circle demonstrates the robustness property established in Lemma 2.1. The length of the black dashed line is the reach of the target set.

Figure 2.2: Visualization of the proximal-PL region and the target set for the optimization problem (2.22).

where μ, β, s, and r collectively satisfy the robustness property in Lemma 2.1 and $\rho_t(\mu, \beta, s)$ is defined in (2.16). Furthermore, assume that S_t is large enough to satisfy the inequalities:

$$\sqrt{\frac{2M_1 \cdot \rho_t(\mu, \beta, s) \cdot (1-\mu s)^{S_t}}{\mu}} \leq r - \bar{r} \tag{2.24a}$$

$$S_t > \frac{\log(\mu) - \log(4M_2)}{\log(1-\mu s)}. \tag{2.24b}$$

Note that this assumption only limits the change in the global minimizer; the overall landscape of the function can change arbitrarily. Under this assumption, we will establish a deterministic dynamic regret bound for Algorithm 2.1. To aid in establishing this bound, we first prove two lemmas:

i. one showing the convergence in terms of the variables $\mathbf{x}$,

ii. another one proving that once the chosen action $\mathbf{x}_t$ is within the target region at time t, all successive actions chosen by the algorithm will also lie within the target region of their respective time.

Lemma 2.2. *Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 2.1. Under Assumptions 2.1 and 2.2, if $\mathbf{x}_t \in \mathcal{T}_t(\mu, \beta, s)$ for any $t \in \{1, 2, \ldots, T-1\}$, then*

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \gamma \|\mathbf{x}_t - \mathbf{x}_t^*\| \tag{2.25}$$

where

$$\gamma = \max_{t=1,\ldots,T} \sqrt{\frac{2M_2(1-\mu s)^{S_t}}{\mu}} < \frac{1}{\sqrt{2}}. \tag{2.26}$$

Proof: From the convergence rate in Theorem 2.1 (specifically equation (2.18)), we have

$$f_t(\mathbf{x}_{t+1}) - f_t^* \leq (1-\mu s)^{S_t} \big[f_t(\mathbf{x}_t) - f_t^* \big]$$

Applying the quadratic growth inequality from Theorem 2.2 and taking the square root of all sides, we obtain

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \sqrt{\frac{f_t(\mathbf{x}_{t+1}) - f_t^*}{\mu/2}} \leq \sqrt{\frac{(1-\mu s)^{S_t} \big(f_t(\mathbf{x}_t) - f_t^* \big)}{\mu/2}}.$$

Then, using the definition of M_2 (2.21), we arrive at

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \underbrace{\sqrt{\frac{2M_2(1-\mu s)^{S_t}}{\mu}}}_{=\gamma_t} \|\mathbf{x}_t - \mathbf{x}_t^*\|. \tag{2.27}$$

Then $\gamma_t < \frac{1}{\sqrt{2}}$ since $S_t > \log(\mu/(4M_2))/\log(1-\mu s)$. □

Algorithm 2.1 Online Projected Gradient Descent with Desirable Initialization

Require: $\mathbf{x}_1 \in \mathcal{T}_1(\mu, \beta, s)$, $0 < s < \min\{\frac{1}{\mu}, \frac{1}{\beta}\}$
1: **for** $t = 1, 2, \ldots, T$ **do**
2: Play $\mathbf{x}_t$
3: Set $\mathbf{z}_0 = \mathbf{x}_t$
4: **for** $i = 1, ..., S_t$ **do**
5: Query $\nabla f_t(\mathbf{z}_{i-1})$
6: Perform projected gradient descent update:
$\mathbf{z}_i = \Pi_{\mathbb{S}}\left[\mathbf{z}_{i-1} - s\nabla f_t(\mathbf{z}_{i-1})\right]$
7: **end for**
8: Set $\mathbf{x}_{t+1} = \mathbf{z}_{S_t}$
9: **end for**

The above lemma proves that given a sufficiently large S_t, we can make γ arbitrarily close to zero, implying that the iterates can become arbitrarily close to the global minimizers at different times. The trade-off is between accuracy and computation time, which is driven by S_t. There is also an intuitive trade-off between the step size s and computation time: smaller step sizes require more algorithmic iterations.

Lemma 2.3. *Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 2.1. Under Assumptions 2.1 and 2.2, if $\mathbf{x}_t \in \mathcal{T}_t(\mu, \beta, s)$ for any $t \in \{1, 2, \ldots, T-1\}$, then $\mathbf{x}_{t+1} \in \mathcal{T}_{t+1}(\mu, \beta, s)$.*

Proof: It is desirable to show that $\|\mathbf{x}_{t+1} - \mathbf{x}_{t+1}^*\| < r$, which ensures that $\mathbf{x}_{t+1} \in \mathcal{B}(\mathbf{x}_{t+1}^*, r)$. By Lemma 2.1 (the robustness property of target sets), we have $\mathcal{B}(\mathbf{x}_{t+1}^*, r) \subset \mathcal{T}_{t+1}(\mu, \beta, s)$. One can write:

$$
\begin{aligned}
\|\mathbf{x}_{t+1} - \mathbf{x}_{t+1}^*\| &\leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + \|\mathbf{x}_t^* - \mathbf{x}_{t+1}^*\| \\
&\leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + \bar{r} \\
&\leq \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| + r - \sqrt{\frac{2M_1\rho_t(\mu,\beta,s)(1-\mu s)^{S_t}}{\mu}} \\
&\leq \sqrt{\frac{2\big(f_t(\mathbf{x}_{t+1}) - f_t^*\big)}{\mu}} + r - \sqrt{\frac{2M_1\rho_t(\mu,\beta,s)(1-\mu s)^{S_t}}{\mu}} \\
&\leq \sqrt{\frac{2(1-\mu s)^{S_t}\big(f_t(\mathbf{x}_t) - f_t^*\big)}{\mu}} + r - \sqrt{\frac{2M_1\rho_t(\mu,\beta,s)(1-\mu s)^{S_t}}{\mu}} \\
&\leq r + \sqrt{\frac{2M_1(1-\mu s)^{S_t}}{\mu}}\left(\sqrt{\|\mathbf{x}_t - \mathbf{x}_t^*\|} - \sqrt{\rho_t(\mu,\beta,s)}\right) \\
&\leq r
\end{aligned}
$$

where the second and third inequalities use Assumption 2.2, the fourth inequality applies Theorem 2.2, the fifth inequality is due to Theorem 2.1, the sixth inequality applies the bounded gradient assumption from Section 2.2.1, and the last inequality is due to (2.16). □

Now, we present a dynamic regret bound for Algorithm 2.1.

Corollary 2.1. *Consider a sequence* $\{\mathbf{x}_t\}_{t=1}^T$ *generated by Algorithm 2.1. Under Assumptions 2.1 and 2.2, the dynamic regret satisfies the following inequality for every constant* $\eta > 0$*:*

$$\boldsymbol{Reg}_T^d(\mathbf{x}_1,...,\mathbf{x}_T) \leq \min\{A(1,T), B(1,T)\}, \tag{2.28}$$

where

$$A(t_1,t_2) := \frac{M_1}{1-\gamma}\sum_{t=t_1+1}^{t_2}\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\| + \frac{M_1\rho_{t_1}(\mu,\beta,s)}{(1-\gamma)} \tag{2.29a}$$

$$B(t_1,t_2) := \frac{L+\eta}{1-2\gamma^2}\sum_{t=t_1+1}^{t_2}\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|^2 + \frac{\sum_{t=t_1}^{t_2}\|\nabla f_t(\mathbf{x}_t^*)\|^2}{2\eta} + \frac{(L+\eta)\rho_{t_1}(\mu,\beta,s)^2}{(2-4\gamma^2)} \tag{2.29b}$$

Proof: The proofs that $A(1,T)$ and $B(1,T)$ each upper-bound the dynamic regret obtained by Algorithm 2.1 will follow the same lines of reasoning as Theorem 1 and Corollary 1 of [90] and Theorem 2 of [126], respectively. In these works, similar results are proved for strongly convex functions. In the nonconvex setting considered in this paper, we will utilize Lemma 2.2 and Lemma 2.3 in our proofs.

Proof of the inequality $\boldsymbol{Reg}_T^d(\mathbf{x}_1,...,\mathbf{x}_T) \leq A(1,T)$*:* By the Intermediate Value Theorem, there exists $\mathbf{y} \in \{\mathbf{z} \mid \mathbf{z} = \omega\mathbf{x}_t + (1-\omega)\mathbf{x}_t^*, 0 \leq \omega \leq 1\}$ such that $f_t(\mathbf{x}) - f_t(\mathbf{x}_t^*) = \nabla f_t(\mathbf{y})^T(\mathbf{x}_t - \mathbf{x}_t^*)$. Therefore, by applying the bounded gradient assumption in Section 2.2.1, we have

$$\mathbf{Reg}_T^d(\mathbf{x}_1,...,\mathbf{x}_T) \leq M_1\sum_{t=1}^{T}\|\mathbf{x}_t - \mathbf{x}_t^*\|. \tag{2.30}$$

Next we establish an upper bound on the summation in (2.30):

$$\begin{aligned}\sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\| &= \|\mathbf{x}_1-\mathbf{x}_1^*\|+\sum_{t=2}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\| \\ &\leq \|\mathbf{x}_1-\mathbf{x}_1^*\|+\sum_{t=2}^{T}\|\mathbf{x}_t-\mathbf{x}_{t-1}^*\|+\sum_{t=2}^{T}\|\mathbf{x}_t^*-\mathbf{x}_{t-1}^*\| \\ &\leq \|\mathbf{x}_1-\mathbf{x}_1^*\|-\gamma\|\mathbf{x}_T-\mathbf{x}_T^*\|+\gamma\sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\|+\sum_{t=2}^{T}\|\mathbf{x}_t^*-\mathbf{x}_{t-1}^*\| \\ \implies \sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\| &\leq \frac{\|\mathbf{x}_1-\mathbf{x}_1^*\|-\gamma\|\mathbf{x}_T-\mathbf{x}_T^*\|}{1-\gamma}+\frac{1}{1-\gamma}\sum_{t=2}^{T}\|\mathbf{x}_t^*-\mathbf{x}_{t-1}^*\| \qquad (2.31)\\ &\leq \frac{\rho_1(\mu,\beta,s)}{1-\gamma}+\frac{1}{1-\gamma}\sum_{t=2}^{T}\|\mathbf{x}_t^*-\mathbf{x}_{t-1}^*\| \qquad (2.32)\end{aligned}$$

The first inequality invokes the triangle inequality. The second inequality applies Lemma 2.2 for each $t=1,\ldots,T$ and re-indexes the summation. This application of Lemma 2.2 is derived by recursively applying Lemma 2.3 to the requirement that $\mathbf{x}_1\in\mathcal{T}_1(\mu,\beta,s)$. We rearrange terms to arrive at (2.31) and then apply the definition of the reach of the target set (2.16) to achieve the final inequality. Combining (2.32) with (2.30) completes the proof.

Proof of the inequality $\boldsymbol{Reg}_T^d(\mathbf{x}_1,\ldots,\mathbf{x}_T)\leq B(1,T)$*:* Beginning with the L-Lipschitz continuity of ∇f_t and basic fact that $(\|\nabla f_t(\mathbf{x}_t^*)\|-\eta\|\mathbf{x}_t-\mathbf{x}_t^*\|)^2\geq 0$, one can write:

$$\begin{aligned}\sum_{t=1}^{T}f_t(\mathbf{x}_t)-f_t^* &\leq \sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t^*)\|\|\mathbf{x}_t-\mathbf{x}_t^*\|+\frac{L}{2}\|\mathbf{x}_t-\mathbf{x}_t^*\|^2 \\ &\leq \frac{1}{2\eta}\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t^*)\|^2+\frac{L+\eta}{2}\sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\|^2. \qquad (2.33)\end{aligned}$$

Following similar steps to those used to bound $\sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\|$ above, we can establish the following bound:

$$\sum_{t=1}^{T}\|\mathbf{x}_t-\mathbf{x}_t^*\|^2\leq\frac{\|\mathbf{x}_1-\mathbf{x}_1^*\|^2}{1-2\gamma^2}+\frac{2\sum_{t=2}^{T}\|\mathbf{x}_t^*-\mathbf{x}_{t-1}^*\|^2}{1-2\gamma^2}. \qquad (2.34)$$

Combining (2.34) with (2.33) and using the definition of the reach of the target set (2.16) completes the proof. □

Observe that the dynamic regret is a function of the temporal variation in the optimal decision, a common measure of variation discussed in the introduction. $A(1,T)$ is a function of the path length, while $B(1,T)$ depends on the squared path length. The (squared) path

length is weighted by a function of γ that is large when γ is close to $1/\sqrt{2}$ and is approximately one when γ is close to 0. Again, this trade-off between the strength of the dynamic regret bound and computation time is driven by S_t. If some function f_t does not have a unique global minimum, as stated in Section 2.2.1, but instead has multiple disconnected global minimizers each satisfying Assumptions 2.1 and 2.2 for their associated target sets, Corollary 2.1 holds under the updated criteria that $\mathbf{x}_1$ is in the union of target sets (each target set corresponding to a global minimizer at $t = 1$).

2.3.2 Scenario 2 - Blind Initialization

The initialization scenario described in Scenario 1 – that a point in the target region is known at the initial time – is difficult to satisfy in practice. The reason is that the decision maker may have no information about how their adversary will design f_1. In this case, it is advantageous to *explore* the landscape of f_t before selecting decision $\mathbf{x}_{t+1}$. Algorithm 2.2 explores by running the projected gradient descent algorithm from multiple initial points, which are sampled uniformly at random from $\mathbb{S}$ and stored in the set $\mathcal{W}_t$.

The goal of exploration is to find a point in a time-varying target set. The decision maker cannot verify when this occurs, however, since they do not have knowledge of the landscape of the function. As a result, Algorithm 2.2 utilizes *memory*, in the form of the set $\mathcal{M}_t$, to make available at time $t+1$ points which may be in the target set at time t. Once a point in a time-varying target set is sampled, memory ensures that the decision maker has at least one initial point (a point in the set $\mathcal{Y}_t$) in the target set of each future time step. Specifically, if $\mathbf{y}_t^k$ is in the target set for time t, the first *while* loop condition paired with the construction of $\mathcal{M}_{t+1}$ guarantees that $\mathbf{z}_*^k$ is used as an initial point at time $t+1$, and the second *while* loop condition ensures that $\mathbf{z}_*^k$ is in the target set for time $t+1$. This tracking guarantee is formalized in the following lemma.

Lemma 2.4. *Consider sequences $\{\mathbf{x}_t\}_{t=1}^T$ and $\{\mathcal{Y}_t\}_{t=1}^T$ generated by Algorithm 2.2. Under Assumptions 2.1 and 2.2, if $\mathbf{z}_0^k \in (\mathcal{T}_t(\mu,\beta,s) \cap \mathcal{Y}_t)$ for any $t \in \{1,2,\dots,T-1\}$, then $\mathbf{z}_*^k \in (\mathcal{T}_{t+1}(\mu,\beta,s) \cap \mathcal{Y}_{t+1})$.*

Proof: The number of iterations I_t^k is at least as large as S_t. Therefore, applying the same logic as the proof of Lemma 2.3, we know that $\mathbf{z}_{I_t^k}^k \in \mathcal{T}_{t+1}(\mu,\beta,s)$. By Theorem 2.1, we have $f_t(\mathbf{z}_0^k) \geq f_t(\mathbf{z}_1^k) \geq \cdots \geq f_t(\mathbf{z}_{S_t}^k) \geq f_t(\mathbf{z}_{I_t^k}^k)$ with $\mathbf{z}_i^k = \mathbf{z}_{i+1}^k$ only if $\mathbf{z}_i^k = \mathbf{x}_t^*$. This implies that $\mathbf{z}_*^k = \mathbf{z}_{I_t^k}^k$. It remains to show that $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$. If $\mathbf{z}_*^k = \mathbf{x}_{t+1}$, then $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$. Otherwise, since $\mathbf{z}_0^k \in \mathcal{T}_t(\mu,\beta,s)$, it holds that $\underline{c}_k \leq f_t^* + \epsilon \leq \underline{c}_K + \epsilon$, which implies $\mathbf{z}_*^k \in \mathcal{M}_{t+1} \subset \mathcal{Y}_{t+1}$. As a result, $\mathbf{z}_*^k \in \mathcal{Y}_{t+1}$, which completes the proof. □

Since Algorithm 2.1 is a deterministic algorithm, the dynamic regret bound established in Corollary 2.1 is deterministic too. Algorithm 2.2 relies on sampling, and therefore its associated regret bound should be probabilistic. In the following culminating theorem, we provide an upper bound on the dynamic regret accrued using Algorithm 2.2 and a lower bound on the probability with which this bound holds.

Algorithm 2.2 Online Projected Gradient Descent with Random Exploration

Require: $\mathbf{x}_1 \in \mathbb{S}$, $\mathcal{M}_1 = \emptyset$, $m_1 = 0$, $0 < s < \min\{\frac{1}{\mu}, \frac{1}{\beta}\}$

for $t = 1, 2, \ldots, T$ **do**
- Play $\mathbf{x}_t$
- Create $\mathcal{W}_t = \{\mathbf{w}_t^1, \ldots, \mathbf{w}_t^q\}$ by uniformly sampling q random points from $\mathbb{S}$
- Set $\mathcal{Y}_t = \mathcal{W}_t \bigcup \mathcal{M}_t \bigcup \{\mathbf{x}_t\} := \{\mathbf{y}_t^1, ..., \mathbf{y}_t^{q+m_t+1}\}$

for $k = 1, 2, \ldots, q + m_t + 1$ **do**
- Initialize $\mathbf{z}_0^k = \mathbf{y}_t^k,\ \mathbf{z}_*^k = \mathbf{y}_t^k,\ \underline{c}_k = \infty,\ \bar{b}_k = -\infty$
- Set $i = 1$

while $\underline{c}_k - \bar{b}_k > \epsilon$ or $i \leq S_t$ **do**
- Query $\nabla f_t(\mathbf{z}_{i-1}^k)$
- Compute $\mathbf{z}_i^k = \Pi_{\mathbb{S}}\left[\mathbf{z}_{i-1}^k - s\nabla f_t(\mathbf{z}_{i-1}^k)\right]$
- Query $c_i^k = f_t(\mathbf{z}_i^k)$

if $c_i^k < \underline{c}^k$ **then**
- $\mathbf{z}_*^k = \mathbf{z}_i^k,\ \underline{c}_k = c_i^k$

end if
- $b_i^k = \Big(f_t(\mathbf{z}_i^k) - (1-\mu s)^i\, f_t(\mathbf{z}_0^k)\Big) / (1-(1-\mu s)^i)$
- Update $\bar{b}_k = \max\{\bar{b}_k, b_i^k\}$
- Update $i = i + 1$

end while
- Set $I_t^k = i$

end for
- Let $K = \operatorname{argmin}_k \underline{c}_k$, and set $\mathbf{x}_{t+1} = \mathbf{z}_*^K$
- Store in memory all other points in $\{\mathbf{z}_*^k\}_{k=1}^{q+m_t+1}$ which could be in the proximal-PL region at time t:
$\mathcal{M}_{t+1} = \{\mathbf{z}_*^k : \underline{c}_k \leq \underline{c}_K + \epsilon, k \in \{1, ..., q+m_t+1\} \setminus K\}$
- $m_{t+1} = |\mathcal{M}_{t+1}|$ (Note: $m_{t+1} \leq qt$)

end for

Theorem 2.3. *Consider a sequence $\{\mathbf{x}_t\}_{t=1}^T$ generated by Algorithm 2.2. Under Assumptions 2.1 and 2.2, the dynamic regret satisfies the following probabilistic bound for all $\bar{T} \in \{1, \ldots, T\}$:*

$$\mathbb{P}\Big[\boldsymbol{Reg}_T^d(\mathbf{x}_1, \ldots, \mathbf{x}_T) \leq \boldsymbol{Reg}_{\bar{T}-1}^d(\mathbf{x}_1, \ldots, \mathbf{x}_{\bar{T}-1}) + \min\{A(\bar{T}, T), B(\bar{T}, T)\}\Big] \tag{2.35}$$
$$\geq 1 - \prod_{t=1}^{\bar{T}} \left(1 - \frac{Vol(\mathcal{T}_t(\mu, \beta, s))}{Vol(\mathbb{S})}\right)^q,$$

where $Vol(\cdot)$ indicates the volume of the set. This theorem relates the dynamic regret at time T to the dynamic regret at an earlier time $\bar{T}$, the variation within the optimal decision sequence after $\bar{T}$, and the relative sizes of the target sets through $\bar{T}$.

Proof: The probability that a point located in the time-varying target set $\mathcal{T}_t(\mu,\beta,s)$ appears in $\mathcal{Y}_t$ by time $\bar{T}$ is related to the volumes of $\mathcal{T}_t(\mu,\beta,s)$ and $\mathbb{S}$ because, at each time step, q initial points are selected from $\mathbb{S}$ uniformly at random. Hence,

$$\begin{aligned}
&\mathbb{P}\left[\mathcal{Y}_t \cap \mathcal{T}_t(\mu,\beta,s) \neq \emptyset \text{ for some } t \in \{1,\dots,\bar{T}\}\right] \qquad (2.36)\\
&\geq \mathbb{P}\left[\mathcal{W}_t \cap \mathcal{T}_t(\mu,\beta,s) \neq \emptyset \text{ for some } t \in \{1,\dots,\bar{T}\}\right]\\
&= \mathbb{P}\left[\bigcup_{t=1}^{\bar{T}} \bigcup_{i=1}^{q} \mathbf{w}_t^i \in \mathcal{T}_t(\mu,\beta,s)\right]\\
&= 1 - \mathbb{P}\left[\mathbf{w}_t^i \notin \mathcal{T}_t(\mu,\beta,s)\ \forall t = 1,\dots,\bar{T}, \forall i = 1,\dots,q\right]\\
&= 1 - \prod_{t=1}^{\bar{T}} \left(1 - \frac{\text{Vol}(\mathcal{T}_t(\mu,\beta,s))}{\text{Vol}(\mathbb{S})}\right)^q \qquad (2.37)
\end{aligned}$$

Now, we will show that if $\mathcal{Y}_t \cap \mathcal{T}_t(\mu,\beta,s) \neq \emptyset$ for some $t \in \{1,\dots,\bar{T}\}$, then the dynamic regret is upper bounded by the expression in (2.35). Applying Lemma 2.4 and Corollary 2.1 yields that

$$\begin{aligned}
\mathbf{Reg}_T^d(\mathbf{x}_1,\dots,\mathbf{x}_T) &= \sum_{t=1}^{\bar{T}-1} (f_t(\mathbf{x}_t) - f_t^*) + \sum_{t=\bar{T}}^{T} (f_t(\mathbf{x}_t) - f_t^*)\\
&= \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1,\dots,\mathbf{x}_{\bar{T}-1}) + \sum_{t=\bar{T}}^{T} (f_t(\mathbf{x}_t) - f_t^*)\\
&\leq \mathbf{Reg}_{\bar{T}-1}^d(\mathbf{x}_1,\dots,\mathbf{x}_{\bar{T}-1}) + \min\{A(\bar{T},T), B(\bar{T},T)\} \qquad (2.38)
\end{aligned}$$

This completes the proof. □

Observe that the strength of this probabilistic bound depends on the landscape of loss functions around the global solution through the volume of the target sets. In particular, one can analyze the role that a "lower-complexity problem" at some time $\bar{T}$ plays in determining the complexity of the entire online nonconvex optimization. As an extreme but important case, suppose that there is a time $\bar{T} \in \{1,\dots,T\}$ such that $f_{\bar{T}}$ is convex. Then the dynamic regret bound (2.35) holds with probability 1 since $\mathcal{T}_t(\mu,\beta,s) = \mathbb{S}$. In other words, the existence of a single convex problem, in between the sequence of nonconvex problems, is enough to break down the NP-hardness of solving nonconvex problems for all future times, under the steadiness of the global solution assumption. On the other hand, if the global solution is extremely "sharp" at all times, it is unrealistic to expect any algorithm with limited computation time to find the global solution. Thus, dynamic regret could be arbitrarily large in this case. Indeed, the target set of a sharp minima is small and therefore the probability of satisfying the dynamic regret bound in (2.35) is low, as expected. If some function f_t does not have a unique global minimum, as stated in Section 2.2.1, but instead has multiple disconnected global minimizers each satisfying Assumptions 2.1 and 2.2 for their associated

target sets, Theorem 2.3 holds with the union of target sets (each target set corresponding to a global minimizer at t) replacing the single target set $\mathcal{T}_t(\mu, \beta, s)$.

Choices for the step size s, number of iterations S_t, and number of samples q, represent trade-offs between regret bound strength and computation time. As discussed in Section 2.3.1, a smaller step size requires more algorithmic iterations to satisfy Assumption 2.2. Increasing the number of iterations may increase the time to execute the while loop (line 8). However, larger values of S_t improve the upper bound on dynamic regret in (2.35) by reducing γ. Increasing the number of random initial points improves the probability with which (2.35) holds, but also increases computation time of in the inner *for* loop, which is executed up to $qt + 1$ times.

2.3.3 Empirical Study of Algorithm 2.2

The objective of this section is to support the results of Section 2.3.2 through numerical analysis. We will illustrate the performance of Algorithm 2.2 on online function sequences which satisfy the assumptions in Section 2.2.1 and Assumptions 2.1 and 2.2. To demonstrate the role that a single comparatively low-complexity problem can play in a sequence of nonconvex problems, we will consider two cases:

A) *"No low-complexity problem":* In this case, $\{f_t : \mathbb{R}^2 \to \mathbb{R}\}_{t=1}^{40}$ each have many local minima over $\mathbb{S} = [-1, 3] \times [-1, 3]$ and the target sets' volumes represent between 2.47% and 4.14% of the feasible space. The geometry of $f_1, \ldots, f_6$, which are representative of the entire sequence, is shown in Figure 2.3.

B) *"Lower-complexity problem at time 4":* In this case, $\{\bar{f}_t : \mathbb{R}^2 \to \mathbb{R}\}_{t=1}^{40}$ is identical to Case A at every time period except $t = 4$. The target set corresponding to $\bar{f}_4$ covers 20.7% of the feasible space. Meanwhile, $\mathbf{x}_4^*$ is the same in both scenarios.

The parameter choices and key problem constants for these two online optimization problems are summarized in Table 2.1.

$\mathbb{S}$	μ	β	s	ϵ	r	$S_t(max)$	L	M_1	$\|\mathbf{x}_t^* - \mathbf{x}_{t-1}^*\|$
$[-1, 3]^2$	0.5	289	0.0031	0.1	0.29	7060	289	140	0.22

Table 2.1: Parameter and constant values for the empirical study of Algorithm 2.2

We conducted 500 trials of Algorithm 2.2 on Case A and Case B for 3 different sampling rates: $q = 1$, $q = 2$, and $q = 5$. Figure 2.4 plots the empirical probability that $\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset$ versus the theoretical lower bound provided in Theorem 2.3. (Note that, by Lemma 2.4, this is the same as the probability that $\mathcal{Y}_t \cap \mathcal{T}_t(\mu, \beta, s) \neq \emptyset$ for some $\bar{t} \in \{1, \ldots, t\}$.) For the same value of q, the two cases are identical for $t = 1, 2, 3$ and diverge at $t = 4$ as a result of the "easy" problem in Case B. These results support Theorem 2.3. A gap between the lower bound and observed likelihood of initializing in the target region

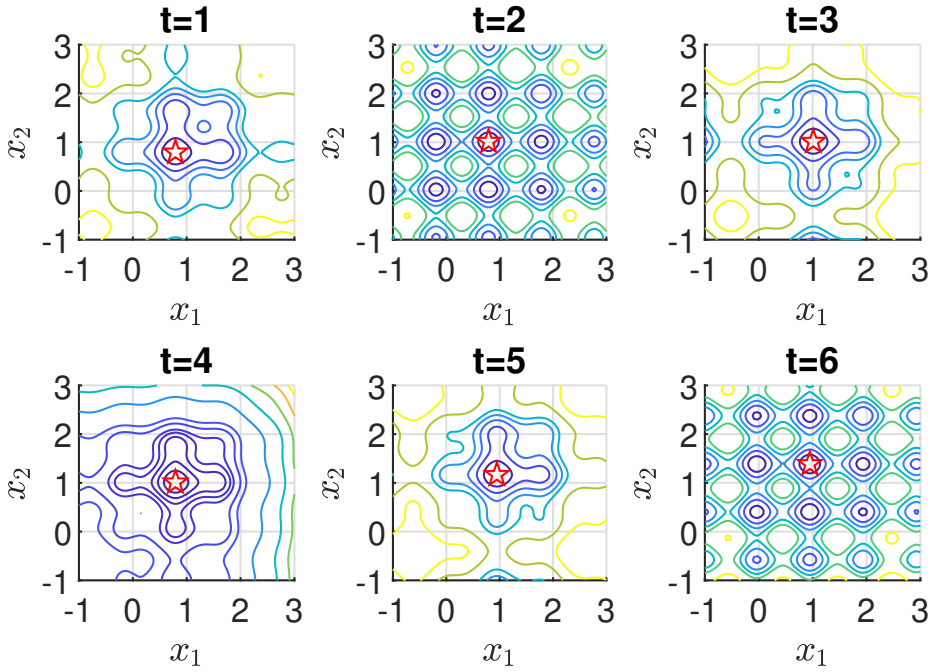

Figure 2.3: Contour plots of $f_1, \ldots, f_6$ for Case A. The red star marks the unique global minimum of each function.

is expected, since the lower bound does not account for the possibility that $\mathbf{x}_t$ or a memory point may be in the subsequent target set. The dynamic regret and optimality gap over time is shown in Figure 2.5. Regret accumulates quickly until the target set is found and then accumulates slowly as Algorithm 2.2 starts tracking the global solution.

2.4 Conclusion

In this paper, we defined proximal-PL regions and target sets, characterized their properties, and used this new knowledge to propose and analyze algorithms for online nonconvex optimization problems. Linear convergence to the global minimizer and quadratic growth are the two key properties of the target sets that we established. Since dynamic regret can be arbitrarily large when there are no restrictions on the loss functions, we constrain consecutive functions to have global solutions which are not too far apart, but do not limit the variation in the loss functions otherwise. In this setting, we propose two online algorithms. Algorithm 2.1 is relevant when the decision maker has a good initial point, and it provides a deterministic dynamic regret upper bound as a function of the temporal variation in the optimal decision sequence. Algorithm 2.2 utilizes exploration and memory to be relevant regardless of the initial point. It provides a probabilistic dynamic regret upper bound, which is also a function of the temporal variation in the optimal decision sequence. The strength

of this probabilistic bound depends on the loss function landscapes. For example, the bound holds with probability 1 in the special case where one of the loss functions in the sequence is convex. Empirical studies support these bounds.

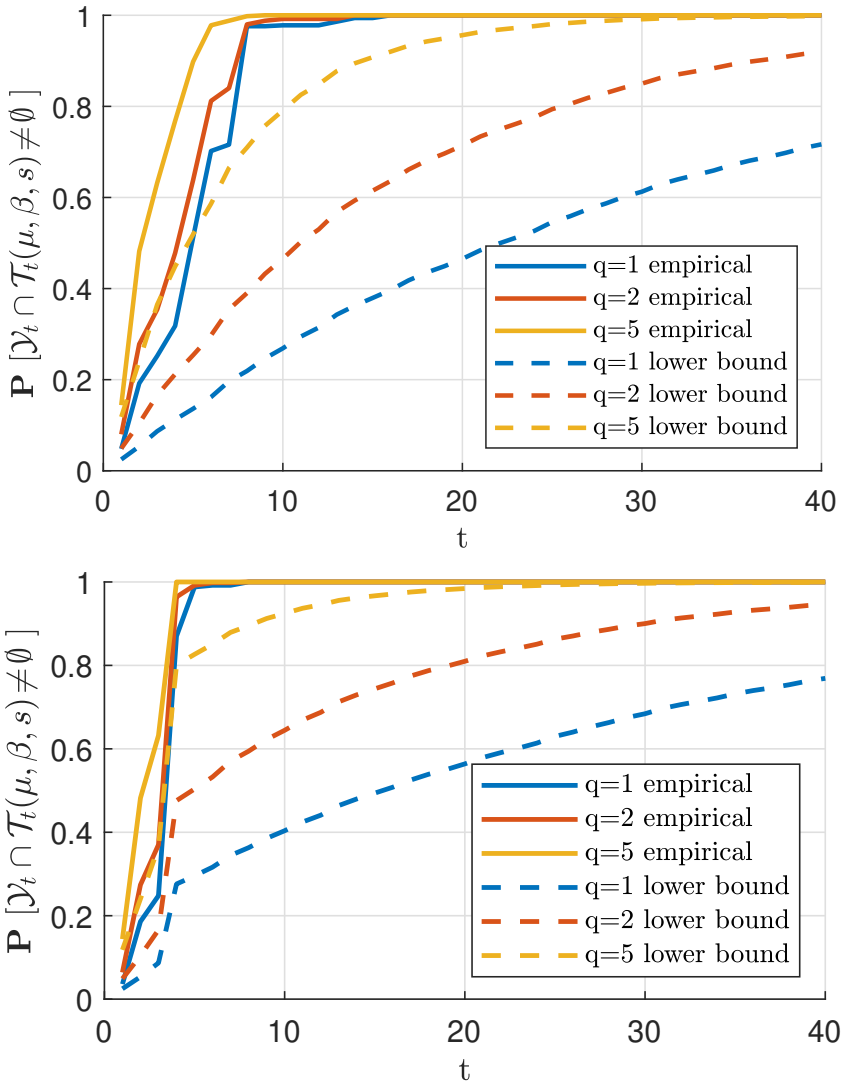

(a) Case A

(b) Case B

Figure 2.4: Empirical validation of Theorem 2.3 probability bound

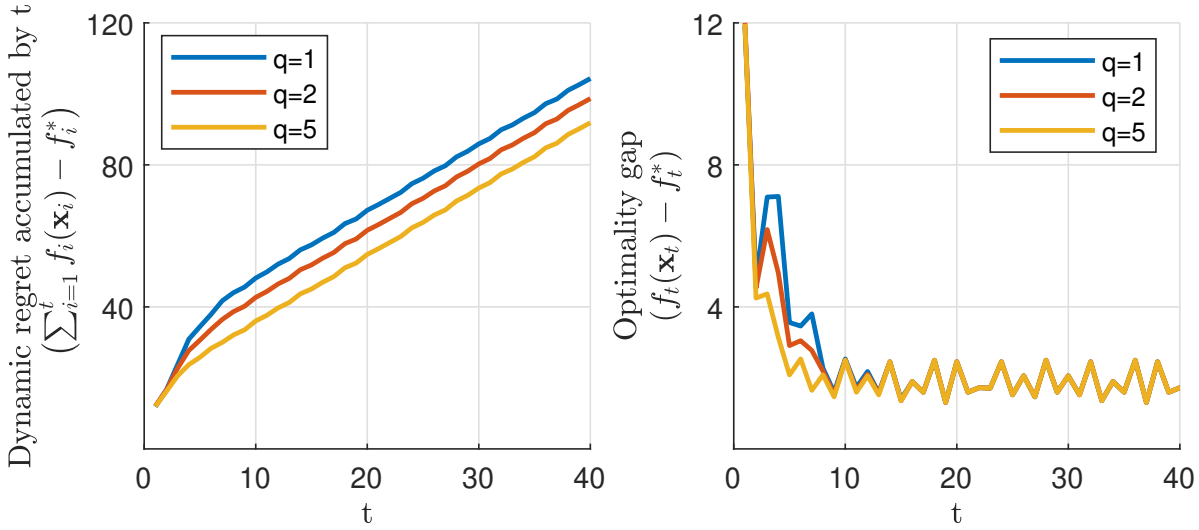

(a) Case A

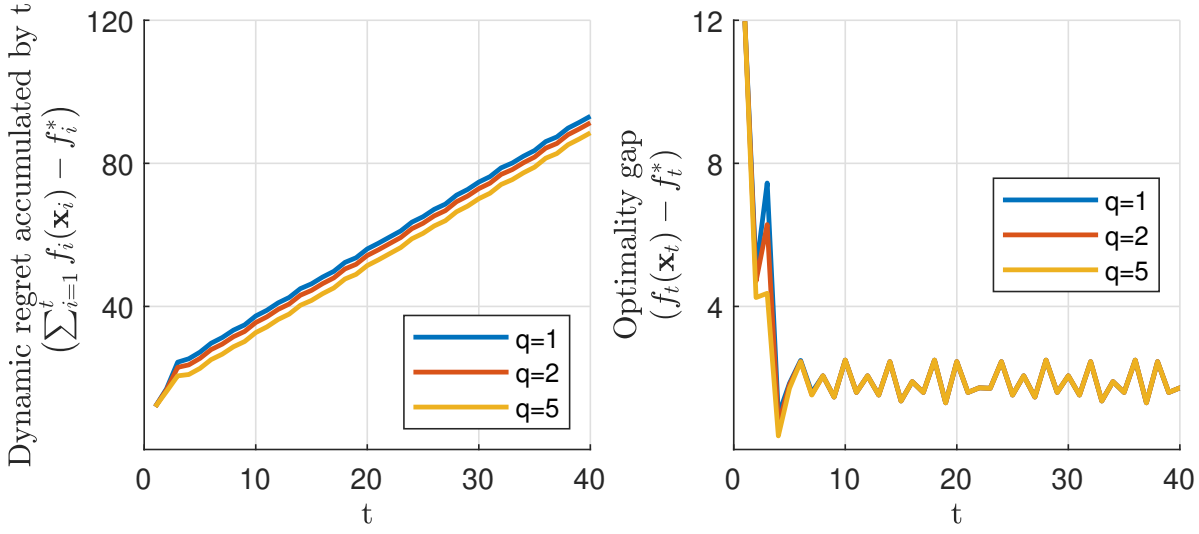

(b) Case B

Figure 2.5: Empirical regret resulting from Algorithm 2.2

Appendix

2.A Proof of Theorem 2.2

Take the function f to be any $f_t,\ t \in \{1,\dots,T\}$. Define the function $g(\mathbf{x}) := \sqrt{f(\mathbf{x}) - f^*}$ and

$$\dot{\mathbf{x}}(\ell) = \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\left(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\right), \quad \forall \ell \geq 0 \tag{2.39}$$

with $\mathbf{x}(0) = \mathbf{x}$. Then, by the fundamental theorem of calculus, we have

$$\sqrt{f(\mathbf{x}) - f^*} = g(\mathbf{x}) - g(\mathbf{x}^*) = -\int_0^\infty \frac{\mathrm{d}}{\mathrm{d}\ell} g(\mathbf{x}(\ell))\,\mathrm{d}\ell = -\int_0^\infty \frac{\nabla f(\mathbf{x}(\ell))}{2g(\mathbf{x}(\ell))}\cdot\dot{\mathbf{x}}(\ell)\,\mathrm{d}\ell \tag{2.40}$$

The following lemma will be used to establish a lower bound on the term inside the integral.

Lemma 2.5. *Consider the projected gradient flow (2.39) with* $\mathbf{x}(0) \in \mathbb{S}$. *There exists a unique solution* $\mathbf{x}(\ell)$ *to this projected dynamical system, and*

$$\left\langle \nabla f(\mathbf{x}(\ell)), \dot{\mathbf{x}}(\ell)\right\rangle = \frac{-1}{2\beta}\cdot\lim_{h\to\infty}\mathcal{D}(\mathbf{x}(\ell),h) - \frac{\beta\left\|\dot{\mathbf{x}}(\ell)\right\|^2}{2}. \tag{2.41}$$

Proof: The existence and uniqueness of the solution of the projected dynamical system (2.39) is guaranteed by [92, Thm. 2.5] under the assumptions in Section 2.2.1. Let $\{\mathbf{x}(\ell)\}_{\ell\geq 0}$ denote the unique solution to (2.39), and

$$\begin{aligned}\overline{\mathbf{x}}_\epsilon(\ell) &:= \Pi_{\mathbb{S}}\left(\mathbf{x}(\ell) - \frac{\epsilon}{\beta}\nabla f(\mathbf{x}(\ell))\right)\\ &= \operatorname*{argmin}_{\mathbf{y}\in\mathbb{S}}\left[\langle \mathbf{y}-\mathbf{x}(\ell), \nabla f(\mathbf{x}(\ell))\rangle\epsilon + \frac{\beta\|\mathbf{y}-\mathbf{x}(\ell)\|^2}{2}\right]\end{aligned}$$

Then, it follows from [53, Sec III Prop. 5.3.5] that

$$\lim_{\epsilon\downarrow 0}\frac{\overline{\mathbf{x}}_\epsilon(\ell) - \mathbf{x}(\ell)}{\epsilon} = \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\left(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\right),$$

where $\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))$ is the tangent cone of $\mathbb{S}$ at $\mathbf{x}(\ell) \in \mathbb{S}$. By the definition of the proximal-gradient,

$$\begin{aligned}
&\lim_{h\to\infty} \mathcal{D}(\mathbf{x}(\ell), h) = \lim_{\epsilon\downarrow 0} \mathcal{D}(\mathbf{x}(\ell), \beta/\epsilon) \\
&= \lim_{\epsilon\downarrow 0} \frac{-2\beta}{\epsilon} \cdot \min_{\mathbf{y}\in\mathbb{S}} \left[\langle \nabla f(\mathbf{x}(\ell)), \mathbf{y}-\mathbf{x}(\ell)\rangle + \frac{\beta}{2\epsilon}\|\mathbf{y}-\mathbf{x}(\ell)\|^2\right] \\
&= \lim_{\epsilon\downarrow 0} \frac{-2\beta}{\epsilon^2} \cdot \min_{\mathbf{y}\in\mathbb{S}} \left[\langle \nabla f(\mathbf{x}(\ell)), \mathbf{y}-\mathbf{x}(\ell)\rangle \epsilon + \frac{\beta}{2}\|\mathbf{y}-\mathbf{x}(\ell)\|^2\right] \\
&= \lim_{\epsilon\downarrow 0} \frac{-2\beta}{\epsilon^2} \left[\langle \nabla f(\mathbf{x}(\ell)), \overline{\mathbf{x}}_\epsilon(\ell)-\mathbf{x}(\ell)\rangle \epsilon + \frac{\beta}{2}\|\overline{\mathbf{x}}_\epsilon(\ell)-\mathbf{x}(\ell)\|^2\right] \\
&= -2\beta \left[\langle \nabla f(\mathbf{x}(\ell)), \dot{\mathbf{x}}(\ell)\rangle + \frac{\beta}{2}\|\dot{\mathbf{x}}(\ell)\|^2\right],
\end{aligned}$$

where the last equation is due to the continuity of $\|\cdot\|^2$. Rearranging the above equation yields the desired result. □

Returning to the proof of Theorem 2.2, next we establish a lower bound on the term inside the integral in (2.40).

$$\begin{aligned}
-\frac{\nabla f(\mathbf{x}(\ell))}{2g(\mathbf{x}(\ell))} \cdot \dot{\mathbf{x}}(\ell) &= -\frac{1}{2g(\mathbf{x}(\ell))} \left\langle \nabla f(\mathbf{x}(\ell)), \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big) \right\rangle \\
&= \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{1}{2\beta} \lim_{h\to\infty} \mathcal{D}(\mathbf{x}(\ell), h) + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big) \right\|^2 \right) \\
&\geq \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{\mathcal{D}(\mathbf{x}(\ell), \beta)}{2\beta} + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big) \right\|^2 \right) \\
&\geq \frac{1}{2g(\mathbf{x}(\ell))} \left(\frac{\mu}{\beta} g(\mathbf{x}(\ell))^2 + \frac{\beta}{2} \left\| \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big) \right\|^2 \right) \\
&\geq \sqrt{\frac{\mu}{2}} \left\| \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big) \right\|.
\end{aligned}$$

The first equality follows from the definition of the gradient flow system and the second equality is due to Lemma 2.5. The first inequality holds because of Lemma 1 of [62]. The second inequality applies the fact that $\mathbf{x}(t)$ satisfies the proximal-PL inequality with the parameters μ and β. The third inequality is the result of the arithmetic-geometric mean inequality. Finally, substituting this lower bound into (2.40) gives

$$\begin{aligned}
\sqrt{f(\mathbf{x}) - f^*} &\geq \sqrt{\frac{\mu}{2}} \left\| \int_0^\infty \Pi_{\mathbb{T}_{\mathbb{S}}(\mathbf{x}(\ell))}\big(-\frac{1}{\beta}\nabla f(\mathbf{x}(\ell))\big)\, \mathrm{d}\ell \right\| \\
&= \sqrt{\frac{\mu}{2}} \left\| \int_0^\infty \dot{\mathbf{x}}(\ell)\, \mathrm{d}\ell \right\| = \sqrt{\frac{\mu}{2}} \|\mathbf{x} - \mathbf{x}^*\|.
\end{aligned}$$

This completes the proof. □

2.B Proof of Lemma 2.1

In this proof, the dependency of the sets $\mathcal{RP}_t^D$ and $\mathcal{RP}_t^C$ on the parameters (μ, β, s) has been omitted in order to simplify notation. The assumption at the end of Section 2.2.2 establishes the existence of $r_1 > 0$ such that $(\mathcal{B}(\mathbf{x}_t^*, r_1) \cap \mathbb{S}) \subset \mathcal{P}_t(\mu, \beta)$. It can be concluded from Proposition 8.5 and Lemma 8.3 in [49], and the initial assumptions that f_t is continuously differentiable and $\mathbb{S}$ is compact, that the projected gradient flow system described in (2.13) converges to the set of critical points of (2.1) and the sublevel sets of f_t are invariant under this system. Define $\bar{l}_t$ as the second-lowest objective value among all critical points of (2.1). Since $f_t^* < \bar{l}_t$, there exists an $r_2 > 0$ such that

$$(\mathcal{B}(\mathbf{x}_t^*, r_2) \cap \mathbb{S}) \subset \mathcal{L}_t(\bar{l}_t) \text{ and}$$
$$(\mathcal{B}(\mathbf{x}_t^*, \min\{r_1, r_2\}) \cap \mathbb{S}) \subset \mathcal{RP}_t^C.$$

Meanwhile, Lemma 2.6 (see Appendix 2.C) and its proof show that there exists an open set $\mathcal{S}_t^D \subset \mathcal{RA}_t^D$ such that $\mathbf{x}_t^* \in \mathcal{S}_t^D \subset (\mathcal{B}(\mathbf{x}_t^*, \epsilon) \cap \mathbb{S})$ for arbitrary ϵ. Take $\epsilon = r_1$. Then by definition of an open set, for some $r_3 > 0$,

$$(\mathcal{B}(\mathbf{x}_t^*, r_3) \cap \mathbb{S}) \subseteq \mathcal{S}_t^D \subset \mathcal{RP}_t^D.$$

If $\alpha_t > f_t^*$, then $\mathcal{L}_t(\alpha_t) \supseteq (\mathcal{B}(\mathbf{x}_t^*, r_4) \cap \mathbb{S})$ for some $r_4 > 0$, and $r = \min\{r_1, r_2, r_3, r_4\}$ satisfies Lemma 1. Thus, it remains to show that $\exists \alpha_t' > f_t^*$ satisfying statement (2.15), since $\alpha_t \geq \alpha_t'$. Define $\bar{r} := \min\{r_1, r_2, r_3\}$ and take $\alpha_t' = f_t^* + \bar{r}\mu/2$. By Theorem 2.2,

$$\|\mathbf{y} - \mathbf{x}_t^*\|^2 \leq \bar{r} \ \ \forall \mathbf{y} \in \mathcal{L}_t(\alpha_t')$$

which is equivalent to

$$\mathcal{L}_t(\alpha_t') \subset (\mathcal{B}(\mathbf{x}_t^*, \bar{r}) \cap \mathbb{S}).$$

Therefore, $\mathcal{L}_t(\alpha_t') = \mathcal{L}_t(\alpha_t') \cap \mathcal{RP}_t^D \subset \mathcal{RP}_t^C$, completing the proof. □

2.C Capture property

Lemma 2.6. *Let f be a continuously differentiable function on a compact, convex set $\mathbb{S}$. Let $\{\mathbf{x}^k\}$ be a sequence of points in $\mathbb{S}$ satisfying $f(\mathbf{x}^{k+1}) \leq f(\mathbf{x}^k)$ generated by the projected gradient descent method $\mathbf{x}^{k+1} = \Pi_{\mathbb{S}}(\mathbf{x}^k - s\nabla f(\mathbf{x}^k))$, which is convergent in the sense that every limit point of such sequences is a stationary point of $f(\mathbf{x})$. Let $\mathbf{x}^*$ be a local minimum of $\min_{\mathbf{x}\in\mathbb{S}} f(\mathbf{x})$, which is the only stationary point within some open set. Then there exists an open set $\mathbb{B}$ containing $\mathbf{x}^*$ such that if $\mathbf{x}^{\bar{k}} \in \mathbb{B}$ for some $\bar{k} \geq 0$, then $\mathbf{x}^k \in \mathbb{B}$ for all $k \geq \bar{k}$ and $\{\mathbf{x}^k\} \to \mathbf{x}^*$.*

Proof: Let $\rho > 0$ be a constant such that

$$f(\mathbf{x}^*) < f(\mathbf{x}), \quad \forall \mathbf{x} \neq \mathbf{x}^* \text{ with } \|\mathbf{x} - \mathbf{x}^*\| \leq \rho.$$

For every $\delta \in [0, \rho]$, define

$$\phi(\delta) = \min_{\{\mathbf{x} | \delta \leq \|\mathbf{x} - \mathbf{x}^*\| \leq \rho\}} f(\mathbf{x}) - f(\mathbf{x}^*).$$

Note that $\phi(\delta)$ is a monotonically non-decreasing function of δ, and that $\phi(\delta) > 0$ for all $\delta \in (0, \rho]$. Given any $\epsilon \in (0, \rho]$, let $r \in (0, \epsilon]$ be such that

$$\|\mathbf{x} - \mathbf{x}^*\| < r \Rightarrow \|\mathbf{x} - \mathbf{x}^*\| + \frac{1}{\beta}\|\nabla f(\mathbf{x})\| < \epsilon.$$

Consider the open set

$$\mathbb{B} = \{\mathbf{x} \in \mathbb{S} \mid \|\mathbf{x} - \mathbf{x}^*\| < \epsilon, \ f(\mathbf{x}) < f(\mathbf{x}^*) + \phi(r)\}.$$

We claim that if $\mathbf{x}^k \in \mathbb{B}$ for some k, then $\mathbf{x}^{k+1} \in \mathbb{B}$. In order to prove the claim, assume that $\mathbf{x}^k \in \mathbb{B}$. Then,

$$\phi(\|\mathbf{x}^k - \mathbf{x}^*\|) \leq f(\mathbf{x}^k) - f(\mathbf{x}^*) < \phi(r),$$

where the first inequality is due to $\phi(\|\mathbf{x}^k - \mathbf{x}^*\|) = \min_{\{\mathbf{x} | \|\mathbf{x}^k - \mathbf{x}^*\| \leq \|\mathbf{x} - \mathbf{x}^*\| \leq \rho\}} f(\mathbf{x}) - f(\mathbf{x}^*) \leq f(\mathbf{x}^k) - f(\mathbf{x}^*)$ and the second inequality is due to the fact that $\mathbf{x}^k \in \mathbb{B}$. Since $\phi(\cdot)$ is monotonically non-decreasing, the above statement implies that $\|\mathbf{x}^k - \mathbf{x}^*\| < r$, which means that

$$\|\mathbf{x}^k - \mathbf{x}^*\| + \frac{1}{\beta}\|\nabla f(\mathbf{x}^k)\| < \epsilon. \tag{2.42}$$

We also know that

$$\begin{aligned}
\|\mathbf{x}^{k+1} - \mathbf{x}^*\| &= \|(\mathbf{x}^{k+1} - \mathbf{x}^k) + (\mathbf{x}^k - \mathbf{x}^*)\| \\
&\leq \|\mathbf{x}^{k+1} - \mathbf{x}^k\| + \|\mathbf{x}^k - \mathbf{x}^*\| \\
&= \|\Pi_{\mathbb{S}}(\mathbf{x}^k - \frac{1}{\beta}\nabla f(\mathbf{x}^k)) - \Pi_{\mathbb{S}}(\mathbf{x}^k)\| + \|\mathbf{x}^k - \mathbf{x}^*\| \\
&\leq \|\big(\mathbf{x}^k - \frac{1}{\beta}\nabla f(\mathbf{x}^k)\big) - \mathbf{x}^k\| + \|\mathbf{x}^k - \mathbf{x}^*\| \qquad (2.43) \\
&= \|\frac{1}{\beta}\nabla f(\mathbf{x}^k)\| + \|\mathbf{x}^k - \mathbf{x}^*\| < \epsilon
\end{aligned}$$

where equation (2.43) follows from the non-expansive property of the projection operator (when projected onto convex sets) and the final inequality follows from applying equation (2.42). Therefore by induction, this implies that if $\mathbf{x}^{\bar{k}} \in \mathbb{B}$ for some $\bar{k}$, we have $\mathbf{x}^k \in \mathbb{B}$ for all $k \geq \bar{k}$. Let $\overline{\mathbb{B}}$ be the closure of $\mathbb{B}$. Since $\overline{\mathbb{B}}$ is compact, the sequence $\{\mathbf{x}^k\}$ will have at least one limit point, which by assumption must be a stationary point of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$. The only stationary point of $\min_{\mathbf{x} \in \mathbb{S}} f(\mathbf{x})$ within $\overline{\mathbb{B}}$ is $\mathbf{x}^*$ since $\|\mathbf{x} - \mathbf{x}^*\| < \epsilon \leq \rho$ for all $\mathbf{x} \in \overline{\mathbb{B}}$. Hence, $\mathbf{x}^k \to \mathbf{x}^*$. □

Chapter 3

Smoothing Property of Load Variation Promotes Finding Global Solutions of Time-Varying Optimal Power Flow

3.1 Introduction

Optimal power flow (OPF) is a large-scale optimization problem at the core of the daily operation of power systems world-wide. OPF aims to find a cost-minimizing operating point for a power system, subject to various operational and security constraints [91]. The OPF problem is challenging because of its nonconvexity and the frequency at which it is solved [8]. Because demand across the system is constantly in flux, the OPF problem is solved every few minutes to match the system's power generation with its latest demand profile. Nonconvex constraints in the AC model of OPF are the main impediment to solving the problem efficiently and optimally. Physical laws govern these constraints, indicating nonconvexity is inherent to the problem. In power systems [13, 128] and in machine learning [129], such nonconvexity is known to give rise to poor local solutions. To realize the vision of sustainable and resilient power grids, there is a pressing need to address the nonconvexity and timescale of both existing and emerging optimization problems for the control and operation of the grid. Since these problems are all built upon the power flow equations, we focus on OPF in this paper.

With the goal of addressing the underlying nonconvexity of the problem, a recent line of research has focused on approximating the problem as a single or sequence of convex optimization problems. These works include quadratic convex [22], second-order conic programming [66], and semidefinite programming [73, 108, 61] relaxations. Despite desirable theoretical guarantees, the convex relaxations of OPF suffer from two major drawbacks: 1) Their global guarantees often come at the expense of higher runtimes or overly compli-

cated implementations; 2) They do not account for the time-varying nature of demand. This time-varying property poses additional constraints on the ramping capabilities of generators, which in turn gives rise to coupled optimization problems.

On the other hand, research on multiperiod OPF, such as [68, 45], and dyanamic OPF, such as [42, 23], endeavors to solve multiple such time-coupled OPF problems simultaneously. This leads to large problem formulations which are still nonconvex in nature. As a result, solution strategies for these problems often rely on the convex relaxations discussed previously in combination with receding horizon approaches or nonlinear programming algorithms, which lack global optimality guarantees [68]. Another drawback is that the data for all time periods must be specified at the outset. In practice, forecasts may not be adequately accurate far in advance.

Real-time OPF is another approach which targets the timescale of OPF. In [110] a real-time algorithm is used to track the optimal solution every few seconds in between traditional OPF updates, which occur on a slower timescale ranging from every 5 to 30 minutes. It uses new measurements of the decision variables' values and constraints at every time step in order to compute a correction and track the optimal solution. The correction is computed by solving a quadratic optimization problem with one iteration of a quasi-Newton algorithm. This has the advantage of responding quickly to fluctuations, but does not replace the need to solve OPF on the traditional timescale. Other faster-timescale approaches to OPF-related problems include [26, 12, 38].

In this work, which is positioned between MPOPF and real-time OPF, we consider time-varying OPF with ramping constraints in an online fashion, where the load profile changes over time. Unlike the previous convexification techniques, we solve the problem sequentially using a simple local-search algorithm. Due to the nonconvex nature of the problem, the local-search algorithm may return a spurious (non-global) local solution, thus leading to a potentially large optimality gap. Previously in [34], we made the observation that for a small system with time-varying demand, the solution trajectories of the time-varying OPF stemming from four initial local solutions could converge over time. Here, we present an extensive empirical study on a larger system with 16 spurious solutions using California load data, and show that all feasible local solution sequences (also called trajectories) converge in cost and value to the best solution. Notably, this phenomenon occurs despite the fact that the problem has multiple point-wise poor local minima at key times. For this system, we show that there is an *escaping period* in which different local solution trajectories converge to a solution with lowest cost, followed by a *tracking period* in which the local trajectories closely track the global solution.

This observation leads to an important phenomenon in time-varying OPF: *load variation enables the local solution trajectories to avoid poor solutions over time.*[1] In other words, despite the highly nonconvex nature of the OPF problem at any given time, our numerical algorithm acts on an implicitly smoothed and well-behaved variant of the problem, thereby

[1]Note that with constant (time-invariant) load, all the local solution trajectories will remain unchanged over time.

avoiding the undesirable local solutions over time. We will formalize this statement in the paper by providing a backward-in-time mapping from the globally optimal solutions of OPF at a given time (namely, end of the escaping period) to the set of desirable initial points. By leveraging its special structure, we show that the proposed backward mapping may act as a stochastic gradient ascent algorithm on an implicitly convexified formulation of the OPF problem, which in turn explains why local solution trajectories could avoid poor solutions over time. This work is the first studying the role of data variation in reducing the complexity of power optimization problems. Since it relies on simple local search methods, the solution techniques have extremely low memory and time complexities and can also be implemented in a distributed setting to accommodate the distributed nature of future grids [5].

3.2 Empirical Study of Time-Varying OPF

In this section, we analyze the local solution trajectories of time-varying OPF primarily for a 39-bus system. A secondary analysis on a 9-bus system is also shared to highlight that the observed behavior is not unique to the 39-bus system. The solution trajectories of time-varying OPF are constructed by sequentially solving a series of optimization problems with time-varying demand levels using a local-search algorithm. California load data and synthetic load scenarios are used to determine demand levels over time. To prevent the solution from changing abruptly over a short period of time, the sequential optimization problems are coupled via so-called *ramping constraints*, as we explain below.

3.2.1 Model and Algorithm Details

To examine the behavior of different local solution trajectories, we consider a modified version of the IEEE 39-bus system, as introduced in [13]. Specifically, the real and reactive power demands are reduced by 50%, voltage limits tightened from +/-6% to +/-5%, and the cost functions associated with all generators are assumed to be linear. The OPF problem for this system with a generation cost-minimizing objective and fixed demand values is known to have 16 local solutions. In this work, we take into account the time-varying nature of the load and scale all demands proportionally to a given load profile. Finally, we introduce the ramping constraints that limit the change in power generation for each generator over time.

Starting from the 16 known initial local solutions, we constructed the sequences of local trajectories using the MATPOWER optimization toolbox [130] and `fmincon` sequential quadratic programming solver[2] in the following procedure. We ran Algorithm 3.1 for all 16 initial local solutions and obtained 16 different solution sequences, which are called *discrete local trajectories* [34].

[2]Note that unlike many interior point methods that require strictly feasible initial points, `fmincon` sequential quadratic programming gives a second-order critical point even if the initial point is not strictly feasible.

Algorithm 3.1 Algorithm for obtaining discrete local trajectories

Input: Power system model with a fixed initial point $\mathbf{x}_0$, demand curve, ramping constraint specifications

Output: Discrete local trajectory $\{\mathbf{x}_t\}_{t=0}^{K}$

1: **Initialization** : $t = 1$
2: **for** every 15-minute time increment over 24 hours **do**
3: Set demand constraints for each bus according to the demand curve at time t.
4: Set generator production limits based on $\mathbf{x}_{t-1}$ and the ramping constraint.
5: Solve the resulting cost-minimization OPF problem with fixed demand and initial point $\mathbf{x}_{t-1}$ using `fmincon`. Upon feasibility, collect the solution as $\mathbf{x}_t$.
6: **end for**
7: **return** $\{\mathbf{x}_t\}_{t=0}^{T}$

3.2.2 Behavior of Discrete Local Trajectories for a 39-bus System with California Data

In this example, the shape of the demand curve is based on the California's net load for an average day in August 2019 [15] (Fig. 3.2.1). The reported actual hourly net load data was interpolated linearly to produce a net load estimate for each 15-minute interval within 24 hours. The curve is normalized and shifted so that time 0 represents 3:00 a.m. Here, the maximum magnitude of allowable change in power generation between two consecutive time steps is 10% of the capacity of each generator. All 16 discrete local trajectories remain feasible throughout the span of twenty-four hours. (This is not guaranteed, as local search may not always find a feasible point or such point may not even exist.) Fig. 3.2.2 shows the point-wise distance between these feasible trajectories and the feasible trajectory with the lowest cost (labeled as *Trajectory 2*). Interestingly, all 16 trajectories converge to *Trajectory 2* within nine hours.

Based on this observation, one may speculate that the problem becomes devoid of spurious local solutions over time. This is not the case for the considered problem. We uniformly searched the feasible region of the problem without ramping constraints and verified that

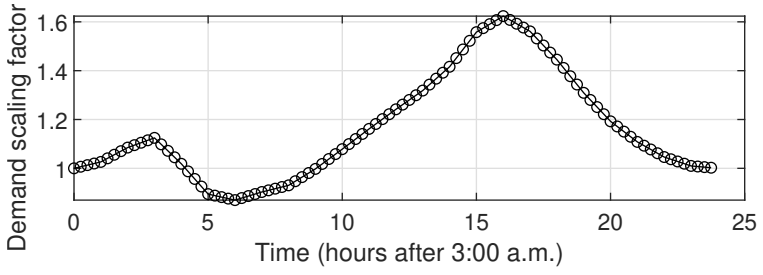

Figure 3.2.1: Average daily net load for California during August 2019 [15]

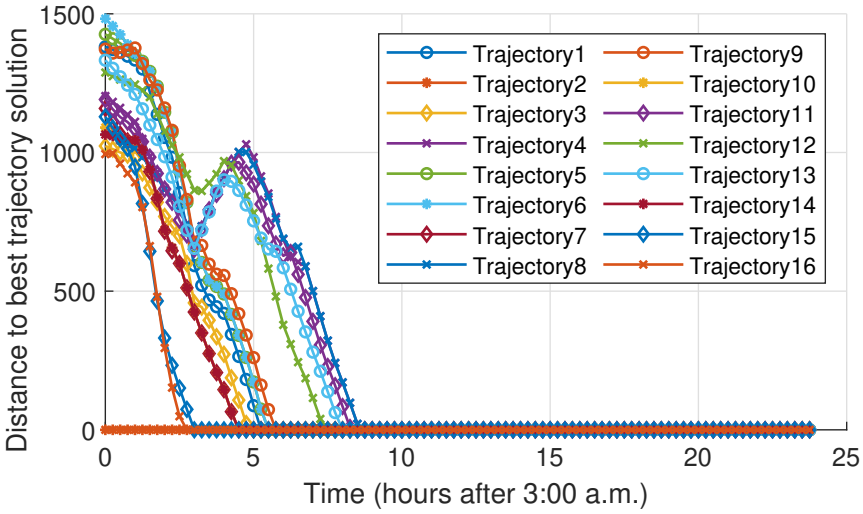

Figure 3.2.2: Solution convergence for points on discrete local trajectories

there are multiple point-wise spurious local solutions for the point-wise (single time instance, without ramping constraints) OPF problem at different times. In particular, there are many local solutions around the *escape time* (hour 9) when the discrete local trajectories merge into one trajectory. Fig. 3.2.3 shows the normalized objective cost values for different discrete local trajectories, alongside the costs of the discovered point-wise local solutions. Despite the existence of multiple sub-optimal operating points at different times, the discrete local trajectories initialized at various local solutions result in the lowest cost values over time. Fig. 3.2.4 examines the active and reactive power generation for two representative generators. This figure shows that the problem has point-wise local solutions with a wide range of generation levels, highlighting the importance of finding the solution with the lowest cost.

Observe that most of the spurious point-wise local solutions have *sharp* and random nature. In other words, they appear at different time-steps with various cost values, and then quickly disappear after a short period of time. This implies that the landscape of OPF may be highly nonconvex at any given time step. However, it can be observed that our numerical algorithm is not affected by such sharp local solutions. To explain this phenomena, we will show in Section 3.4 that the data variation enables the solver to act on a *smoothed version* of the problem that is devoid of sharp local minima.

3.2.3 Behavior of Discrete Local Trajectories for a 9-bus System with California Data

In this example, we consider a modified version of the IEEE 9-bus system with 4 known local solutions to the OPF problem, as introduced in [13]. Specifically, the active and reactive power demands are reduced by 40% and the lower bounds on reactive power compensation are

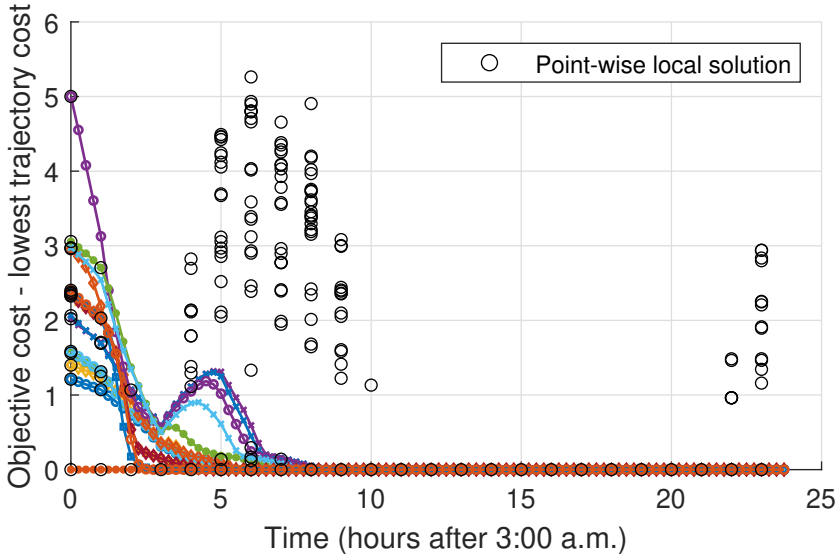

Figure 3.2.3: Cost for points on discrete local trajectories and point-wise local solutions (for a single instance of OPF), relative to the cost of the best trajectory

tightened to -5 Mvar. The demand data is a normalized and shifted version of California's net load for an average day in May 2019 [15] (Fig. 3.2.5a). Fig. 3.2.5b shows the relative objective cost values for different discrete local trajectories, alongside the costs of the discovered point-wise local solutions, produced using Algorithm 1 with a 5% ramping constraint. Again, we observe that the load variation enables all trajectories to converge to the optimal trajectory.

3.2.4 Impact of Load Variation on 39-bus System

Next, we consider discrete local trajectories for three different load profiles on the same 39-bus system. Isolating the impact of load variation enables insight into how variation creates trajectories that avoid poor solutions, as occurred in the previous examples. The three demand curves used are sinusoidal functions with amplitudes representing 5%, 10% and 12% deviation from the initial load, as shown in the left column of Fig. 3.3.1. The ramping constraint (i.e., maximum magnitude of allowable change in power generation between two consecutive time steps) is 5% of the capacity of each generator. In each scenario, all 16 discrete local trajectories remain feasible throughout the time horizon (100 steps).

The results show that larger magnitudes of data variation lead to fewer poor solutions over time. At 5% variation, 4 trajectories remain at 4 different poor solutions, while the remaining 12 trajectories converge to the best solution. At 10% variation, 3 trajectories converge to the same poor solution, while the remaining 13 trajectories converge to the best solution. At 12% variation, all 16 trajectories converge to the best known solution. These results are displayed in the center column of Fig. 3.3.1, which shows the distance between each trajectory and

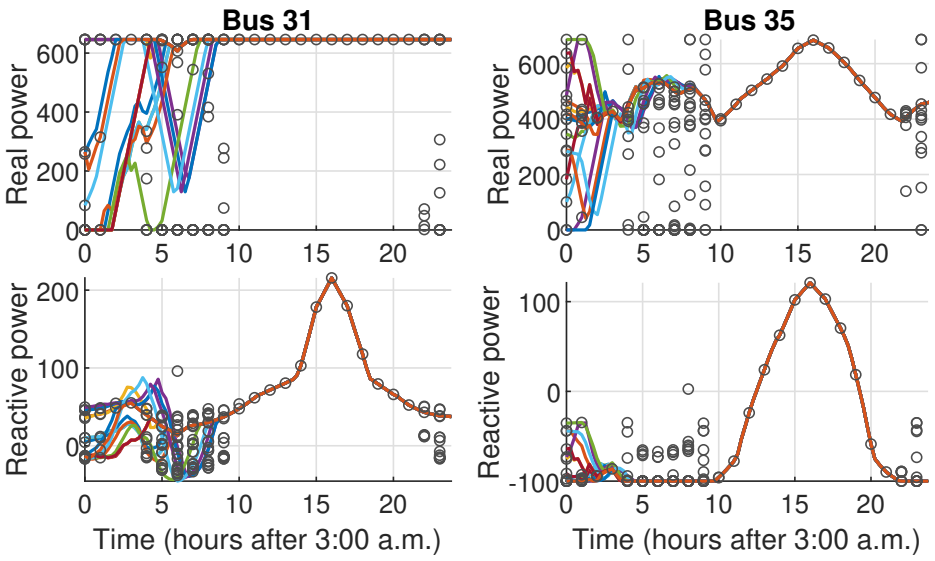

Figure 3.2.4: Real and reactive power output of select generators: points on discrete local trajectories and point-wise local solutions

the trajectory with the lowest cost, along with discovered point-wise local solutions. The search for point-wise local solutions is done every fourth time step due to the significant computational effort required to repeatedly solve the problem from a range of initial points. Fig. 3.3.1 (right column) compares the number of point-wise local solutions with the number of distinct[3] trajectories over time. In these three cases, *the number of distinct trajectories decreases until it plateaus at the minimum number of point-wise local solutions found over the entire period.* This offers one potential explanation of how load variation creates trajectories that escape poor solutions: In exploring a range of static problems, you may encounter one or more times at which the problem has a favorable landscape[4]. At such times, the coupled problem may escape a poor solution. Eventually, the number of poor trajectories is limited by the number of spurious point-wise local solutions of the most favorable landscape.

3.3 Mathematical Analysis of Time-Varying OPF

The case study in Section 3.2 reveals an important property of the time-varying OPF problem: In the *escaping period*, different discrete local trajectories converge to the operating

[3]Solutions are considered distinct if the real or reactive power output at any generator differs by at least 1 MW or 1 MVAr, respectively, or if the voltage magnitude or angle at any bus differs by at least 10^{-3} p.u. (345V) or 10^{-3} radians, respectively.

[4]The number of spurious point-wise local solutions is an indicator of how difficult a given static OPF problem is. If only one point-wise local solution is found, the problem may be convex. However, the search is not exhaustive, so other local minima with small regions of attraction may exist.

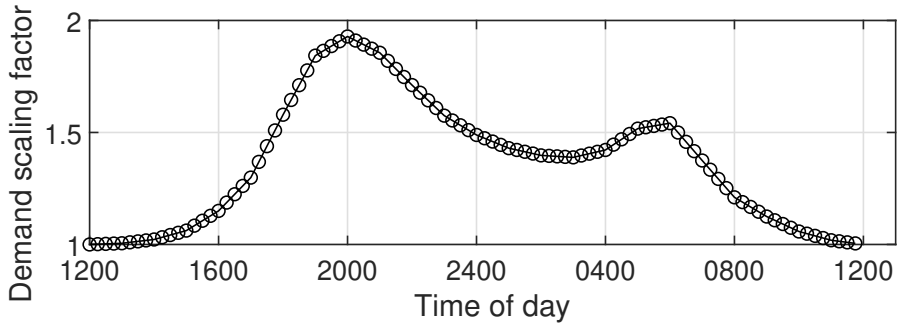

(a) Average daily net load for California during May 2019 [15]

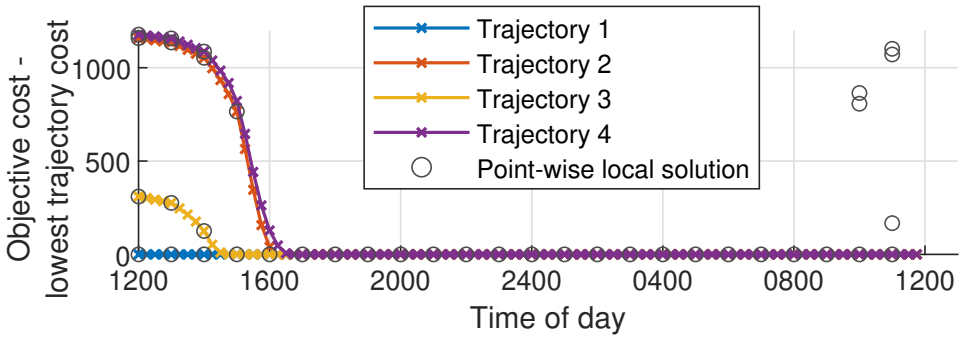

(b) Cost for points on discrete local trajectories and point-wise local solutions (for a single instance of OPF), relative to the cost of the best trajectory

Figure 3.2.5: Data and results for an empirical study on a 9-bus system

point with the lowest cost. Then, in the *tracking period*, the discrete local trajectories track these globally optimal operating points, even if the load profile changes gradually over time. Such *tracking period* has been studied in [86, 111], but the striking feature of power systems is the existence of escaping periods.

To better understand this phenomenon, we analyze the problem structure mathematically. First, we reformulate the time-varying OPF as an unconstrained optimization problem to enable the analysis. Using the derived unconstrained problem, we introduce a backward mapping that characterizes the dynamics of the discrete local trajectories over time. We show that the convergence of different local trajectories can be explained by the expansive property of this backward mapping. Finally, in Section 3.4 we draw a novel connection between our derived mapping and stochastic gradient ascent and use this insight to explain that the behavior of the trajectories may be driven by some low-complexity averaged model over a period, rather than the high-complexity OPF problems at each step.

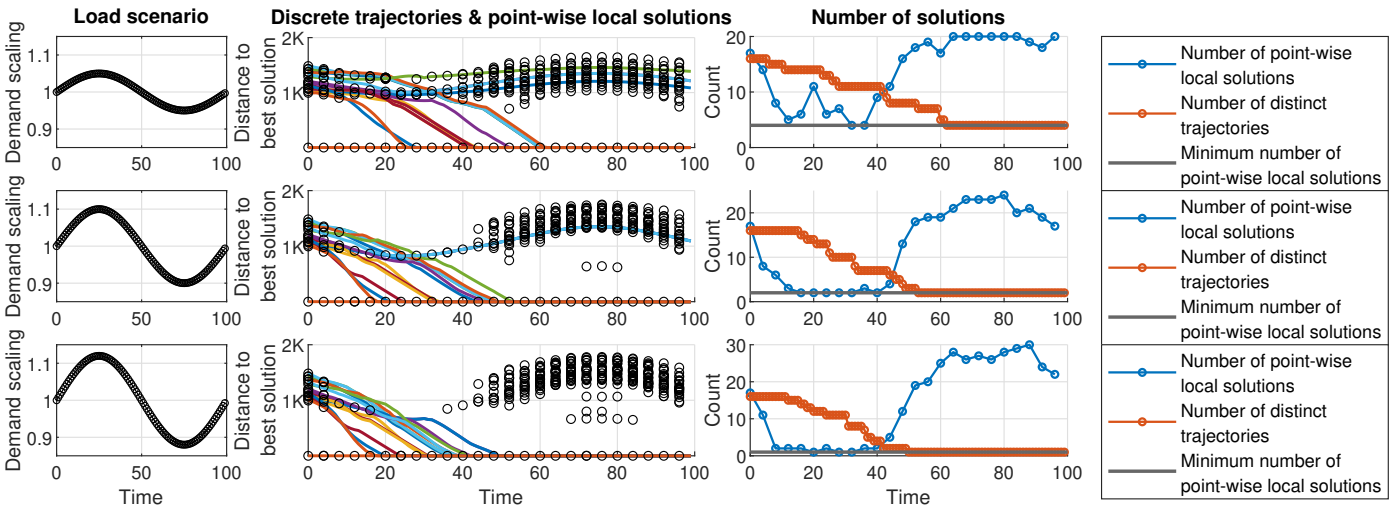

Figure 3.3.1: Three scenarios highlighting the role of load variation. The three plots for each scenario are (left to right): load profile input, resulting discrete local trajectories and point-wise local solutions (for a single instance of OPF), comparison of the number of point-wise local solutions to the number of distinct discrete local trajectories.

3.3.1 Unconstrained Model for OPF with Fixed Demand

The AC model of OPF in a single time instance with fixed and predefined demand values can be written compactly as an optimization problem with both equality and inequality constraints:

$$\min_{\mathbf{x}\in\mathbb{R}^p} f(\mathbf{x}) \tag{3.1a}$$

$$\text{s.t. } h(\mathbf{x}) = \mathbf{d} \in \mathbb{R}^n \tag{3.1b}$$

$$g(\mathbf{x}) \leq \mathbf{0} \in \mathbb{R}^m \tag{3.1c}$$

Here, $\mathbf{x}$ is the concatenation of the voltage angle and magnitude at each bus, as well as the real and reactive power generation outputs for each generator. The equality constraint (3.1b) ensures that the generated power meets the demand, where $\mathbf{d}$ is the vector of real and reactive demand at each bus, and respects the underlying structure and physical constraints of the network. The remaining constraints in the problem—including the upper and lower bounds on the voltage magnitudes and degrees, power generation, and line flows—are captured by the inequality constraint (3.1c). It is easy to verify that $p > n$. We refer the reader to [91], [13] and [73] for more information on the exact formulation of the problem. Note that $f(\mathbf{x})$, $h(\mathbf{x})$, and $g(\mathbf{x})$ are continuously differentiable (piecewise linear cost functions can be reformulated as such).

In order to analyze this optimization problem theoretically, it is desirable to convert it to an unconstrained problem. First, we enforce the inequality constraints (3.1c) through a

penalty in the objective function:

$$\min_{\mathbf{x}\in\mathbb{R}^p} f(\mathbf{x}) + \beta \sum_{i=1}^{m} \left([g_i(\mathbf{x})]^+\right)^2 \tag{3.2a}$$

$$\text{s.t. } h(\mathbf{x}) = \mathbf{d} \in \mathbb{R}^n \tag{3.2b}$$

where $\beta > 0$ serves as the penalization parameter, $g_i(\cdot)$ is the i^{th} element of $g(\cdot)$, and $[y]^+$ denotes $\max(y, 0)$. This choice of quadratic penalty function is inexact, meaning that problem (3.2) is an approximation of problem (3.1). However, as β increases, each global minimizer of (3.1) approaches a global minimizer of (3.2) under mild regularity conditions [9]. Second, we use the implicit function theorem [9] to complete the transformation to an unconstrained model. Consider a feasible point $\mathbf{x}_\star$ satisfying the Karush-Kuhn-Tucker (KKT) conditions for (3.2). Assuming that constraint qualifications hold at $\mathbf{x}_\star$, this vector can be partitioned into two sub-vectors ${\mathbf{x}_\star}^B \in \mathbb{R}^n$ and ${\mathbf{x}_\star}^R \in \mathbb{R}^{p-n}$ such that the Jacobian of $h(\mathbf{x}_\star)$ with respect to $\mathbf{x}^B$ is invertible. Therefore, the implicit function theorem guarantees the existence of a unique differentiable function $\phi(\cdot)$ such that $\mathbf{x}^B = \phi(\mathbf{x}^R)$ in a local neighborhood of $\mathbf{x}_\star$. Given such function, Problem (3.2) can be re-written as (see [9]):

$$\min_{\mathbf{x}^R\in\mathbb{R}^{p-n}} f(\phi(\mathbf{x}^R), \mathbf{x}^R) + \beta \sum_{i=1}^{m} \left([g_i(\phi(\mathbf{x}^R), \mathbf{x}^R)]^+\right)^2 \tag{3.3}$$

Enforcing the equality constraint (3.1b) directly using the implicit function theorem instead of through penalization will be advantageous when we move to the time-varying setting. Namely, it avoids amplifying the demand variation as scaling by a large penalization parameter would do. This is not an issue for the inequality constraint (3.1c) because it does not vary in time.

Remark 3.1. *Note that* (3.3) *cannot be formulated explicitly, due to the unknown nature of the local solution* $\mathbf{x}_\star$ *and the function* $\phi(\mathbf{x}^R)$*. Instead, this formulation serves as an intermediate step to analyze the behavior of discrete local trajectories over time. In other words, one would solve the OPF problem directly in practice, and the surrogate problem* (3.3) *is designed to understand the properties of OPF.*

3.3.2 Unconstrained Model for Time-Varying OPF

The above analysis reveals that, under some technical conditions, the OPF problem with fixed load can be modeled as an unconstrained optimization problem (with a controllable approximation error). In this subsection, we extend our analysis to time-varying OPF where demand changes over time and the problem must respect ramping constraints. As previously stated, ramping constraints ensure that the solution does not change too drastically from one time step to the next. One way to softly impose ramping constraints is through a proximal method, which penalizes the distance between the current and previous solutions in the

objective function of the optimization [31]. Time-varying OPF with K equally-spaced time steps $t_0 = 0, t_1 = \Delta t, ..., t_K = K\Delta t$ ($\Delta t > 0$) can be written as the following sequence of optimization problems:

$$\min_{\mathbf{x}^{R_k} \in \mathbb{R}^{p-n}} f_{t_k}(\phi_{t_k}(\mathbf{x}^{R_k}), \mathbf{x}^{R_k}) \quad + \quad \alpha \left\| \mathbf{x}^{R_k} - \mathbf{x}_{\star t_{k-1}}^{R_k} \right\|_2^2 \quad + \quad \beta \sum_{i=1}^{m} ([g_i \left(\phi_{t_k}(\mathbf{x}^{R_k}), \mathbf{x}^{R_k}\right)]^+)^2 \quad (3.4)$$

for $k = 1, ..., K$, where $\alpha > 0$ is a penalization parameter and $\mathbf{x}_{\star t_{k-1}} = \left[\left(\mathbf{x}_{\star t_{k-1}}^{B_k} \right)^\top \left(\mathbf{x}_{\star t_{k-1}}^{R_k} \right)^\top \right]^\top$ denotes an arbitrary local solution to Problem (3.4) obtained at time t_{k-1}. In light of its dependence on $\mathbf{x}^{R_k}$, $\mathbf{x}^{B_k}$ is not regularized in this approximated model. Due to the time-varying nature of the demand, the functions $f_{t_k}(\cdot)$ and $\phi_{t_k}(\cdot)$ may change over time, hence they are indexed by time step.

To simplify the analysis, assume that the partition (B_k, R_k) does not change over time, i.e., we have $B_k = B$ and $R_k = R$ for $k = 1, ..., K$. Then problem (3.4) can be written as

$$\min_{\mathbf{z} \in \mathbb{R}^{p-n}} \underbrace{F_k(\mathbf{z}) + \beta \sum_{i=1}^{m} \left([G_{k,i}(\mathbf{z})]^+ \right)^2}_{\Gamma_k(\mathbf{z})} + \alpha \left\| \mathbf{z} - \mathbf{z}_{k-1} \right\|_2^2 \quad (3.5)$$

for $k = 1, ..., K$, where $\mathbf{z} = \mathbf{x}^{R_k}$, $\mathbf{z}_{k-1} = \mathbf{x}_{\star t_{k-1}}^{R_k}$, $F_k(\mathbf{z}) := f_{t_k}(\phi_{t_k}(\mathbf{z}), \mathbf{z})$, and $G_{k,i}(\mathbf{z}) := g_i\left(\phi_{t_k}(\mathbf{z}), \mathbf{z}\right)$. If the partition changes, then the time interval $[0, K\Delta t]$ should be divided into sub-intervals, each with a constant partitioning of $\mathbf{x}$. In this case, the argument presented in Section 3.4 applies to each sub-interval.

3.3.3 Backward-In-Time Mapping

The above analysis reveals that a local-search algorithm used to solve the time-varying OPF implicitly aims to recover a stationary point of the unconstrained problem (3.5). Therefore, we focus on (3.5) in our subsequent analysis. Consider a given time step $T\Delta t$, representing the end of the escaping period. Then, a sequence of stationary points $\{\mathbf{z}_k\}_{k=1}^T$ for (3.5) satisfies

$$0 = \nabla \Gamma_k(\mathbf{z}_k) + 2\alpha(\mathbf{z}_k - \mathbf{z}_{k-1}) \quad (3.6)$$

for every $k = 1, 2, \ldots, T$ (where ∇ is the gradient operator). Note that $\Gamma_k(\cdot)$ is differentiable. Therefore, given the solution $\mathbf{z}_{k-1}$, this equation defines an implicit nonlinear formula for obtaining $\mathbf{z}_k$ which cannot be written in closed form. However, going backward in time, one can express $\mathbf{z}_{k-1}$ in terms of $\mathbf{z}_k$:

$$\mathbf{z}_{k-1} = \mathbf{z}_k + \frac{1}{2\alpha} \nabla \Gamma_k(\mathbf{z}_k) := M_k(\mathbf{z}_k) \quad (3.7)$$

This gives rise to the following end-to-end backward mapping from $\mathbf{z}_T$ to the initial point $\mathbf{z}_0$ via the composition operator $\circ$:

$$\mathbf{z}_0 = M_1 \circ M_2 \circ \cdots \circ M_T(\mathbf{z}_T) \tag{3.8}$$

Provided that the mapping $M_1 \circ \cdots \circ M_T(\mathbf{z}_T)$ is expansive enough when applied to a small neighborhood of a global solution of OPF at time $T\Delta t$, a large set of initial points (even multiple local solutions of OPF at time 0) are guaranteed to converge to that small neighborhood of the globally optimal solution of the problem at time $T\Delta t$. This expansive nature of the mapping implies the escape of spurious local solutions between time 0 and time $T\Delta t$. The global solutions at future times after $T\Delta t$ will be tracked successfully if the data variation is not too high [86]. This expansive property can be observed in the empirical study conducted in Section 3.2 on the modified IEEE 39-bus and 9-bus system under both California load data and synthetic sinusoidal loads.

3.4 Connection to Stochastic Gradient Ascent

This section aims to explain how data variation plays a key role in escaping spurious local solutions of time-varying OPF. Specifically, we will show that the backward mapping (3.7) can be treated as a variant of stochastic gradient ascent on a smoothed version of the function $\Gamma_T(\mathbf{z})$. This gives rise to the following important observation:

A certain level of stochasticity in $\{\Gamma_k(\mathbf{z})\}_{k=1}^T$ over time may enable the stationary points $\{\mathbf{z}_k\}_{k=1}^T$ to escape "sharp" local minima over time.

To explain this phenomenon, we first introduce the smoothing property of the stochastic gradient descent (SGD) algorithm.

Smoothing property of SGD: Recently, [65] proposed an alternative viewpoint to SGD and its ability to avoid spurious sharp local minima. Given an initial point $\mathbf{z}_0$, suppose our goal is to find a global minimum of a (time-invariant) function $\Gamma(\mathbf{z})$ using SGD. Accordingly, the iterations of SGD can be written

$$\mathbf{z}_{k+1} = \mathbf{z}_k - \eta(\nabla\Gamma(\mathbf{z}_k) + \omega_k) \qquad \forall k \in \{0, 1, 2, \dots\} \tag{3.9}$$

where ω_t is a bounded random variable with zero mean and η is a predefined step size. Upon defining $\tilde{\mathbf{z}}_k = \mathbf{z}_k - \eta\nabla\Gamma(\mathbf{z}_k)$, one can write the above iterations (3.9) in terms of the intermediate sequence $\{\tilde{\mathbf{z}}_k\}$:

$$\tilde{\mathbf{z}}_{k+1} = \tilde{\mathbf{z}}_k - \eta\omega_k - \eta\nabla\Gamma(\tilde{\mathbf{z}}_k - \eta\omega_k), \forall k \in \{0, 1, 2, \dots\} \tag{3.10}$$

To analyze the average behavior of SGD, consider the evolution of $\mathbb{E}_{\omega_k}(\tilde{\mathbf{z}}_{k+1})$, where the expectation is taken over ω_k conditioned on $\{\omega_0, \dots, \omega_{k-1}\}$. Hence,

$$\mathbb{E}_{\omega_k}[\tilde{\mathbf{z}}_{k+1}] = \tilde{\mathbf{z}}_k - \eta\nabla\mathbb{E}_{\omega_k}\left[\Gamma(\tilde{\mathbf{z}}_k - \eta\omega_k)\right], \forall k \in \{0, 1, 2, \dots\} \tag{3.11}$$

Therefore, on average, SGD acts as the exact gradient descent on the surrogate function $\mathbb{E}_{\omega_k}[\Gamma(\tilde{\mathbf{z}}_k - \eta\omega_k)]$. Comparing this function with $\Gamma(\mathbf{z})$, one can verify that the former is a smoothed version of the latter, where the smoothness is due to the convolution of $\Gamma(\mathbf{z})$ with the probability density function of the random variable ω_k. As illustrated in [65], such convolution may give rise to (one-point) strong convexity of $\mathbb{E}_{\omega_k}[\Gamma(\tilde{\mathbf{z}}_k - \eta\omega_k)]$ with respect to the globally optimal solution, which in turn guarantees the convergence of $\{\tilde{\mathbf{z}}_k\}$ (and hence $\{\mathbf{z}_k\}$) to a small neighborhood around the global solution, even in the presence of sharp local minima. A key takeaway from this observation is that $\Gamma(\mathbf{z})$ can possess multiple sharp, poor local minima, and yet its smoothed version $\mathbb{E}_{\omega_k}[\Gamma(\tilde{\mathbf{z}}_k - \eta\omega_k)]$ may be devoid of such solutions.

Time-varying optimization and time-varying OPF: Returning to time-varying OPF and the backward mapping (3.7), we assume that the variation in $\{\nabla\Gamma_k(\mathbf{z})\}_{k=1}^T$ follows a stochastic process indexed by the time k. In particular, we write $\nabla\Gamma_k(\mathbf{z}) - \nabla\Gamma_{k+1}(\mathbf{z}) = \zeta_k(\mathbf{z}) + \omega_k$, where $\zeta_k(\mathbf{z})$ is a deterministic, time-varying function and ω_k is a bounded random variable with zero mean. Such assumption is realistic in power systems, where demand can be modeled as a deterministic, time-varying function capturing the average demand behavior, together with an additive stochastic term accounting for its random nature. The iteration (3.7) is equivalent to

$$\begin{aligned}\mathbf{z}_k =& \mathbf{z}_{k+1} + \frac{1}{2\alpha}\nabla\Gamma_T(\mathbf{z}_{k+1}) \\ &+ \frac{1}{2\alpha}\sum_{\tau=k+1}^{T-1} \underbrace{\left(\nabla\Gamma_\tau(\mathbf{z}_{k+1}) - \nabla\Gamma_{\tau+1}(\mathbf{z}_{k+1})\right)}_{\zeta_\tau(\mathbf{z}_{k+1}) - \omega_\tau}\end{aligned} \tag{3.12}$$

which can be written as the following dynamical model:

$$\mathbf{z}_k = \mathbf{z}_{k+1} + \frac{1}{2\alpha}\nabla\Gamma_T(\mathbf{z}_{k+1}) + \frac{1}{2\alpha}\nu_{k+1}(\mathbf{z}_{k+1}) \tag{3.13a}$$

$$\nu_{k+1}(\mathbf{z}_{k+1}) = \nu_{k+2}(\mathbf{z}_{k+1}) + \zeta_{k+1}(\mathbf{z}_{k+1}) - \omega_{k+1} \tag{3.13b}$$

where $\nu_{k+1}(\mathbf{z}_{k+1})$ is referred to as the *variation process*. In particular, (3.13b) defines explicit dynamics for the variation process comprised of three parts. The first term $\nu_{k+2}(\mathbf{z}_{k+1})$ captures the correlation between the variation processes at times t_{k+1} and t_{k+2}. The second term $\zeta_{k+1}(\mathbf{z}_{k+1})$ captures the *bias* that is added to the variation process at time t_{k+1}. Lastly, the third term $\omega_{k+1} \sim W(\mathbf{z}_{k+1})$ is an independent noise injected into the variation process at time t_{k+1} (also referred to as *effective noise*). Comparing (3.13) with (3.9), one can verify that (3.13) reduces to stochastic gradient ascent if $\nu_{k+2}(\mathbf{z}_{k+1}) + \zeta_{k+1}(\mathbf{z}_{k+1}) = 0$. Therefore, if ω_{k+1} dominates the first two terms, (3.13) resembles an approximate version of stochastic gradient ascent applied to $\Gamma_T(\mathbf{z})$; otherwise, it is a *biased* and *correlated* version of SGD [20]. Similar to (3.11), this implies that, on average, the points generated via the backward mapping (3.7) would be close to the iterations of the gradient ascent on the smoothed version of $\Gamma_T(\mathbf{z})$. Now, assume that despite the possible existence of multiple spurious and sharp local

minima in $\{\Gamma_k(\mathbf{z})\}_{k=1}^T$, the smoothed version of $\Gamma_T(\mathbf{z})$ after convolution is strongly convex. This together with the expansive nature of gradient ascent on strongly convex functions [103] yields that the end-to-end backward mapping (3.8) is expansive, and the discrete local trajectories can escape poor local solutions over time. We formalize and rigorously analyze this intuition in the next subsection.

3.4.1 Theoretical Analysis of Dynamics

For simplicity of notation, we define $\eta = \frac{1}{2\alpha}$. Furthermore, suppose that $\mathbf{z}^*$ denotes the globally minimum point of $\Gamma_T(\mathbf{z})$. Without loss of generality, $\|\mathbf{v}\|$ is used to refer to the 2-norm of the vector $\mathbf{v}$. We make the following assumptions for the dynamical model (3.13):

Assumption 3.1 (Smoothness)**.** *The following statements hold:*

- *The function $\Gamma_T(\mathbf{z})$ is L-smooth, i.e., we have*

$$\|\nabla\Gamma_T(\mathbf{x}) - \nabla\Gamma_T(\mathbf{y})\| \leq L\|\mathbf{x}-\mathbf{y}\| \ \forall \mathbf{x},\mathbf{y} \in \mathbb{R}^{p-n}. \tag{3.14}$$

- *The functions $\zeta_\tau(\mathbf{z})$ are l-Lipschitz for $\tau = 1, \cdots, T-1$, i.e., we have*

$$\|\zeta_k(\mathbf{x}) - \zeta_k(\mathbf{y})\| \leq l\|\mathbf{x}-\mathbf{y}\| \ \forall \mathbf{x},\mathbf{y} \in \mathbb{R}^{p-n}. \tag{3.15}$$

Assumption 3.2 (Implicit Convexity)**.** *There exists $\mathbf{z}^*$ such that the following statements hold:*

- *(One-point strong convexity of convolution) For every $\mathbf{y}$, there exists $c > 0$ such that*

$$\langle \mathbf{z}^* - \mathbf{y}, -\nabla\mathbb{E}_{\omega\sim W(\mathbf{y})}\left[\Gamma_T(\mathbf{y}-\eta\omega)\right]\rangle \geq c\|\mathbf{y}-\mathbf{z}^*\|^2 \tag{3.16}$$

- *(Bounded one-point curvature of convolution) For every $\mathbf{y}$, there exists $c' > 0$ such that*

$$\langle \mathbf{z}^* - \mathbf{y}, -\sum_{\tau=k+1}^{T-1} \mathbb{E}_{\omega\sim W(\mathbf{y})}\left[\zeta_\tau(\mathbf{y}-\eta\omega)\right]\rangle \geq -c'\|\mathbf{y}-\mathbf{z}^*\|^2 \tag{3.17}$$

for every $k \in \{0, \ldots, T-2\}$.

The existence of L and l which satisfy Assumption 3.1 can be verified for the unconstrained model of the time-varying OPF. Meanwhile, Assumption 3.2 implies that the *convoluted* variant of the objective function at time T is one-point strongly convex. We note that such assumption may not be easily verifiable for the time-varying OPF. However, our simulations strongly support the fact that most of the spurious local solutions in time-varying OPF have a sharp nature, and therefore, they are likely to be absent in the convoluted (smoothed) landscape of the problem.

Under these two assumptions, we present the main theorem of this paper.

Theorem 3.1. *Suppose that $c \geq c'$ and there exists $r \geq 1$ such that $\|\omega_t\| \leq r$ for every t. Define $\lambda := \eta(c - c')$, and assume that $2\eta^2 L < 1$. Then, under Assumptions 3.1 and 3.2, the following inequality holds:*

$$\|\mathbf{z}_T - \mathbf{z}^*\|^2 \leq \frac{1}{1-2\eta^2 L}\left(\mathcal{D} + \frac{\mathbb{E}\left[\|\mathbf{z}_0 - \mathbf{z}^*\|^2\right]}{(1+\lambda)^{T-1}} + \frac{8\eta^2 r^2 T^2}{(1+\lambda)^{T-1}}\right) \tag{3.18}$$

where

$$\mathcal{D} = \left(\frac{4}{\lambda} + \frac{4}{\lambda^2}\right)\eta^3 r^2 l + 16\left(1 + \frac{1}{\lambda}\right)^2 \frac{\eta^2 r^2 (1+2\lambda)^2}{\lambda^2} \tag{3.19}$$

A proof for Theorem 3.1 is provided in appendix 3.A. A number of observations can be made based on this theorem. Not surprisingly, the provided bound on $\|\mathbf{z}_T - \mathbf{z}^*\|$ depends on the *accuracy of the initial point* $\|\mathbf{z}_0 - \mathbf{z}^*\|$. However, the effect of this initial accuracy diminishes exponentially fast with respect to T. Moreover, as $T \to \infty$, the following asymptotic inequality holds:

$$\|\mathbf{z}_T - \mathbf{z}^*\|^2 \leq \frac{\mathcal{D}}{1 - 2\eta^2 L} \tag{3.20}$$

which is independent of the initial point. Another implication of this asymptotic bound is that, for any value of T, Theorem 3.1 can only guarantee the convergence of $\mathbf{z}_T$ to a neighborhood of $\mathbf{z}^*$. This is not surprising if we consider the non-diminishing nature of η and its connection to SGD, as delineated in the introduction of Section 3.4. Indeed, similar results on SGD show that, with non-diminishing step-sizes, the iterations of the algorithm may only converge to a neighborhood of the globally optimal solution [65]. Finally, it is worthwhile to study how $\mathcal{D}$ depends on different parameters of problem, namely η, r, l, L, and $c - c'$. Equation (3.19) reveals that $\mathcal{D}$ is a decreasing function of $c - c'$. Combined with Assumption 3.2, this implies that one-point strong convexity of $\Gamma_t(\mathbf{z})$ for $t = 1, \ldots, T$ has a favorable effect on $\|\mathbf{z}_T - \mathbf{z}^*\|$. Similarly, it can be seen from (3.18) and (3.19) that $\|\mathbf{z}_T - \mathbf{z}^*\|$ decreases as l, L, and the noise values' magnitude (characterized by r) shrink. However, notice that Assumption 3.2 may not be satisfied for small values of noise. Finally, $\mathcal{D}$ does not have a monotone behavior with respect to η. In particular, it can be verified that $\mathcal{D} \to \infty$ if $\eta \to \infty$ or $\eta \to 0^+$. Recalling (3.5) and $\eta = \frac{1}{2\alpha}$, this implies that over- or under-regularization may lead to large values for $\|\mathbf{z}_T - \mathbf{z}^*\|$. This observation is in line with Example 1 of [34], which shows that both small and large regularization may cause the solution trajectory to remain trapped at spurious local solutions of a time-varying optimization.

3.4.2 Illustrative Example on a 2-bus System

In this subsection, we analyze the effect of the load variation on the landscape of a 2-bus system. Our goal is to visualize the smoothing effect of the load variation on the objective

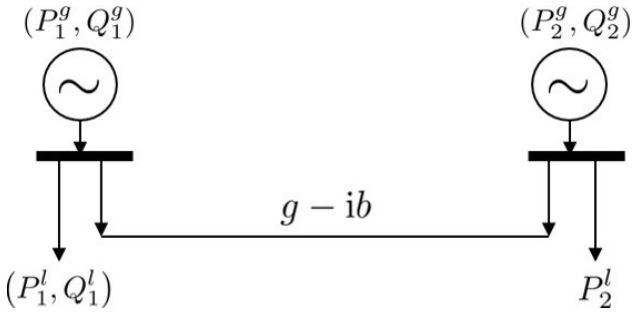

Figure 3.4.1: The 2-bus system. Here, $\mathrm{i} = \sqrt{-1}$.

function, thereby verifying the assumption on the implicit one-point strong convexity of the convoluted objective function. Consider the simple 2-bus system illustrated in Figure 3.4.1. Assume that both buses are equipped with generators, and they are connected via a single line with admittance $g - \mathrm{i}b$. The time-varying load connected to the first bus has both active and reactive power demands, while the time-varying load connected to the second bus is purely active. At any given time k, the point-wise OPF (without ramping constraints) can be formulated as follows[5]:

$$\min\ f_1(P_1^g) + f_2(P_2^g) \tag{3.21a}$$
$$\text{s.t. } P_1^g - P_{1;k}^l = |v_1|^2 g + |v_1||v_2| b \sin(\Delta\theta) - |v_1||v_2| g \cos(\Delta\theta) \tag{3.21b}$$
$$P_2^g - P_{2;k}^l = |v_2|^2 g + |v_1||v_2| b \sin(\Delta\theta) - |v_1||v_2| g \cos(\Delta\theta) \tag{3.21c}$$
$$Q_1^g - Q_{1;k}^l = |v_1|^2 g - |v_1||v_2| g \sin(\Delta\theta) - |v_1||v_2| b \cos(\Delta\theta) \tag{3.21d}$$
$$Q_2^g = |v_2|^2 g - |v_1||v_2| g \sin(\Delta\theta) - |v_1||v_2| b \cos(\Delta\theta) \tag{3.21e}$$
$$V^{\min} \le |v_1| \le V^{\max}, \quad V^{\min} \le |v_2| \le V^{\max} \tag{3.21f}$$
$$P_1^{\min} \le P_1^g \le P_1^{\max}, \quad P_2^{\min} \le P_2^g \le P_2^{\max} \tag{3.21g}$$
$$Q_1^{\min} \le Q_1^g \le Q_1^{\max}, \quad Q_2^{\min} \le Q_2^g \le Q_2^{\max} \tag{3.21h}$$

where P_i^g, Q_i^g, $|v_i|$, $\Delta\theta$ are the variables for active power generation, reactive power generation, voltage magnitude at bus i, and angle difference between buses 1 and 2 respectively. Moreover, $P_{i;k}^l$, $Q_{i;k}^l$ are the active and reactive load parameters at bus i and time k, respectively. To simplify our subsequent analysis, we assume that the voltage magnitudes at both buses are equal to the their nominal values, i.e., $|v_1| = |v_2| = 1$. Therefore, according to (3.21b)-(3.21e), the variables $(P_1^g, P_2^g, Q_1^g, Q_2^g)$ can be written in terms of the angle differences $\Delta\theta$. In other words, $P_1^g = p_1(\Delta\theta, P_{1;k}^l)$, $P_2^g = p_2(\Delta\theta, P_{2;k}^l)$, $Q_1^g = q_1(\Delta\theta, Q_{1;k}^l)$,

[5]For simplicity, we omit the apparent power flow limits on the line connecting the two buses. Moreover, to streamline our subsequent analysis, we avoid the index k for the variables.

$Q_2^g = q_2(\Delta\theta)$ where

$$\begin{aligned}
p_1(\Delta\theta, P_{1;k}^l) &= P_{1;k}^l + g + b\sin(\Delta\theta) - g\cos(\Delta\theta)\\
p_2(\Delta\theta, P_{2;k}^l) &= P_{2;k}^l + g + b\sin(\Delta\theta) - g\cos(\Delta\theta)\\
q_1(\Delta\theta, Q_{1;k}^l) &= Q_{1;k}^l + g - g\sin(\Delta\theta) - b\cos(\Delta\theta)\\
q_2(\Delta\theta) &= g - g\sin(\Delta\theta) - b\cos(\Delta\theta)
\end{aligned}$$

Based on these simplifications, the OPF at time k can be re-written as

$$\begin{aligned}
\min\ & f_1(p_1(\Delta\theta, P_{1;k}^l)) + f_2(p_2(\Delta\theta, P_{2;k}^l)) && \text{(3.22a)}\\
\text{s.t. } & P_1^{\min} \le p_1(\Delta\theta, P_{1;k}^l) \le P_1^{\max}, && \text{(3.22b)}\\
& P_2^{\min} \le p_2(\Delta\theta, P_{2;k}^l) \le P_2^{\max} && \text{(3.22c)}\\
& Q_1^{\min} \le q_1(\Delta\theta, Q_{1;k}^l) \le Q_1^{\max}, && \text{(3.22d)}\\
& Q_2^{\min} \le q_2(\Delta\theta, Q_{2;k}^l) \le Q_2^{\max} && \text{(3.22e)}
\end{aligned}$$

Moreover, suppose that the upper and lower bounds on the active and reactive power generations are chosen such that all inequality constraints in (3.22) remain inactive, except for lower bound on the reactive power generation, i.e., $Q_1^{\min} \le q_1(\Delta\theta, Q_{1;k}^l)$. Similar to (3.1), we convert (3.22) to an unconstrained optimization by removing this constraint, and instead, penalizing its violation in the objective function. Based on these modifications, we arrive at the following nonconvex and unconstrained optimization problem:

$$\min_{\Delta\theta}\ \Gamma_k(\Delta\theta) = f_1\left(p_1(\Delta\theta, P_{1;k}^l)\right) + f_2\left(p_2(\Delta\theta, P_{2;k}^l)\right) + \beta\left(\left[Q_{\min} - q_1(\Delta\theta, Q_{1;k}^l)\right]^+\right)^2 \quad (3.23)$$

Suppose that $g - \mathrm{i}b = 0.01 - \mathrm{i}0.1$ and $Q_{\min} = -0.181$. Moreover, suppose that $f_1(P_1^g) = 2(P_1^g)^2 + 2P_1^g + 1$ and $f_2(P_2^g) = 0.1(P_2^g)^2 + 0.1P_2^g + 1$. Finally, the penalization parameter β is set to 500. Figure 3.4.2a illustrates the objective function at the final time T as a function of $\Delta\theta$ for the choices of $P_{1;T}^l = P_{2;T}^l = 0.5$, and $Q_{1;T}^l = Q_{2;T}^l = 0$. Note that the objective function has one global minimum, one strict local minimum, and one local maximum within the interval $-2 \le \Delta\theta \le 1.5$.

Next, we illustrate the effect of load variation on the landscape of this optimization problem and verify Assumption 2. We empirically compute the function $\mathbb{E}_{\omega\sim W(\Delta\theta)}\left[\Gamma_T(\Delta\theta - \eta\omega)\right]$ introduced in Assumption 3.2 when the active and reactive loads are chosen according to the following rules:

- $P_{1;k}^l$ and $P_{2;k}^l$ are chosen uniformly at random from the interval $[0.005, 0.55]$.
- $Q_{2;k}^l = 0$ and $Q_{1;k}^l$ is chosen uniformly at random from the interval $[-0.02, 0.18]$.

Setting $\eta = 2$, for every $k = 0, 1, \ldots, N = 10,000$ we randomly generate the active and reactive load values based on the aforementioned rules, and compute $\Gamma_k(\Delta\theta)$ and $\nabla\Gamma_k(\Delta\theta)$.

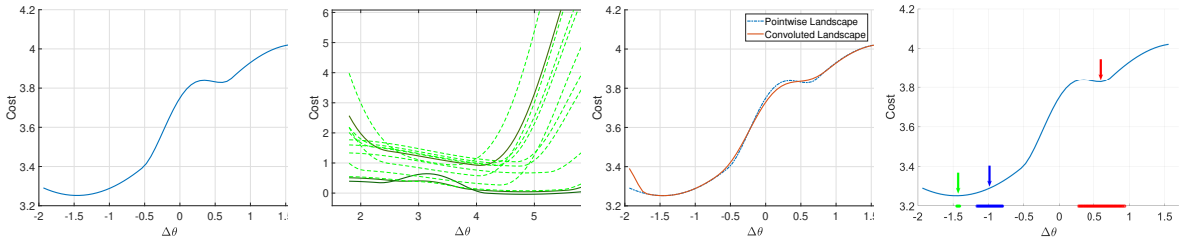

(a) (b) (c) (d)

Figure 3.4.2: (a) The objective function at $t = T$, (b) instances of the objective function for different values of the load, (c) the convoluted and pointwise objective functions, (d) realizations of $\Delta\theta - \eta\omega$ showing the effective noise of the load variation at different $\Delta\theta$ points.

Figure 3.4.2b shows realizations of $\Gamma_k(\Delta\theta)$ for different values of k. Then, for every $k = 0, 1, \ldots, N-1$, we compute the gradient difference $\nabla\Gamma_k(\Delta\theta) - \nabla\Gamma_{k+1}(\Delta\theta)$, capturing the effects of the bias $\zeta_k(\Delta\theta)$ and the effective noise $w_k \sim W(\Delta\theta)$. Since the load distribution is the same at every time, we have $\mathbb{E}[\Gamma_k(\Delta\theta)] = \mathbb{E}[\Gamma_{k+1}(\Delta\theta)]$. Hence $\zeta_k(\Delta\theta) = 0$ for every k. Finally, we approximate $\mathbb{E}_{\omega\sim W(\Delta\theta)}[\Gamma_T(\Delta\theta - \eta\omega)]$ with its empirical counterpart $\frac{1}{N}\sum_{k=0}^{N-1}\Gamma_T(\Delta\theta - \eta\omega_k(\Delta\theta))$.[6] The resulting function for $-2 \le \Delta\theta \le 1.5$ is depicted in Figure 3.4.2c. It can be seen that, unlike $\Gamma_T(\Delta\theta)$, the convoluted function is devoid of spurious local minimum. In fact, it is one-point strongly convex, thereby verifying Assumption 3.2 on the implicit convexity of the convoluted objective function.

3.4.3 The Effect of Expected Gradient on The Noise Variance

Another interpretation of the smoothing effect of the noise is based on the average behavior of the objective function. The *expected objective function* takes the expectation directly over the load's randomness, whereas the convoluted objective function's expectation is taken over a random perturbation of the variables. Because the demand distribution has a direct and immediate physical interpretation, while the effective noise does not, the expected objective function and its gradient are easier to compute or approximate. We will show that the variance of the effective noise $\mathbb{E}_{\omega\sim W(\Delta\theta)}[\|w\|^2]$ at a given point $\widehat{\Delta\theta}$ depends on the gradient of the expected objective function. In other words, a large gradient of the expected objective function at $\widehat{\Delta\theta}$ leads to a high variance $\mathbb{E}_{\omega\sim W(\widehat{\Delta\theta})}[\|w\|^2]$, which in turn yields a smoother $\mathbb{E}_{\omega\sim W(\widehat{\Delta\theta})}\left[\Gamma_T(\widehat{\Delta\theta} - \eta\omega)\right]$. Figure 3.4.2d precisely shows this behavior. In particular, the local minimum $\Delta\theta = 0.6$ of $\Gamma_T(\Delta\theta)$ disappears in $\mathbb{E}_{\omega\sim W(\Delta\theta)}[\Gamma_T(\Delta\theta - \eta\omega)]$ due to the high variance of the additive noise ω at $\Delta\theta = 0.6$ (shown with red circles). On the other hand,

[6] Note that, due to the law of large numbers, the empirical average converges to the expected value as N tends to infinity.

the additive noise at the global minimum $\Delta\theta = -1.4$ is infinitesimal due to the fact that the gradient of the average function remains close to zero at $\Delta\theta = -1.4$. We will now formalize this intuition.

To better elucidate the relationship between the effective noise variance and the expected gradient of the objective function, consider an n-bus system with the following properties:

- Every bus i is equipped with a generator.
- The upper and lower bound constraints on the reactive power generations, and the upper bound constraints on the apparent power flows at different lines are inactive.
- The voltage magnitudes are set to their nominal values.

The above assumptions are made to simplify our subsequent presentation. Note that the problem is still highly nonconvex due to the nonconvex power balance equations and the upper and lower bounds on the active power generations. Let $p_{i;k}(\theta) = P_i^g - P_{i;k}^l$ be the net power injection at bus i at time k, where $\theta \in \mathbb{R}^{N-1}$ is a vector collecting the angles at different buses, except for the slack bus. Then the unconstrained objective function can be defined as $\Gamma_k(\theta) = \sum_i^{n_g} c_i(p_{i;k}(\theta) + P_{i;k}^l)$, where $c_i(p_{i;k}(\theta) + P_{i;k}^l)$ is a linear combination of the cost of generation and the penalties on the violation of the lower and upper bound constraints on the active power generation at generator i. Moreover, suppose that $P_{i;k}^l = \bar{P}_i + \gamma_i$, where $\bar{P}$ is a vector collecting the *nominal* loads, and γ is an isotropic random vector with a known distribution $\mathcal{P}$ such that $\mathbb{E}[\gamma_1] = \cdots = \mathbb{E}[\gamma_n] = \bar{\gamma} \neq 0$. In other words, the variations in the load are *biased*. For simplicity of presentation, we abuse the notation and write $\Gamma(\theta; \bar{P} + \gamma^k) = \Gamma_k(\theta)$, where $\gamma^k \sim \mathcal{P}$ is a realization of the randomness in the load at time k. Define the linearization of $\Gamma(\theta; \bar{P} + \gamma)$ around $\bar{P}$ as

$$\Gamma_{\text{lin}}(\theta; \bar{P} + \gamma) = \Gamma(\theta; \bar{P}) + \nabla_P \Gamma(\theta; \bar{P})^\top \gamma \tag{3.24}$$

For small values of γ, the linearized function $\Gamma_{\text{lin}}(\theta; \bar{P} + \gamma)$ is a good approximation of $\Gamma(\theta; \bar{P} + \gamma)$. In particular, under mild conditions on Γ, the Mean Value theorem implies that $\Gamma(\theta; \bar{P} + \gamma) = \Gamma_{\text{lin}}(\theta; \bar{P} + \gamma) + O(\gamma^2)$. Note that while Γ_{lin} is linear in terms of γ, it is potentially nonconvex with respect to θ. Define effective noise of the linearized functions as

$$\omega_{\text{lin}}^k(\theta; \bar{P}, \gamma^k, \gamma^{k-1}) = \nabla_\theta \Gamma(\theta; \bar{P} + \gamma^k) - \nabla_\theta \Gamma(\theta; \bar{P} + \gamma^{k-1}) \tag{3.25}$$

for every $k = 1, \ldots, T$. Again, ω_{lin}^k is an accurate approximation of the true effective noise, provided γ is sufficiently small. Note that the bias term in (3.25) is zero since the right-hand side of (3.25) has zero mean. Moreover, we can drop the time index k, since the distribution of $\omega_{\text{lin}}^k(\theta; \bar{P}, \gamma^k, \gamma^{k-1})$ does not depend on k, as γ^k and γ^{k-1} are independent and identically distributed. With these definitions, we present our next proposition.

Proposition 3.1. *Suppose that at time T, the objective function of the time-varying OPF corresponds to $\Gamma(\theta; \bar{P})$ with an stationary point $\tilde{\theta}$. Then,*

$$\mathbb{E}_{\gamma,\tilde{\gamma}\sim\mathcal{P}}\left[\|\omega_{\text{lin}}(\tilde{\theta}; \bar{P}, \gamma, \tilde{\gamma})\|^2\right] \geq \frac{2\text{Var}_{\gamma\sim\mathcal{P}}(\gamma)}{N}\left(\frac{\left\|\mathbb{E}_{\gamma\sim\mathcal{P}}\left[\nabla_\theta \Gamma_{\text{lin}}(\tilde{\theta}; \bar{P}+\gamma)\right]\right\|}{\bar{\gamma}}\right)^2 \tag{3.26}$$

The proof of this proposition is deferred to appendix 3.B. Note that a larger variance of the effective noise leads to a higher smoothing effect, which in turn facilitates the satisfaction of Assumption 2. In essence, Proposition 3.1 implies that this smoothing effect (captured by the variance of the effective noise) is controlled by the average behavior of the objective function. In particular, suppose that the point $\tilde{\theta}$ is *not* a stationary point of the expected objective function. Therefore, we have $\left\|\mathbb{E}_{\gamma\sim\mathcal{P}}\left[\nabla_\theta \Gamma_{\text{lin}}(\tilde{\theta}; \bar{P}+\gamma)\right]\right\| > 0$, and the above proposition implies that the *generalized variance* of the effective noise at $\tilde{\theta}$ increases with the norm of the gradient of the expected function at $\tilde{\theta}$, thereby leading to a higher smoothing effect of the load variation and the elimination of the spurious local minima. This partly explains the high variance of the effective noise at the local minimum of the objective function for the 2-bus system described in Subsection 3.4.2, and the elimination of its spurious local minimum.

Based on our results, it is possible to eliminate the spurious local solutions in a point-wise OPF problem by adding a synthetically generated noise to the load, thereby elevating the data variation in the problem. This effect of random perturbation in the load values can be observed in Fig. 3.4.2c, where it is shown that randomness in the load can eliminate the spurious local minimum and maximum, while keeping the global minimum intact.

However, in practice deriving a class of variation sequences which guarantee convergence is not tractable, due to the nonconvex relationship between the load variation and the "effective noise". This is not surprising, considering the NP-hardness of the time-varying OPF in its worst case. However, even without such a guarantee, computing a discrete OPF trajectory for a load sequence which starts and ends with the load of the target problem may often succeed for a straightforward choice of load variation such as a sinusoidal function, uniform variation, or random walk. Fig. 3.4.3 shows two examples on the modified 39-bus system of scaling load by a uniform random walk for 100 time steps with a 20% ramping constraint. As with the sinusoidal load in Fig. 3.3.1, we observe that some or all initializations lead to the optimal solution over time, depending on the specific variation.

3.5 Conclusion

This paper studies time-varying optimal power flow (OPF) problems, in which a set of optimization problems are solved sequentially due to load data variation over time. The solution to each OPF is obtained using local search initialized at the solution to the previous OPF. We offer a case study on a 39-bus system under California data, where the OPF at

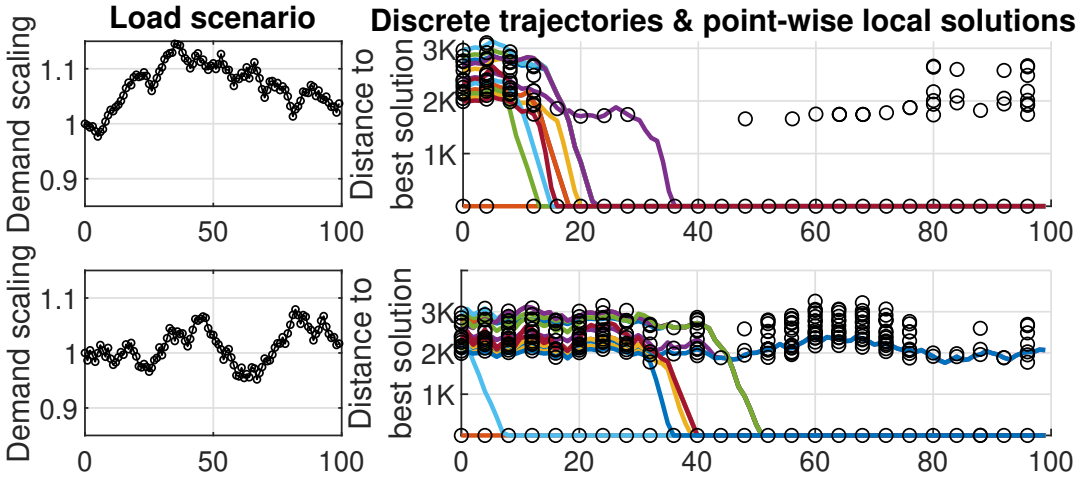

Figure 3.4.3: Numerical studies on the modified 39-bus system with a synthetic load created by a uniform random walk. This type of noise injection could be used to help find the global solution to a static OPF problem.

the initial time has 16 locally optimal solutions leading to 16 solution trajectories. We show that, in this experiment, all trajectories converge to the best solution trajectory, even though OPF has many local minima at most times.

To understand this highly desirable property, we analyze the optimization landscape of the time-varying OPF. Our developed theory is based on the underlying structure of time-varying OPF problems. Despite the generality of our theoretical results, its application relies on assumptions which may not be satisfied for all power systems. Developing a more interpretable set of conditions for our developed theory is left as an important direction of future research. In this work, we introduce the notions of escaping period and tracking period, examine the role of data variation and the easiest intermediate problem, study the behavior of the time-varying OPF during the escaping period via a backward-in-time mapping, and relate it to the SGD algorithm. By modeling the data variation as a biased noise, we prove that enough data variation enables escaping poor solutions of time-varying OPF over time.

Appendix

3.A Proof of Theorem 3.1

For simplicity of notation, we reverse the order of time, changing $T-t$ to t. Then, the dynamics (3.13) can be written

$$\mathbf{z}_t = \mathbf{z}_{t-1} + \eta\nabla\Gamma_0(\mathbf{z}_{t-1}) + \eta\sum_{k=1}^{t-1}\zeta_k(\mathbf{z}_{t-1}) - \eta\sum_{k=1}^{t-1}\omega_k \tag{3.27}$$

We will extensively use the following sequences of *intermediate* points in our analysis:

$$\mathbf{y}_t = \mathbf{z}_t + \eta\nabla\Gamma_0(\mathbf{z}_t) + \eta\sum\nolimits_{k=1}^{t}\zeta_k(\mathbf{z}_t) \tag{3.28}$$

$$\tilde{\mathbf{y}}_t = \mathbf{y}_t - \eta\sum\nolimits_{k=1}^{t-1}\omega_k \tag{3.29}$$

It is easy to verify that the above definitions together with (3.27) gives rise to the following recursive equation:

$$\mathbf{y}_t = \mathbf{y}_{t-1} - \eta\sum_{k=1}^{t-1}\omega_k + \eta\nabla\Gamma_0\left(\mathbf{y}_{t-1} - \eta\sum_{k=1}^{t-1}\omega_k\right) + \eta\sum\nolimits_{k=1}^{t}\zeta_k\left(\mathbf{y}_{t-1} - \eta\sum_{k=1}^{t-1}\omega_k\right) \tag{3.30}$$

which in turn implies

$$\mathbf{y}_t = \tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1} + \eta\nabla\Gamma_0\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right) + \eta\sum_{k=1}^{t}\zeta_k\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right) \tag{3.31}$$

Define the filtration $\mathcal{F}_{t-1} = \sigma\{\omega_1, \ldots, \omega_{t-2}\}$ and the following stochastic process:

$$G_t = (1+\lambda)^{-t}\left(\|\mathbf{y}_t - \mathbf{z}^*\|^2 - \frac{2(b_1 + b_2 t + b_3 t^2)}{\lambda}\right) \tag{3.32}$$

where $b_1 := 2\eta^3 r^2 L$, $b_2 := 2\eta^3 r^2 l$, and $b_3 := \frac{4\eta^2 r^2 (1+2\lambda)^2}{\lambda}$. Our next lemma provides a lower bound on $\mathbb{E}[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}]$ in terms of $\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2$.

Lemma 3.1. *The following inequality holds:*

$$\mathbb{E}\left[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}\right] \geq (1+\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - (b_1 + b_2 t + b_3 t^2) \tag{3.33}$$

Proof. Based on (3.31), one can write

$$\begin{aligned}
&\mathbb{E}\left[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}\right] \\
&= \mathbb{E}\left[\left\|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^* - \eta\omega_{t-1} + \eta\nabla\Gamma_0\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right) + \eta\sum_{k=1}^{t}\zeta_k\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right)\right\|^2\Big|\mathcal{F}_{t-1}\right] \\
&\geq \|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^*\|^2 + \eta^2\mathbb{E}\left[\|\omega_{t-1}\|^2|\mathcal{F}_{t-1}\right] \\
&+ \mathbb{E}\left[\|\eta\nabla\Gamma_0\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right) + \eta\sum_{k=1}^{t}\zeta_k\left(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}\right)\|^2|\mathcal{F}_{t-1}\right] \\
&- 2\eta\left(\mathbb{E}\left[\langle\eta\omega_{t-1}, \nabla\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})\rangle|\mathcal{F}_{t-1}\right] + \mathbb{E}\left[\langle\eta\omega_{t-1}, \sum_{k=1}^{t}\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})\rangle|\mathcal{F}_{t-1}\right]\right) \\
&+2\eta\Big\langle\mathbf{z}^* - \tilde{\mathbf{y}}_{t-1}, -\nabla\mathbb{E}[\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})|\mathcal{F}_{t-1}]\Big\rangle \\
&+ 2\eta\Big\langle\mathbf{z}^* - \tilde{\mathbf{y}}_{t-1}, -\sum_{k=1}^{t}\mathbb{E}\left[\zeta_k(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})|\mathcal{F}_{t-1}\right]\Big\rangle \\
&\geq \|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^*\|^2 \underbrace{-2\eta\mathbb{E}[\langle\eta\omega_{t-1}, \nabla\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})\rangle|\mathcal{F}_{t-1}]}_{A} \\
&\underbrace{-2\eta\mathbb{E}\left[\langle\eta\omega_{t-1}, \sum_{k=1}^{t}\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})\rangle|\mathcal{F}_{t-1}\right]}_{B} + \underbrace{2\eta\langle\mathbf{z}^* - \tilde{\mathbf{y}}_{t-1}, -\nabla\mathbb{E}[\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})|\mathcal{F}_{t-1}]\rangle}_{C} \\
&\underbrace{+2\eta\langle\mathbf{z}^* - \tilde{\mathbf{y}}_{t-1}, -\sum_{k=1}^{t}\mathbb{E}[\zeta_k(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1})|\mathcal{F}_{t-1}]\rangle}_{D}
\end{aligned} \tag{3.34}$$

Next, we will provide a separate lower bound for each term in the above inequality. First, due to Assumption 3, we have

$$C \geq 2\eta c\|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^*\|^2 \text{ and } D \geq -2\eta c'\|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^*\|^2 \tag{3.35}$$

Furthermore, one can write

$$\begin{aligned}
A &= -2\eta\mathbb{E}[\langle\eta\omega_{t-1}, \nabla\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}) - \nabla\Gamma_0(\tilde{\mathbf{y}}_{t-1})\rangle|\mathcal{F}_{t-1}] \\
&\geq -2\eta\mathbb{E}[\|\eta\omega_{t-1}\|\|\Gamma_0(\tilde{\mathbf{y}}_{t-1} - \eta\omega_{t-1}) - \nabla\Gamma_0(\tilde{\mathbf{y}}_{t-1})\||\mathcal{F}_{t-1}] \\
&\geq -2\eta^3 r^2 L
\end{aligned} \tag{3.36}$$

where the first equality is due to the fact that $\mathbb{E}[\langle \eta\omega_{t-1}, \Gamma_0(\tilde{\mathbf{y}}_{t-1})\rangle|\mathcal{F}_{t-1}] = 0$. Similarly, we can write $B \geq -2\eta^3 r^2 lt$. This implies that

$$\begin{aligned}\mathbb{E}[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}] \geq& (1 + 2\eta(c - c'))\|\tilde{\mathbf{y}}_{t-1} - \mathbf{z}^*\|^2 - 2\eta^3 r^2(L + lt)\\ =&(1+2\lambda)\|\tilde{\mathbf{y}}_{t-1}-\mathbf{z}^*\|^2-(b_1+b_2t)\end{aligned} \tag{3.37}$$

This together with the definition of $\tilde{\mathbf{y}}_{t-1}$ gives rise to the following chain of inequalities

$$\begin{aligned}\mathbb{E}[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}] \geq&(1 + 2\lambda)\left\|\mathbf{y}_{t-1} - \eta\sum_{k=1}^{t-2}\omega_k - \mathbf{z}^*\right\|^2 - (b_1 + b_2t)\\ \geq&(1 + 2\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - 2(1 + 2\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|\left\|\eta\sum_{k=1}^{t-2}\omega_k\right\| - (b_1 + b_2t)\\ \geq&(1 + 2\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - 2\eta r(1+2\lambda)t\|\mathbf{y}_{t-1} - \mathbf{z}^*\| - (b_1 + b_2t)\end{aligned} \tag{3.38}$$

Now we consider two cases:

- If $\|\mathbf{y}_{t-1} - \mathbf{z}^*\| \geq \frac{2\eta r(1+2\lambda)t}{\lambda}$, then one can write

$$\mathbb{E}[\|\mathbf{y}_t-\mathbf{z}^*\|^2|\mathcal{F}_{t-1}]\geq(1+\lambda)\|\mathbf{y}_{t-1}-\mathbf{z}^*\|^2-(b_1+b_2t) \tag{3.39}$$

- If $\|\mathbf{y}_{t-1} - \mathbf{z}^*\| < \frac{2\eta r(1+2\lambda)t}{\lambda}$, then one can write

$$\mathbb{E}[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}] \geq (1 + 2\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - \frac{4\eta^2 r^2(1 + 2\lambda)^2t^2}{\lambda} - (b_1 + b_2t) \tag{3.40}$$

Combining the above two inequalities leads to

$$\mathbb{E}[\|\mathbf{y}_t-\mathbf{z}^*\|^2|\mathcal{F}_{t-1}] \geq (1+\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - (b_1+b_2t+ b_3t^2) \tag{3.41}$$

□

The next lemma is at the crux of our proof for Theorem 3.1.

Lemma 3.2. *The following two statements hold:*

i. G_t *is a submartingale with a vanishing drift. More precisely, it satisfies the following inequality*

$$\mathbb{E}[G_t|\mathcal{F}_{t-1}]\geq G_{t-1}-\frac{1}{(1+\lambda)^{(t-1)}}\left(\frac{2b_2+2b_3(2t-1)}{\lambda}\right) \tag{3.42}$$

ii. $\mathbb{E}[G_t] \geq G_0 - \left(\frac{2}{\lambda} + \frac{2}{\lambda^2}\right) b_2 - \left(\frac{4}{\lambda}\left(1 + \frac{1}{\lambda}\right)^2\right) b_3$

Proof. One can write

$$\mathbb{E}[G_t|\mathcal{F}_{t-1}] = (1+\lambda)^{-t} \cdot \left(\mathbb{E}[\|\mathbf{y}_t - \mathbf{z}^*\|^2|\mathcal{F}_{t-1}] - \frac{2(b_1+b_2t+b_3t^2)}{\lambda}\right) \tag{3.43}$$

Invoking Lemma 3.1 leads to

$$\begin{aligned}
\mathbb{E}[G_t|\mathcal{F}_{t-1}] \geq & (1+\lambda)^{-t}\Big((1+\lambda)\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - (b_1+b_2t+b_3t^2) - \frac{2(b_1+b_2t+b_3t^2)}{\lambda}\Big) \\
= & (1+\lambda)^{-(t-1)}\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - (1+\lambda)^{-(t-1)}\left(\frac{2(b_1+b_2t+b_3t^2)}{\lambda}\right) \\
= & (1+\lambda)^{-(t-1)}\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2 - (1+\lambda)^{-(t-1)}\left(\frac{2(b_1+b_2(t-1)+b_3(t-1)^2)}{\lambda}\right) \\
& - (1+\lambda)^{-(t-1)}\left(\frac{2(b_2+b_3(2t-1))}{\lambda}\right) \\
= & G_{t-1} - (1+\lambda)^{-(t-1)}\left(\frac{2b_2+2b_3(2t-1)}{\lambda}\right)
\end{aligned} \tag{3.44}$$

This completes the proof of the first part. To prove the second part, we use the result of the first part together with the tower property of the expectation to write

$$\mathbb{E}[G_t] \geq G_0 - \underbrace{\left(\frac{2b_2}{\lambda}\sum\nolimits_{k=0}^{t-1}(1+\lambda)^{-k}\right)}_{A} - \underbrace{\left(\frac{4b_3}{\lambda}\sum\nolimits_{k=0}^{t-1}(k+1)(1+\lambda)^{-k}\right)}_{B}$$

It is easy to verify that

$$A \leq \left(\frac{2}{\lambda}+\frac{2}{\lambda^2}\right)b_2 \quad \text{and} \quad B \leq \left(\frac{4}{\lambda}\left(1+\frac{1}{\lambda}\right)^2\right)b_3 \tag{3.45}$$

This completes the proof. □

Proof of Theorem 3.1: From the second statement of Lemma 3.2, one can write

$$\|\mathbf{y}_0 - \mathbf{z}^*\|^2 \leq \left(\frac{2}{\lambda}+\frac{2}{\lambda^2}\right)b_2 + \left(\frac{4}{\lambda}\left(1+\frac{1}{\lambda}\right)^2\right)b_3 + (1+\lambda)^{-(t-1)}\,\mathbb{E}\left[\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2\right] \tag{3.46}$$

On the other hand, one can write

$$\begin{aligned}
\mathbb{E}\left[\|\mathbf{z}_t - \mathbf{z}^*\|^2\right] &= \mathbb{E}\left[\left\|\mathbf{y}_{t-1} - \mathbf{z}^* - \eta\sum\nolimits_{k=1}^{t-1}\omega_k\right\|^2\right] \\
&\geq \mathbb{E}\left[\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2\right] - 2\eta rt\mathbb{E}\left[\|\mathbf{y}_{t-1} - \mathbf{z}^*\|\right]
\end{aligned} \tag{3.47}$$

Inequality (3.47) together with some simple algebra reveals that

$$\mathbb{E}\left[\|\mathbf{y}_{t-1} - \mathbf{z}^*\|^2\right] \leq 2\mathbb{E}\left[\|\mathbf{z}_t - \mathbf{z}^*\|^2\right] + 16\eta^2r^2t^2 \tag{3.48}$$

Combining the above inequality with (3.46) results in

$$\|\mathbf{y}_0 - \mathbf{z}^*\|^2 \leq \left(\frac{2}{\lambda} + \frac{2}{\lambda^2}\right) b_2 + \left(\frac{4}{\lambda}\left(1 + \frac{1}{\lambda}\right)^2\right) b_3 + 2\,(1+\lambda)^{-(t-1)}\,\mathbb{E}\left[\|\mathbf{z}_t - \mathbf{z}^*\|^2\right] + 16\eta^2 r^2 t^2 (1+\lambda)^{-(t-1)} \tag{3.49}$$

Finally, it only remains to characterize the relationship between $\|\mathbf{y}_0 - \mathbf{z}^*\|^2$ and $\|\mathbf{z}_0 - \mathbf{z}^*\|^2$. To this goal, one can write

$$\begin{aligned}
\|\mathbf{y}_0 - \mathbf{z}^*\|^2 &= \|\mathbf{z}_0 - \mathbf{z}^* + \eta \nabla f_0(\mathbf{z}_0)\|^2 \\
&\geq \|\mathbf{z}_0 - \mathbf{z}^*\|^2 - 2\eta \langle \mathbf{z}_0 - \mathbf{z}^*, \eta \nabla f_0(\mathbf{z}_0)\rangle \\
&= \|\mathbf{z}_0 - \mathbf{z}^*\|^2 - 2\eta \langle \mathbf{z}_0 - \mathbf{z}^*, \eta \nabla f_0(\mathbf{z}_0) - \eta \nabla f_0(\mathbf{z}^*)\rangle \\
&\geq \|\mathbf{z}_0 - \mathbf{z}^*\|^2 - 2\eta^2 \|\mathbf{z}_0 - \mathbf{z}^*\| \|\nabla f_0(\mathbf{z}_0) - \nabla f_0(\mathbf{z}^*)\| \\
&\geq (1 - 2\eta^2 L)\|\mathbf{z}_0 - \mathbf{z}^*\|^2
\end{aligned} \tag{3.50}$$

where the last inequality is due to Assumption 3.1. Combining (3.50) with (3.49) concludes the proof. □

3.B Proof of Proposition 3.1

Due to the definition of $\Gamma_{\text{lin}}(\theta; \bar{P} + \gamma)$ in (3.24), one can write

$$\begin{aligned}
\mathbb{E}_{\gamma \sim \mathcal{P}}\left[\nabla_\theta \Gamma_{\text{lin}}(\tilde{\theta}; \bar{P} + \gamma)\right] &= \sum_{i=1}^{N} \nabla_{p_i(\theta)} c_i(p_i(\tilde{\theta}) + \bar{P}_i) \nabla_\theta p_i(\tilde{\theta}) \\
&\quad + \sum_{i=1}^{N} \nabla_{p_i(\theta)} \nabla_{P_i} c_i(p_i(\tilde{\theta}) + \bar{P}_i) \mathbb{E}[\gamma_i] \nabla_\theta p_i(\tilde{\theta}) \\
&= \sum_{i=1}^{N} \nabla_{p_i(\theta)} \nabla_{P_i} c_i(p_i(\tilde{\theta}) + \bar{P}_i) \mathbb{E}[\gamma_i] \nabla_\theta p_i(\tilde{\theta})
\end{aligned} \tag{3.51}$$

where the second equality follows from the assumption $\nabla_\theta \Gamma_{\text{lin}}(\tilde{\theta}; \bar{P}) = 0$. Let us define the vector $v_i = \nabla_{p_i(\theta)} \nabla_{P_i} c_i(p_i(\tilde{\theta}) + \bar{P}_i) \nabla_\theta p_i(\tilde{\theta})$. Therefore, one can write

$$\left(\sum_{i=1}^{N} \|v_i\|\right)^2 \geq \left\|\sum_{i=1}^{N} v_i\right\|^2 = \frac{\left\|\mathbb{E}_{\gamma \sim \mathcal{P}}\left[\nabla_\theta \Gamma_{\text{lin}}(\tilde{\theta}; \bar{P} + \gamma)\right]\right\|^2}{\bar{\gamma}^2} \tag{3.52}$$

On the other hand, a simple calculation reveals that

$$\begin{aligned}
\omega_{\text{lin}}^k\left(\theta; \bar{P}, \gamma, \tilde{\gamma}\right) &= \nabla_\theta \Gamma(\theta; \bar{P} + \gamma) - \nabla_\theta \Gamma(\theta; \bar{P} + \tilde{\gamma}) \\
&= \sum_{i=1}^{N} \nabla_{p_i(\theta)} \nabla_{P_i} c_i(p_i(\tilde{\theta}) + \bar{P}_i)(\gamma_i - \tilde{\gamma}_i) \nabla_\theta p_i(\tilde{\theta})
\end{aligned} \tag{3.53}$$

Upon defining the matrix $V = [v_1, v_2, \ldots, v_N]$, one can verify that $\omega^k_{\text{lin}}(\theta; \bar{P}, \gamma, \tilde{\gamma}) = V(\gamma - \tilde{\gamma})$, which implies that

$$\begin{aligned}\mathbb{E}_{\gamma,\tilde{\gamma}\sim\mathcal{P}}\left[\left\|\omega_{\text{lin}}\left(\tilde{\theta}; \bar{P}, \gamma, \tilde{\gamma}\right)\right\|^2\right] &= \mathbb{E}_{\gamma,\tilde{\gamma}\sim\mathcal{P}}\left[\|V(\gamma - \tilde{\gamma})\|^2\right] \\ &= 2\text{Var}_{\gamma\sim\mathcal{P}}(\gamma)\text{trace}\left(VV^\top\right) \\ &= 2\text{Var}_{\gamma\sim\mathcal{P}}(\gamma)\text{trace}\left(V^\top V\right) \\ &= 2\text{Var}_{\gamma\sim\mathcal{P}}(\gamma)\left(\sum_{i=1}^N \|v_i\|^2\right) \end{aligned} \tag{3.54}$$

This implies that

$$\mathbb{E}_{\gamma,\tilde{\gamma}\sim\mathcal{P}}\left[\left\|\omega_{\text{lin}}\left(\tilde{\theta}; \bar{P}, \gamma, \tilde{\gamma}\right)\right\|^2\right] \geq \frac{2\text{Var}_{\gamma\sim\mathcal{P}}(\gamma)}{N}\left(\sum_{i=1}^N \|v_i\|\right)^2$$

The above inequality combined with (3.52) completes the proof. □

Chapter 4

Hybrid Power Plant Bidding in Models of Future Electricity Systems

4.1 Introduction

In the quest to deeply decarbonize electricity sources, batteries and other energy storage technologies are touted as a means of increasing electricity grid flexibility to balance the variability of renewable power sources, namely solar and wind. One option for deploying an energy storage resource ("ESR"), particularly battery storage, is to install it alongside solar panels and wind turbines. So-called hybrid power plants ("hybrids") – those operating both renewable generation and storage resources at the same site – have drawn interest from developers and policy-makers. At the end of 2021, almost 36 GW of generation and 3.2 GW (8.1 GWh) of storage were operating as hybrids across nearly 300 plants [11]. These represent the first wave of projects of the 285 GW of solar, 17 GW of wind, and 207 GW proposed as hybrids in interconnection queues. The trend toward hybrids is especially strong in California, where 95% and 42% of proposed solar and wind capacity, respectively, are part of hybrid power plants [100] and the independent system operator expects 5 GW and 2.8 GW of generation and storage capacity, respectively, built as hybrids by the end of 2025 [55].

Understanding how hybrids will impact wholesale electricity markets and system reliability is of interest to many stakeholders, including developers, system operators, and regulators. Developers want to understand where markets will be attractive for hybrids in the short, medium, and long term, which market participation option(s) they should purse, and how to operate a hybrid for maximum profits given market and renewable generation uncertainty. System operators need to decide what options hybrids will have for participating in their markets, and they seek information on how the various options will impact prices and system reliability in order to make these decisions. In recent years the Federal Energy Regulatory Commission has identified a need to define participation models for storage (Order 841 [37]), a closely related asset.

Researchers try to gain insight into such stakeholder questions by simulating hypothetical

future electricity markets under different assumptions. Production cost models are a key tool used to evaluate the economic impact of decisions of dispatching a fleet of generators [56]. Building a production cost model requires visibility into generators' marginal operating costs, which define their market bids. For fossil fuel-based generators, operating costs are closely tied to commodity prices; for renewable generators, marginal operating costs are near zero. However, resources with the ability to store energy have an effective marginal cost based on expectations of future prices. This makes it challenging to emulate hybrid operators' views of their marginal costs in any given hour. The primary goal of this paper is to enable researchers to incorporate hybrids that participate in the market through price-quantity bidding curves, similar to those of conventional resources, into production cost models with fidelity.

4.1.1 Participation Models for Hybrid Power Plants

A participation model defines the interface between an electricity market and a market participant, in this case hybrids. It includes the information to be exchanged, such as market bids, generation forecasts, and telemetry, and the responsibilities of each party. The four general hybrid participation model options defined by [46] are shown in Figure 4.1.1. This paper exclusively focuses on model (b) in which the hybrid is a "single, self-managed resource." That is, the hybrid power plant appears as a single combined system to the market and is responsible for managing its dispatch schedule and storage state-of-charge levels using market bids for combined resource. This option provides hybrid operators a great deal of flexibility to maximize the plant's value, limits the market operator's computational burden and places the responsibility for risk on the hybrid participant. A summary of wholesale market participation options for hybrid resources in each ISO and RTO in the United States is available in [54].

4.1.2 Capability Needed to Enable Research Agenda

The single, self-managed resource participation option is the most challenging for researchers to represent in a production cost model, because it requires adopting an asset owner's point of view and integrating the two resources' capabilities. Further, this challenge is distinct from that of operating a hybrid under this model in a market. In the simulation environment, there is limited data available, since it represents a hypothetical futuristic market being simulated for the first time. Actual hybrid participants, on the other hand, would have a long record of historical market prices, the performance of their past bids, and, in some markets, the past bids of other market participants to inform their decisions. The goal is to create bids which are consistent with those a hybrid participant would submit, without this knowledge of historical market outcomes which would inform the hybrid's strategy in practice. Developing this capability is one step toward allowing researchers to analyze and plan future power system market designs in which hybrids represent a large share of the system capacity and to compare the benefits and risks of using a single, self-managed participation model for hybrids to those of using a separate resource or market-managed model.

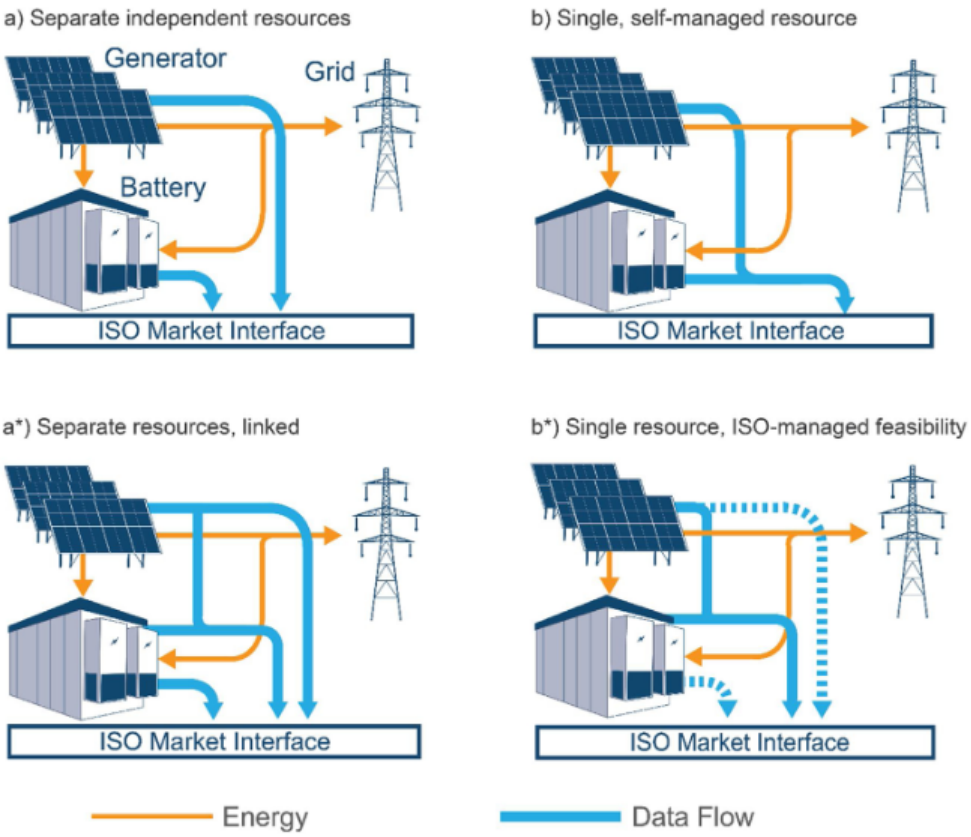

Figure 4.1.1: Spectrum of participation models for hybrid resources, reproduced from [46].

4.1.3 Context

This work funded by the Department of Energy's Office of Energy Efficiency & Renewable Energy and Office of Electricity as part of a project examining how participation models will impact hybrid dispatch, hybrid revenue, system efficiency and system reliability. Polaris' Power System Optimizer was the production cost modeling software used in this project. The subsequent sections will discuss specific requirements, challenges and design decisions made in the context of this project in addition to broadly applicable models, algorithms and analytic techniques.

4.1.4 Contributions

This work offers two main contributions to the challenge of representing hybrids operating as a single, self-managed resource in day-ahead wholesale electricity markets. The first contribution addresses the challenge of decision-making under uncertainty. This challenge is inherent to self-management, as market prices and renewable generator production are unknown in advance, and is one hybrid operators will face today and in the future. It is addressed with a comprehensive stochastic optimization model which reflects the decision-

making process of a hybrid operator in the day-ahead market. This mixed-integer linear programming model is informed by an extensive review of relevant literature and conversations with industry experts. The second contribution addresses the research challenge of incorporating a decision-making model into a futuristic system with limited historical data. This challenge is critical to storage resources, which rely on expectations of market prices in order to arbitrage between high and low price periods. It is addressed with a detailed methodology for incorporating the proposed model into production cost models of new electricity markets. The approach includes the creation of scenarios representing renewable generation and market prices, selection of modeling parameters, and a heuristic for addressing the possibility of price spikes not reflected in the scenarios.

4.2 Hybrid Bidding Problem Formulation

The objective of this section is to design a decision-making model for a hybrid determining its hourly bid curves for a day-ahead market under a single, self-managed resource participation model. The model must reliably produce realistic bid curves based on the type and volume of information available when simulating a new electricity market. Section 4.2.1 states the criteria which the hybrid bidding model must satisfy. Section 4.2.2 evaluates relevant literature on the operation of hybrids, stand-alone storage and related assets and concludes that a single model completely satisfactory for this setting has yet to be proposed. Sections 4.2.4 and 4.2.5 present the proposed model formulation in detail. Finally, section 4.2.6 discusses the model and relates it to those in the literature and section 4.2.7 covers how to extend the base model to a handful of special cases.

4.2.1 Model Requirements

The day-ahead bidding model should base decisions on a complete view of the hybrid's capabilities and the energy market opportunities. Specifically, it must:

1. set hourly price-dependent bid curves for a day-ahead energy market for a profit-seeking PV-battery or wind-battery hybrid in a "single, self-managed resource" participation model,

2. produce bid curves with a valid structure, that is, consist of a limited number of non-decreasing marginal price steps,

3. reflect the technical capabilities and physical limitations of the generation and storage resources and their shared point of interconnection to the grid without imposing unnecessary restrictions,

4. consider day-ahead and intraday markets, as well as consequences for deviating from market dispatch, when making day-ahead decisions,

5. account for uncertainty in renewable generation and market prices,
6. accommodate positive and negative prices,
7. consider the value of storage at the end of the 24-hour bidding horizon,
8. depend only on the type and volume of information available when simulating a new electricity market, and
9. be computationally tractable.

In addition to the above requirements, an ideal model would also:

10. flexibly manage risk according to an individual developer's risk tolerance,
11. easily generalize to other technologies beyond PV-battery and wind-battery, and,

Note that this model is focused on the energy market, and it need not consider ancillary services or capacity markets. However, a straightforward extension of the proposed model to include capacity markets is shared in Section 4.2.7.

4.2.2 Literature Review

Numerous approaches to market participation for hybrids and related resources have been proposed in the literature. Since renewable energy generators typically have near zero marginal cost, a major challenge of hybrid bidding centers on the ESR's unknown opportunity costs. References on stand-alone storage, therefore, can provide insights into addressing this challenge. Table 4.2.1 summarizes key characteristics of recent works on how to offer stand-alone storage in electricity markets. Hydro storage systems are the longest-studied storage technology. When hydro reservoirs have stochastic inflows, the system model is more similar to a hybrid than an ESR. Bidding strategies for hydro-electric producers are surveyed by [109], where the problems classified as medium term are most relevant to this setting. Virtual power plants (VPPs) and microgrids consist of multiple resources (e.g., generation, storage, demand response) that may bid into the market as one, like a hybrid in a single, self-managed resource participation model, though they differ from hybrids in their physical coupling. Table 4.2.2 summarizes key characteristics of recent works on how to offer hybrids, VPPs and microgrids as a single resource in electricity markets, along with select works on hydro storage bidding.

Reviewing Table 4.2.2, there are five approaches which appear to be candidates. That is, approaches which offer price-dependent offer curves instead of self-scheduling, consider at least two sequential energy markets when making day-ahead decisions, and include uncertainty in both day-ahead prices and generation levels. However, none of these approaches completely satisfy the criteria in section 4.2.1. [81] requires a model of the environmental states underlying the system, how these states evolve over time, and the distribution of prices given the current state, which is unrealistic given the limited information available in the

simulation environment considered here (requirement 8). The remaining candidates do not meet requirement 2, since decreasing offer curves may be produced by [99], and neither [80], [41], nor [58] limits the number of price steps in the bid curve. [41] also does not meet requirement 8 because it requires scenarios on the ESR's power generation/consumption, akin to wind generation scenarios, which, in this paper's setting, is not information the bidder is considered to be able to predict in advance of solving this scheduling problem.

Ref.	**Year**	**Bid structure**	**Day-ahead price formation**	**# energy markets**	**Uncertain day-ahead prices**	**Modeling approach**	**Solution method**
[4]	2014	SS	T	2	-	SO, H	LP, CP
[60]	2015	Curve	T	2	✓	SO	Monotone Appx. DP
[117]	2015	SS	T	1	✓	SO, RO	MILP
[52]	2016	SS	T	1	✓	SO	NLP
[78]	2016	SS	M	1	✓	SO	MILP
[89]	2016	Curve	T	2	✓	SO	MILP
[105]	2017	Curve	T	1	✓	IGDT	MINLP
[116]	2017	Curve	M	1	✓	Bi-level SO	MILP
[14]	2018	Curve	T	2	-	H	-
[21]	2018	SS	T	2	✓	SO	DP
[25]	2018	SS	M	1	-	NE	CP
[69]	2018	Curve	T	2	✓	SO	MILP
[93]	2018	Curve	M	1	-	Bi-level SO	MILP
[114]	2018	SS	M	1	-	DO	QP, MILP
[3]	2019	Curve	T	1	✓	SO	MILP
[24]	2019	SS	T	1	-	DO	LP
[71]	2019	SS	T	1	✓	SO	LP
[87]	2020	SS	T	2	-	DO	MILP
[120]	2020	SS	T	1	✓	SO	H, Stoch. Dual DP
[74]	2021	Curve	M	2	-	Bi-level DO	MILP

Abbreviations: SS=self-schedule, T=price-taker, M=price-maker, H=heuristic, (D/R/S)O=(deterministic/robust/stochastic) optimization, NE=Nash equilibrium, IGDT=information gap decision theory, (C/D/MI/N/L/Q)P=(convex/dynamic/mixed-integer/non-/linear/quadratic) programming

Table 4.2.1: Overview of recent works on stand-alone storage bidding or scheduling

Ref.	Year	Bid structure	Day-ahead price formation	# energy markets	Uncertain day-ahead prices	Uncertain generation	Modeling approach	Solution method
[67]	2003	SS	T	1	-	-	DO	DP
[40]	2008	SS	T	2	✓	✓	SO	MILP
[28]	2012	SS	T	1	-	-	H	-
[85]	2012	SS	T	1	-	-	DO	DP
[81]	2013	Curve	T	2	✓	✓	SO	ADDP
[104]	2013	SS	T	2	✓	✓	SO	MILP
[113]	2013	SS	T	1	✓	✓	RO	LP
[80]	2015	Curve	T	2	✓	✓	SO, RO	MILP
[82]	2015	SS	T	1	-	✓	H	-
[98]	2015	SS	T	1	✓	✓	RO	MILP
[99]	2016	Curve	T	2	✓	✓	RO	LP
[104]	2016	SS	T	2	✓	✓	SO	MILP
[29]	2017	Curve	M	2	-	✓	SO	MILP
[43]	2017	SS	T	2	✓	✓	SO	MILP
[94]	2017	SS	T	2	-	-	DO	LP
[63]	2018	SS	T	2	✓	✓	SE	MLCP
[1]	2018	SS	T	1	-	✓	SO	DP
[7]	2018	SS	T	2	✓	✓	RO	MILP
[64]	2018	SS	T	1	-	-	DO	DP
[18]	2019	SS	T	2	-	✓	SO	NLP
[44]	2019	SS	T	3	-	✓	DO	LP
[2]	2020	SS	T	2	-	✓	SO	NLP
[27]	2020	SS	T	2	-	-	DO	MILP
[41]	2020	Curve	T	3	✓	✓	SO	MILP
[84]	2020	SS	T	1	-	✓	H	-
[121]	2020	SS	T	1	-	✓	DRO	SOCP
[125]	2020	SS	T	1	✓	✓	SO	MILP
[58]	2020	Curve	T	3	✓	✓	SO	MILP
[88]	2021	SS	T	1	-	✓	SO	MILP
[119]	2021	SS	T	1	✓	✓	RO	MILP

Abbreviations: SS=self-schedule, T=price-taker, M=price-maker, H=heuristic,
(D/R/S)O=(deterministic/robust/stochastic) optimization, SE=Stochastic equilibrium,
DRO=distributionally RO, (AD)DP=(approximate dual) dynamic programming
(MI/N)LP=(mixed-integer/non-)linear programming
MLCP=mixed linear complementarity problem, SOCP=second-order conic programming

Table 4.2.2: Overview of recent works on hybrid bidding or scheduling

4.2.3 Notation

In the following sections, a hybrid bidding model is formulated to satisfy the criteria listed in Section 4.2.1. The model will make use of the following notation:
Subscripts

$b \in \mathcal{B} = \{1, 2, \ldots, N_B\}$	Index of day-ahead bidding curve breakpoint.
$s \in \mathcal{S}$	Index of scenario, where a scenario is the set of parameters defined in (4.1).
$t \in \mathcal{T} = \{1, 2, \ldots, N_T\}$	Index of time.

For convenience, a given time step t and the time period from 1 to N_T will at times be referred to as an "hour" and "day", respectively, however N_T is not required to be 24 and the unit t represents can be arbitrary.
Superscripts

cha	Index for charging mode of the ESR.
dis	Index for discharging mode of the ESR.
DA	Index for day-ahead market.
IN	Index for intraday market.
SC	Index for schedule across all markets.
DV	Index for deviations from schedule across all markets.
min	Index specifying the minimum allowable value for a variable.
max	Index specifying the maximum allowable value for a variable.

Variables (also occasionally used as superscripts)

C	State of charge of the ESR.
E	Power produced by the ESR.
G	Power produced by the generator.
$\bar{G}$	Potential power from the generator that is curtailed instead of produced.
H	Power produced by the hybrid.
$MODE$	Indicator of ESR operating status.
$\epsilon^{+/-}$	Deviation of the hybrid above/below the scheduled power output.
ϵ	Total deviation of the hybrid from the scheduled power output.
θ	Supplementary variable to calculate CVaR.
ϕ	Supplementary variable to calculate CVaR.
x	Volume component of the (y, x) price-volume pairs which define the hybrid's day-ahead market bidding curve.

Parameters

$N_S = \lvert\mathcal{S}\rvert$	Number of scenarios.
$N_T = \lvert\mathcal{T}\rvert$	Number of time periods.
$N_B = \lvert\mathcal{B}\rvert$	Number of bidding curve breakpoints.
gen	Maximum power output for the generator, given environmental conditions.
π_s	Probability of scenario s occurring.
$\eta^{+/-}$	Ratio for calculating positive/negative deviation prices.
poi	Point of interconnection power capacity.
$\psi \in \{1, 0\}$	Indicator of whether or not the ESR may charge from the grid.
ρ^{DA}	Price of energy in the day-ahead electricity market.
ρ^{ceil}	Price ceiling in the day-ahead electricity market.
ρ^{IN}	Difference between the price of energy in the intraday and day-ahead markets.
$\rho^{+/-}$	Price of positive/negative deviations from the combined market schedule.
ζ	Risk-aversion factor.
$\sigma \in (0, 1)$	Confidence level.
Λ	Factor limiting energy volume in the intraday, relative to day-ahead, market.
$\mu \in (0, 1)$	Efficiency of the ESR.
C_{start}	Initial ESR state of charge.
O	Operating cost per unit of energy.
y	Price component of the (y, x) price-volume pairs which define the hybrid's day-ahead market bidding curve.

4.2.4 Hybrid Power Plant Bidding Model Formulation

The goal of the hybrid operator is to maximize its total operating profits while managing risk due to uncertainty. Uncertainty appears in the model to satisfy requirement 5 as a set of scenarios, where any given scenario s consists of

$$\pi_s \text{ and } \left\{\rho^{DA}_{t,s},\ \rho^{IN}_{t,s},\ gen_{t,s},\ \eta^{+}_{t,s},\ \eta^{-}_{t,s},\ \mu^{cha}_{t,s},\ \mu^{dis}_{t,s},\ C^{max}_{t,s},\ C^{min}_{t,s}\right\}_{t\in\mathcal{T}}. \tag{4.1}$$

Implicit in this scenario definition is the assumption that the installed generation capacity, G^{max}_t, and rated charging and discharging power capacities of the storage, $E^{cha,max}_t$ and $E^{dis,max}_t$, are deterministic. If this is not the case, it is straightforward to index these quantities by s too. The hybrid's objective is formulated as

$$\max_{\Gamma} \sum_{s\,\in\,\mathcal{S}} \left(\pi_s\big(\underbrace{revenue_s - cost_s}_{profit_s}\big)\right) + \zeta\underbrace{\left(\theta - \frac{1}{1-\sigma}\sum_{s\,\in\,\mathcal{S}}\pi_s\phi_s\right)}_{CVaR}, \tag{4.2}$$

where

$$revenue_s := \sum_{t=1}^{N_T} \rho_{t,s}^{DA} H_{t,s}^{DA} + \left(\rho_{t,s}^{DA} + \rho_{t,s}^{IN}\right) H_{t,s}^{IN} + \rho_{t,s}^{+}\epsilon_{t,s}^{+} \tag{4.3}$$

$$cost_s := \sum_{t=1}^{N_T} \rho_{t,s}^{-}\epsilon_{t,s}^{-} + O_t^G G_{t,s}^{SC} + O_t^{E^{SC,cha}} E_{t,s}^{SC,cha} + O_t^{E^{SC,dis}} E_{t,s}^{SC,dis} \tag{4.4}$$

$$\begin{aligned}\Gamma := \Big\{ & \{x_{t,b}\}_{t=1,\, b=1}^{t=N_T,\, b=N_B}, \theta, \{\phi_s\}_{s=1}^{N_S}, \\ & \{H_{t,s}^{DA}, H_{t,s}^{IN}, H_{t,s}^{SC}, G_{t,s}^{DA}, G_{t,s}^{IN}, G_{t,s}^{SC}, E_{t,s}^{DA}, E_{t,s}^{IN}, E_{t,s}^{SC}, E_{t,s}^{SC,dis}, E_{t,s}^{SC,cha}, \\ & MODE_{t,s}^{dis}, MODE_{t,s}^{cha}, C_{t,s}^{SC}, \epsilon_{t,s}^{+}, \epsilon_{t,s}^{-}, \epsilon_{t,s}\}_{t=1,\, s=1}^{t=N_T,\, s=N_S} \Big\}\end{aligned} \tag{4.5}$$

and the deviation prices are defined as follows:

$$\rho_{t,s}^{+} = \begin{cases} \eta_{t,s}^{+} \cdot \min\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}, & \text{if } \min\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\} \geq 0 \\ \eta_{t,s}^{-} \cdot \min\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}, & \text{otherwise} \end{cases} \tag{4.6}$$

$$\rho_{t,s}^{-} = \begin{cases} \eta_{t,s}^{+} \cdot \max\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}, & \text{if } \max\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\} \leq 0 \\ \eta_{t,s}^{-} \cdot \max\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}, & \text{otherwise} \end{cases} \tag{4.7}$$

The objective function (4.2) balances performing well on average with limiting the extent of poor performance in the worst cases. It does so by maximizing the expected profit from participating in all three markets, satisfying requirement 4, plus the weighted conditional value-at-risk (CVaR), which is equal to the expected profit of the least profitable $(1-\sigma) \times 100\%$ of scenarios [102]. The weight ζ placed on the CVaR term and the choice of σ are design parameters which should be selected based on the risk tolerance of the hybrid owner, satisfying requirement 10. In general, it is desirable to limit one's exposure to deviation prices, which are used to settle differences between the delivered power and the schedule determined by the day-ahead and intraday markets. The following lemma states that, given appropriate choices of $\eta_{t,s}^{+}$ and $\eta_{t,s}^{-}$, this model maintains the property that deviations always close against the participant.

Lemma 4.1. *If $\eta_{t,s}^{+} < 1$ and $\eta_{t,s}^{-} > 1$, then the day-ahead and intraday markets are at least as attractive as the deviation prices at time t of scenario s. Specifically, the price for power above the scheduled output is less than then the day-ahead and real-time market prices and the price for power below the scheduled output is greater than the day-ahead and real-time market prices. That is,*

$$\rho_{t,s}^{+} \leq \min\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\} \text{ and} \tag{4.8}$$

$$\rho_{t,s}^{-} \geq \max\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}, \tag{4.9}$$

with equality if and only if the right-hand side is zero.

Proof: This proof will show that the desired inequalities hold for all possible values of $\rho_{t,s}^{DA}$ and $\rho_{t,s}^{IN}$ by considering three main cases and then the edge cases. To simplify the notation, define $\bar{\rho}_{t,s} = \max\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}$ and $\underline{\rho}_{t,s} = \min\{\rho_{t,s}^{DA}, \rho_{t,s}^{DA} + \rho_{t,s}^{IN}\}$.

i. $\underline{\rho}_{t,s} > 0$: In this case, $\rho_{t,s}^{+} = \eta_{t,s}^{+} \cdot \underline{\rho}_{t,s} < \underline{\rho}_{t,s}$ and $\rho_{t,s}^{-} = \eta_{t,s}^{-} \cdot \bar{\rho}_{t,s} > \bar{\rho}_{t,s}$.

ii. $\bar{\rho}_{t,s} > 0$ and $\underline{\rho}_{t,s} < 0$: In this case, $\rho_{t,s}^{+} = \eta_{t,s}^{-} \cdot \underline{\rho}_{t,s} < \underline{\rho}_{t,s}$ and $\rho_{t,s}^{-}$ is as in Case 1.

iii. $\bar{\rho}_{t,s} < 0$: In this case, $\rho_{t,s}^{+}$ is as in Case 2 and $\rho_{t,s}^{-} = \eta_{t,s}^{+} \cdot \bar{\rho}_{t,s} > \bar{\rho}_{t,s}$.

Thus, it is proven that (4.8) and (4.9) hold strictly if $\underline{\rho}_{t,s} \neq 0$ and $\bar{\rho}_{t,s} \neq 0$. Finally, by inspection of (4.6) and (4.7), it is apparent that $\rho_{t,s}^{+} = 0$ if $\underline{\rho}_{t,s} = 0$ and $\rho_{t,s}^{-} = 0$ if $\bar{\rho}_{t,s} = 0$. □

The hybrid operator is subject to constraints imposed by the hybrid's physical infrastructure, market rules, and their risk management strategy. The first set of constraints primarily pertain to the hybrid's physical infrastructure and govern the power output of the hybrid, H, the generator component, G, the storage component, E. These constraints are as follows:

$$\max\{-\psi E_t^{cha,max}, -poi\} \leq H_{t,s}^{SC} \leq poi \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.10}$$
$$H_{t,s}^{SC} = H_{t,s}^{DA} + H_{t,s}^{IN} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.11}$$
$$H_{t,s}^{k} = G_{t,s}^{k} + E_{t,s}^{k} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S},\ \forall k \in \{SC, DA, IN\} \tag{4.12}$$
$$0 \leq G_{t,s}^{SC} \leq G_t^{max} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.13}$$
$$G_{t,s}^{SC} = G_{t,s}^{DA} + G_{t,s}^{IN} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.14}$$
$$-E_t^{cha,max} \leq E_{t,s}^{SC} \leq E_t^{dis,max} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.15}$$
$$E_{t,s}^{SC} = E_{t,s}^{DA} + E_{t,s}^{IN} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.16}$$

The above constraints ensure that, in accordance with requirement 3, for each scenario and each time interval, the schedule for the hybrid and each component is the sum of its activity in the day-ahead and intraday market (4.11, 4.14, 4.16), the physical upper and lower bounds are respected in the combined-market schedule (4.10, 4.13, 4.15), and the hybrid is truly the sum of its parts (4.12).

The second set of constraints also supports requirement 3 by governing the storage component's operation:

$$MODE_{t,s}^{k} \in \{0, 1\} \qquad \forall t \in \mathcal{T}, \forall s \in \mathcal{S}, \forall k \in \{cha,\ dis\} \tag{4.17}$$
$$MODE_{t,s}^{cha} + MODE_{t,s}^{dis} \leq 1 \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.18}$$
$$0 \leq E_{t,s}^{SC,k} \leq E_t^{k,max} \cdot MODE_{t,s}^{k} \qquad \forall t \in \mathcal{T}, \forall s \in \mathcal{S}, \forall k \in \{cha,\ dis\} \tag{4.19}$$
$$E_{t,s}^{SC} = E_{t,s}^{SC,dis} - E_{t,s}^{SC,cha} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.20}$$

$$C^{min}_{t,s} \leq C^{SC}_{t,s} \leq C^{max}_{t,s} \qquad \forall t \in \mathcal{T} \cup \{final\},\ \forall s \in \mathcal{S} \qquad (4.21)$$

$$C^{SC}_{1,s} = C_{start} \qquad \forall s \in \mathcal{S} \qquad (4.22)$$

$$C^{SC}_{t,s} = C^{SC}_{t-1,s} + \mu^{cha}_{t-1,s} \cdot E^{SC,cha}_{t-1,s} - \frac{1}{\mu^{dis}_{t-1,s}} \cdot E^{SC,dis}_{t-1,s} \qquad \forall t = 2, \ldots, N_T,\ \forall s \in \mathcal{S} \qquad (4.23)$$

$$C^{SC}_{final,s} = C^{SC}_{N_T,s} + \mu^{cha}_{N_T,s} \cdot E^{SC,cha}_{N_T,s} - \frac{1}{\mu^{dis}_{N_T,s}} \cdot E^{SC,dis}_{N_T,s} \qquad \forall s \in \mathcal{S} \qquad (4.24)$$

Constraints (4.17) and (4.18) encode the inability of the ESR to simultaneously charge and discharge. Constraints (4.19) and (4.20) decompose the ESR's power output into the power output associated with the discharging mode and the power input associated with the charging mode, at most one of which is non-zero. The storage state-of-charge at the beginning of each time interval is managed by constraints (4.21-4.24). This requires that the model define a realistic operating schedule for the ESR for each scenario. In many applications, the parameters $C^{min}_{t,s}$ and $C^{max}_{t,s}$ may be fixed across time and scenarios, with $C^{min}_{t,s} = 0$ and $C^{max}_{t,s}$ equal to the ESR's energy capacity. However, an operator may wish to set $C^{min}_{t,s}$ above zero in most hours, since operating near zero state-of-charge can be detrimental to a battery's lifespan. Another use of these parameters is to require the state-of-charge at the end of the time horizon to be a specific value or within a range, which can be achieved through the choice of $C^{min}_{final,s}$ and $C^{max}_{final,s}$.

The next set of constraints determines the positive and negative imbalances which will be settled using deviation prices and are defined as follows:

$$\epsilon_{t,s} = gen_{t,s} - G^{SC}_{t,s} - \bar{G}_{t,s} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \qquad (4.25)$$

$$0 \leq \bar{G}_{t,s} \leq gen_{t,s} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \qquad (4.26)$$

$$0 \leq \epsilon^{k}_{t,s} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S},\ \forall k \in \{+, -\} \qquad (4.27)$$

$$\epsilon_{t,s} = \epsilon^{+}_{t,s} - \epsilon^{-}_{t,s} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \qquad (4.28)$$

The total deviation for a specific time interval and scenario is the difference between the generator's output assuming no curtailment occurs, the scheduled generator contribution, and the amount of curtailed power, as specified in constraints (4.25, 4.26). This total deviation is decomposed into the magnitude of the positive or negative deviation (corresponding to injecting more or less power than scheduled, respectively) in constraints (4.27, 4.28), so that the appropriate prices can be applied in the objective function.

While market rules vary, the following constraints reflect universal market principles in support of requirements 1 and 2:

$$H_{t,s}^{DA} = x_{t,b} \quad \text{if} \quad y_{t,b-1} < \rho_{t,s}^{DA} \leq y_{t,b} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.29}$$

$$x_{t,b} \leq x_{t,b+1} \qquad \forall t \in \mathcal{T},\ \forall b = 1, \ldots, N_B - 1 \tag{4.30}$$

$$\max\{-\psi E_t^{cha,max}, -poi\} \leq H_{t,s}^{DA} \leq \min\{E_t^{dis,max} + G_t^{max}, poi\} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.31}$$

$$-\Lambda G_{t,s}^{DA} \leq G_{t,s}^{IN} \leq \Lambda G_{t,s}^{DA} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.32}$$

$$-\Lambda E_t^{cha,max} \leq E_{t,s}^{IN} \leq \Lambda E_t^{dis,max} \qquad \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{4.33}$$

Constraint (4.29) ensures that there are at most N_B points on the bid curve for each time interval, and scenarios are assigned to a bid curve segment based on their day-ahead price for that time. The *if* statement in this constraint does *not* affect its linearity, because the price components of the bidding pairs, $\{y_{tb}\}_{t\in\mathcal{T},b\in\mathcal{B}}$, are not variables in this optimization problem. These price points are fixed in advanced based on the scenarios, through a procedure explained in section 4.2.5, since determining prices and volumes simultaneously would create a nonconvex decision problem [81]. The requirement that the bids be non-decreasing and non-anticipative (requirement 2) is enforced by constraint (4.30). (More details on the bid curve structure are found in section 4.2.5. Typically, market participants (who are not virtual bidders) must submit bids which do not exceed their physical operating limits. For a hybrid, this means that day-ahead bid quantities must not exceed the point of interconnection limits, as well as fall between the negative of the ESR's rated charging power capacity and the combined rated power capacity of the ESR and generator (4.31). Constraints (4.32) and (4.33) limit the degree of participation in the intraday market. This may be enforced by a market operator to ensure that the bulk of transactions are settled in the day-ahead market and to limit bids seeking to arbitrage between the day-ahead and intraday markets. It could also be imposed by the hybrid operator to limit their exposure to the more volatile intraday market in a coarse way. If this is not relevant to a market or the hybrid's strategy, these two constraints must be dropped.

The final two constraints, which follow, are necessary to calculate the CVaR:

$$0 \leq \phi_s \qquad \forall s \in \mathcal{S} \tag{4.34}$$

$$\theta - \phi_s - profit_s \leq 0 \qquad \forall s \in \mathcal{S} \tag{4.35}$$

4.2.5 Selecting Price Points $\{y_{t,b}\}_{t\in\mathcal{T},b\in\mathcal{B}}$ Based on Jenks Natural Breaks Partitioning

While the above model selects the volume components $x_{t,b}$ of the price-volume bidding pairs, this section details the approach to selecting the price points $y_{t,b}$. As depicted in Figure 4.2.1 and described in constraint (4.29), these price points partition the scenarios for a specified time interval, such that the planned power output of the hybrid is the same for all scenarios within a class of the partition. The proposed procedure for selecting price points for each time interval is based on the day-ahead price scenarios, $\{\pi_s, \{\rho_{t,s}^{DA}\}_{t\in\mathcal{T}}\}_{s\in\mathcal{S}}$, and has three phases:

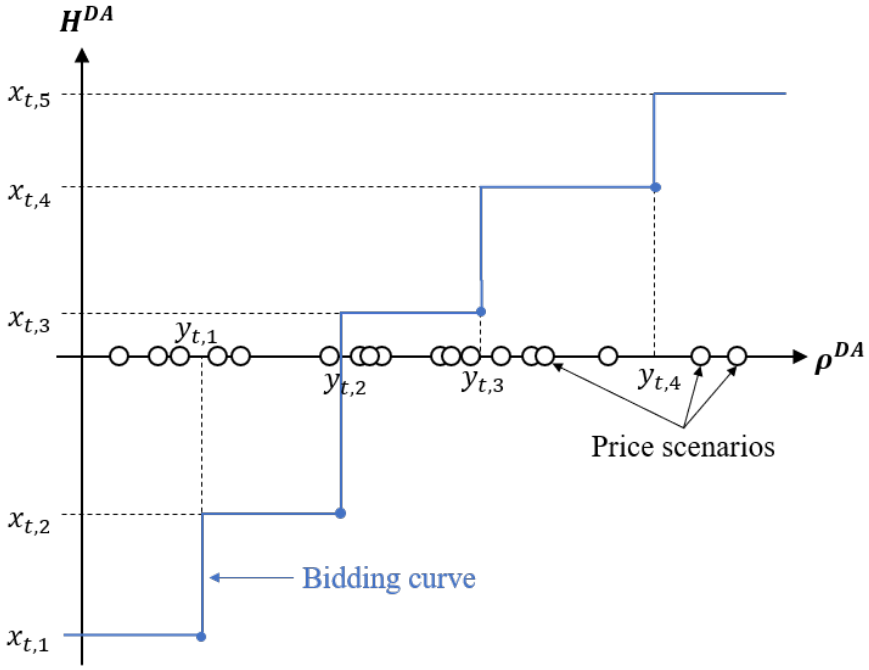

Figure 4.2.1: Bidding curve illustration for hour t highlighting the role of price points in the grouping of scenarios.

i. **Partition Phase:** Using Jenks natural breaks algorithm ("Jenks") [59], partition the day-ahead price scenarios $\{\rho_{t,s}^{DA}\}_{s\in\mathcal{S}}$ into classes $\{P_{t,b}\}_{b\in\mathcal{B}}$ for every $t \in \mathcal{T}$, which are ordered in the sense that

$$\bar{P}_{t,b} := \max_{\rho \in P_{t,b}} \rho < \min_{\rho \in P_{t,b+1}} \rho := \underline{P}_{t,b+1}, \ \forall b = 1, \ldots, N_B - 1, \ \forall t \in \mathcal{T}. \tag{4.36}$$

Jenks minimizes the within-class variation so that the price scenarios forced to share a common bid quantity via constraint (4.29) are as similar as possible. This is an advantageous property which does not hold if you divide the price domain evenly or assign the same number of price scenarios to each class, the approach used by [81]. The author believes this to be a new application of Jenks, which was originally developed for coloring choropleth maps.

ii. **Adjustment Phase:** There is a critical difference between prices above and below the generator's operating cost, O_t^G. If prices are above this threshold, the generator will not be curtailed. If prices are below this threshold, the generator may be curtailed if it is not possible or profitable to charge the ESR. Therefore, all prices within the same class should be uniformly above or uniformly below O_t^G. This phase adjusts the Jenks Natural Breaks partition, if necessary, using Algorithm 4.1 for every $t \in \mathcal{T}$ to ensure this property holds.

Algorithm 4.1 Adjustment phase procedure for selecting price points

Require: $\{P_{t,b}\}_{b\in\mathcal{B}}$, O_t^G
for $b = 1, \ldots, N_B$ **do**
  if $\min_{\rho\in P_{t,b}} \rho < O_t^G < \max_{\rho\in P_{t,b}} \rho$ **then**
    Identify prices above the threshold: $A = \{\rho \in P_{t,b} \mid \rho \geq O_t^G\}$
    Identify prices below the threshold: $B = \{\rho \in P_{t,b} \mid \rho < O_t^G\}$
    if $b = 1$ **then**
      Reassign prices above the threshold to the next class: $P_{t,b} = B$, $P_{t,b+1} = P_{t,b+1} \cup A$
    else if $b = N_B$ **then**
      Reassign prices below the threshold to the prior class: $P_{t,b} = A$, $P_{t,b-1} = P_{t,b-1} \cup B$
    else
      Reassign prices in A or B based on which has lower combined probability:
      if $\sum_{s|\rho_{t,s}^{DA}\in B} \pi_s \leq \sum_{s|\rho_{t,s}^{DA}\in A} \pi_s$ **then**
        $P_{t,b} = A,\ P_{t,b-1} = P_{t,b-1} \cup B$
      else
        $P_{t,b} = B,\ P_{t,b+1} = P_{t,b+1} \cup A$
      end if
    end if
  end if
end for

iii. **Selection Phase:** Select prices $y_{t,b}$ for the bidding curves in such a way that they divide the classes. That is the selected prices must satisfy the following inequalities:

$$y_{t,0} < \underline{P}_{t,1} \quad \text{and} \quad \bar{P}_{t,b} \leq y_{t,b} < \underline{P}_{t,b+1} \quad \forall b = 1, \ldots, N_B - 1,\ t \in \mathcal{T}. \tag{4.37}$$

Specifically, we chose to set the price points equidistant between adjacent classes, unless it divides a class above the marginal cost of generation from one below it.

$$y_{t,b} = \begin{cases} \underline{P}_{t,1} & b = 0 \\ \rho^{ceil} & b = N_B \\ O_t^G & \bar{P}_{t,b} < O_t^G \leq \underline{P}_{t,b+1} \text{ and } b \geq 1 \\ (\bar{P}_{t,b} + \underline{P}_{t,b+1})/2 & \text{otherwise} \end{cases} \tag{4.38}$$

A variation on this approach is to set the price points proportionally to the probability of the adjacent classes, thereby generalizing the bid associated with high-probability price scenarios to a wider range of market prices.

4.2.6 Discussion

First, it should be verified that the model meets the criteria stated in section 4.2.1. Requirements 1-5 and 10 were addressed through the model formulation. As all decision variables

are bounded as a result of requirement 3 and the objective function is linear, the problem will remain bounded even if there is a mixture of positive and negative prices, either across time periods or across markets for the same time period, in accordance with requirement 6. Requirement 7 is covered in section 4.3.4. The uncertainty model here is well-suited for simulation environments, as demanded by requirement 8, since it can easily accommodate more or fewer scenarios depending on how much information is available. Section 4.3.5 addresses requirement 9 on tractability. The ESR is characterized by the rates and efficiency with which it can charge and discharge energy, its energy storage capacity, and operating (i.e., degradation) cost. Appropriate definitions of these parameters can reflect a multitude of storage technologies. The generator has a similarly general characterization ensuring requirement 11.

Compared with the existing literature, this model can be viewed as advancing the one proposed in [41] in two significant ways. First, the bid curves here will have at most N_B marginal price steps. In contrast, [41], along with [80] and [58], may produce bid curves with as many breakpoints as there are price scenarios. Price scenarios should reflect the underlying uncertainty distribution, and restricting the number of scenarios to N_B in order to comply with market rules limiting the number of price steps competes with this priority. In [81], scenarios are grouped into classes to limit the number of breakpoints by assigning the smallest N_S/N_B scenarios to the first class, the next smallest N_S/N_B scenarios to the second class, etc., regardless of their distribution. Compared with this approach, the one in section 4.2.5 is more sophisticated and conscious of discontinuities in the optimal generator production and curtailment around O_t^G. Second, this model provides the hybrid operator greater flexibility for the two components and between day-ahead and intraday markets without sacrificing the principle that bids and planned operating schedules respect all physical limits. For example, [41] constrains $G_{t,s}^{IN} \geq 0$ for all $t \in \mathcal{T}$ and $s \in \mathcal{S}$, while this model allows the generator to plan to adjust downward in the intraday market relative to day-ahead. This flexibility could be used for a scenario in which no generation is expected in hour t, but $G_t^{DA} > 0$ due to other scenarios in the same price class. A second example of this enhanced flexibility is allowing the generator to curtail its output. While this flexibility allows the hybrid to arbitrage between day-ahead and real-time markets, it can only do so at a low volume, unlike virtual bidders, due to the previously stated principle that $H^{SC}, G^{SC}, E^{SC,cha/dis}$, and H^{DA} all must reflect the plant's physical limits.

4.2.7 Model Cases and Extensions

The model in section 4.2.4 is flexible in types of hybrids and bids it can represent and easily extensible to additional constraints and objectives. Guidance on modeling a number of common situations is found below.

i. Self-scheduled bids: In this setting, the hybrid operator does not submit day-ahead bid curves to the market, rather they submit hourly volumes which they plan to produce regardless of the market price. The proposed bidding model can be tailored to create

self-scheduled bids by setting $N_B = 1$, $y_{t,1} = \min_{s \in \mathcal{S}}\{\rho^{DA}_{t,s}\}$, and $y_{t,2} = \max_{s \in \mathcal{S}}\{\rho^{DA}_{t,s}\}+1$. Then, $x_{t,1}$ is the self-scheduled bid for hour t.

ii. AC versus DC coupling: Some hybrids are built with separate DC/AC converters for the generation and storage components, these are known as AC-coupled, while others only utilize one DC/AC converter in a DC-coupled configuration. The proposed model can represent either configuration, if the appropriate input parameters are used. For an AC-coupled hybrid, *gen* scenarios should represent the (potentially clipped) power after the inverter and, for a DC-coupled hybrid, *gen* is the power before the DC/AC inverter. For a DC-coupled hybrid, *poi* should be redefined as the minimum of the point of interconnection power capacity and the shared DC/AC inverter power rating.

iii. ESR cycle limit: The proposed model addresses wear-and-tear on the ESR through its associated operating cost parameters, $O^{E^{SC,dis}}$ and $O^{E^{SC,cha}}$. However, one may wish to directly limit the ESR use to a specific number of cycles to reflect a battery's warranty restrictions, for example. This can be accomplished by adding the following constraints to the model:

$$\sum_{t=1}^{N_T} E^{SC,dis}_{t,s} \leq C^{max} \times \{\text{cycle limit}\} \quad \forall s \in \mathcal{S} \tag{4.39}$$

$$\sum_{t=1}^{N_T} E^{SC,cha}_{t,s} \leq C^{max} \times \{\text{cycle limit}\} \quad \forall s \in \mathcal{S}. \tag{4.40}$$

iv. Capacity market: The proposed model exclusively focuses on energy markets. If the hybrid also participates in a capacity market which pays an additional amount, ρ^{CAP}, per unit produced during peak demand hours, this can be incorporated into the objective function with the following update to equation (4.3):

$$revenue_s := \sum_{t=1}^{N_T} \rho^{DA}_{t,s} H^{DA}_{t,s} + \left(\rho^{DA}_{t,s} + \rho^{IN}_{t,s}\right) H^{IN}_{t,s} + \rho^{+}_{t,s}\epsilon^{+}_{t,s} + \rho^{CAP}_{t,s} H^{SC}_{t,s} \tag{4.41}$$

4.3 Simulating Hybrid Bids Within a Future Electricity Market

Obtaining bids from the stochastic optimization model in section 4.2 for use in a production cost model requires the development of scenarios and choice of model hyperparameters. The production cost model considered is one of a hypothetical system, so there is no historical record of market prices or the hybrid's past bids. This section will present an approach to setting all model inputs while discussing various trade-offs that researchers should be aware of. The approach is designed to reflect the accuracy of information which would

typically be available to a market participant, even if the researcher has access to better data. Further, the approach aims to consistently produce bids which perform well and are not overly sensitive to small changes in the model inputs. The approach is applied to the project mentioned in section 4.1.3 and the specific parameter choices made for this project are shared.

Sections 4.3.1, 4.3.2 and 4.3.3 cover the formation of scenarios, while section 4.3.4 discusses selecting hyperparameters for the model, including the number of scenarios. Implementation of scenario formation algorithms and the decision-making model is covered in section 4.3.5. Finally, section 4.3.6 proposes a heuristic for adding elasticity to the bid curves to account for prices beyond the range considered by the price scenarios.

4.3.1 Generation Scenarios (*gen*)

This section describes how the generation scenarios are constructed to reflect the range of plausible generation levels and their relative likelihood. As a reminder, the *gen* parameter being built here represents the maximum power output possible for a generation given the environmental conditions, i.e., the generation in the absence of curtailment.

Information available

In this setting, the researcher has access to the maximum generation levels which will be used in the production cost model and day-ahead forecasts of these values for a time period of at least several months, ideally longer. The bidder is allowed access to the day-ahead forecast for the optimization horizon (periods $1, 2, \ldots, N_T$) and the empirical distribution of forecasts and forecast errors for two consecutive periods for the full data set (Table 4.3.1).

Forecast F	Forecast Error E_1	Forecast Error for Preceding Period E_2
9.0	0.5	0.8
0.0	0.0	0.0
3.0	-2.0	-1.0
...	...	...

Table 4.3.1: Illustration of the empirical distribution of generation forecast accuracy (forecast error = actual - forecast) and sequential correlations available to the bidder.

Monte Carlo scenario generation

Instead of relying on a single day-ahead generation forecast to inform decision-making, a set of plausible scenarios are built around the forecast using the empirical distribution of forecast errors. The approach described in this section accounts for two factors correlated with forecast error: the forecast itself and the error in the preceding period. To understand the correlation between forecast value and forecast error, consider a solar plant. This plant has a forecasted generation of zero for many hours each day (i.e., the nighttime hours) and very

small forecast error during such hours. However, when the forecasted generation is half of the plant's capacity there is uncertainty due to weather. So, plausible scenarios for these two hours (one with a forecast of zero and an hour with a moderate forecast value) should not be based on the same error distribution. Also, errors in consecutive periods are positively correlated, as demonstrated in Figure 4.3.1. To address these two factors, the generation scenario procedure outlined in Algorithm 4.2 samples uniformly from the distribution of errors in periods with similar forecast values and similar error values in the preceding period. An example set of scenarios created with this procedure is shown in Figure 4.3.2a. Depending on amount and type of data available and plant-specific knowledge, the distribution could be based on additional factors such as time of day, error trend, or season.

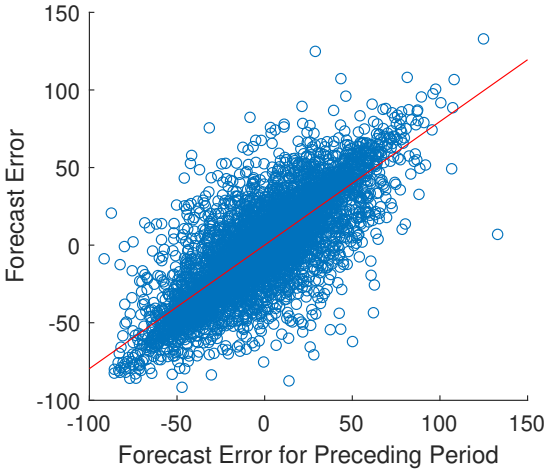

Figure 4.3.1: Hourly data and best-fit line for a wind plant in New York for 2019. Note the correlation between consecutive forecast errors.

Algorithm 4.2 Generation scenario s construction

Require: Forecasted generation values $f_1, , f_2, \ldots, f_{N_T}$
Require: Hyperparameters Δ_f and Δ_e which define similarity thresholds for forecast and error values, respectively
Require: Generator capacity G_t^{max}
Require: Data set described in Table 4.3.1

1: Select error value for period 1 based on forecast*: $e_1 \sim E_1 \mid F \in [f_1 - \Delta_f, f_1 + \Delta_f]$
2: Adjust e_1 to ensure feasibility: $e_1 = \min(G_t^{max}, \max(0, e_1 + f_1)) - f_1$
3: **for** $t = 2, \ldots, N_T$ **do**
4: Select error value for period t based on forecast and error in preceding period*: $e_t \sim E_1 \mid (F \in [f_t - \Delta_f, f_t + \Delta_f], E_2 \in [e_{t-1} - \Delta_e, e_{t-1} + \Delta_e])$
5: Adjust e_t to ensure feasibility: $e_t = \min(G_t^{max}, \max(0, e_t + f_t)) - f_t$
6: **end for**
7: **return** $gen_{t,s} = f_t + e_t \ \forall t \in \mathcal{T}$

**If no points in the data set meet the conditional criteria, temporarily increase Δ_e and/or Δ_f until the distribution is nonempty.*

Scenario reduction with fast forward algorithm

It is straight-forward and computationally inexpensive to create a large number of equally likely generation scenarios using Algorithm 4.2. However, the stochastic optimization prob-

lem in Section 4.2 will become computationally challenging if the number of scenarios is too large. Further, many of the scenarios may be similar and the breadth and frequency of possible outcomes could be well-represented by a smaller number of scenarios with varying probabilities. This is the objective of scenario reduction methods.

The Fast Forward Selection (FFS) algorithm [47] is used in this project to perform scenario reduction on the generation scenarios. FFS is considered the "state-of-the-art" [115], and has been used in past research on energy storage bidding in electricity markets [69], because it is effective and efficient [33]. FFS attempts to solve the optimal reduction problem though a greedy heuristic which recursively selects a single "best" scenario until the desired number of scenarios have been selected. An example outcome of FFS is shown in Figure 4.3.2. Alternative methods of scenario reduction include simultaneous backward reduction [47], Latin hyper-cube sampling [81, 106], submodularity-based reduction [115], importance sampling [95], and heuristic search [77].

Scenario weighting

The scenarios selected by FFS have an associated likelihood relative to each other, however the probability of these scenarios relative to the original forecast still needs to be addressed. In this project, the day-ahead forecast was assigned a probability of 0.8, and the selected scenario probabilities were each scaled by 0.2 so that the set has a combined probability of 0.2. These values (0.8, 0.2) were selected based on a combination of experiments testing how often the generated and selected scenarios were closer than the forecast to the actual generation profile (closer in the sense that the ℓ^2-norm of the difference is smaller) and engineering insight. A modified approach would be re-weight the forecast and scenarios before performing scenario reduction, and then include the forecast in the set of scenarios provided to FFS. This modified approach would not guarantee that the original forecast is used as a generation scenario in the decision-making model defined in Section 4.2 and was rejected for this reason.

4.3.2 Price Scenarios

Creating meaningful price scenarios in a simulation environment is a unique challenge, because prices are an output of the production cost model and therefore there is not a record of past prices. Price scenarios should predict both the shape of the price profile (e.g., when the highest and lowest price hours will occur) and the magnitude of prices. The latter is important because these scenarios will be used to set the price component of the bids, as discussed in section 4.2.5. This section provides a methodology for creating price scenarios in a simulation environment where bids for the entire study horizon must be submitted at one time and there is no opportunity to use price results from one day to inform price scenarios for a future day.

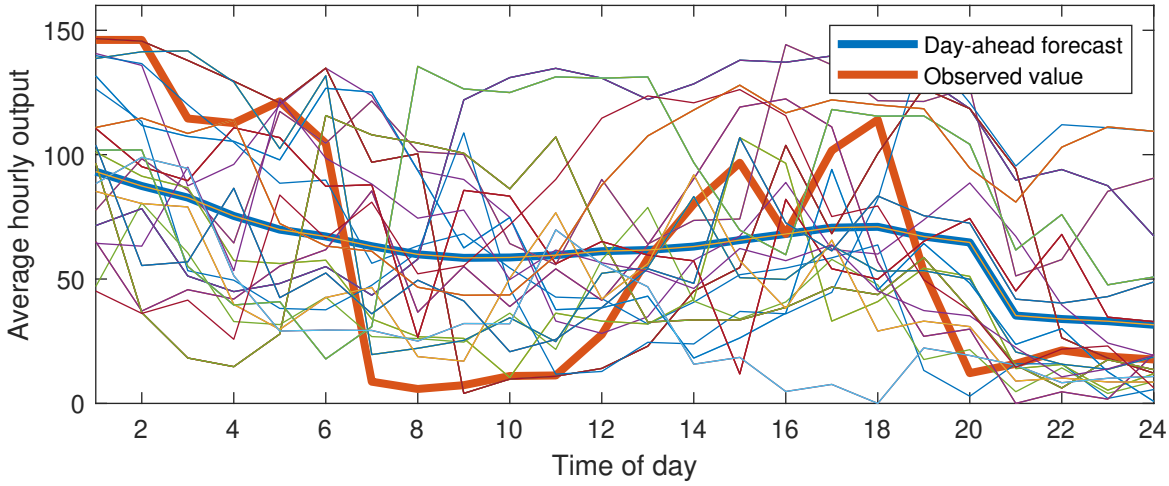

(a) Twenty-five generation scenarios (thin lines) generated via Algorithm 4.2

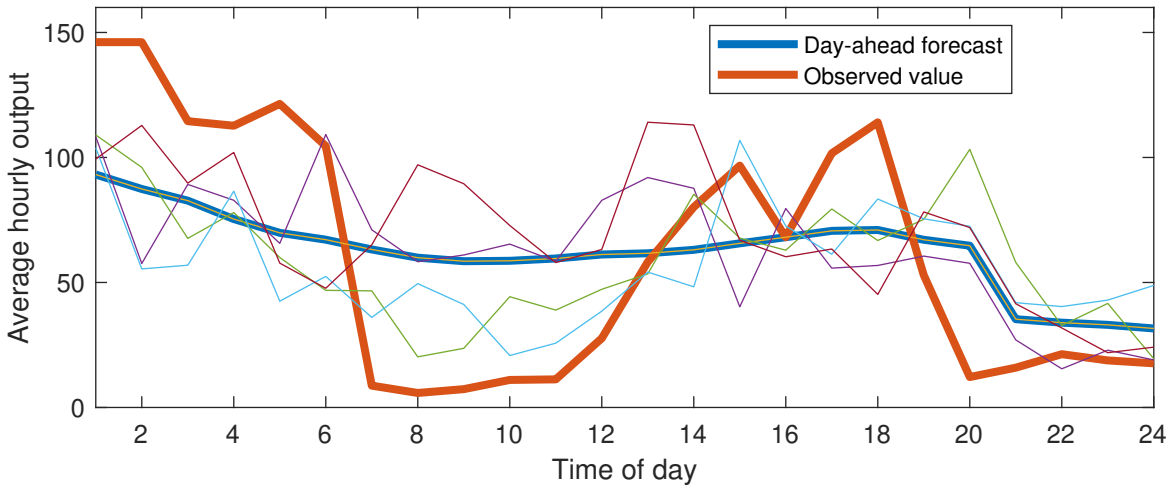

(b) Five generation scenarios (thin lines) chosen by FFS

Figure 4.3.2: Generation scenarios compared with the day-ahead generation forecast (thick blue line) and realized generation level (thick orange line) for a wind plant in New York.

Information available

In this setting, the researcher has access to the simulated day-ahead and intraday market prices for a system which is identical, other than that the hybrids are in a system operator-managed "separate independent resources" (abbreviated "2R") participation model. This is a version of option (a) in Figure 4.1.1 in which the assets are dispatched by the system operator to minimize overall system cost, as opposed to being bid into the market by the hybrid operator. These "2R prices" are available for the optimization horizon (periods $1, 2, \ldots, N_T$) and at least a few weeks prior. The approach is based on the assumption that market prices will be similar under both participation models. The bidder is not allowed to use the 2R prices for the optimization horizon directly as a price scenario, but they may use this profile to design price scenarios reflective of those which a knowledgeable market participant could hope to produce.

Day-ahead price scenarios meeting state-of-the-art forecast accuracy (ρ^{DA})

A set of plausible price scenarios is built around the day-ahead "2R profile" by taking price profiles of recent days and incrementally improving them until they meet a state-of-the-art electricity price forecasting metric, where the 2R profile is treated as the ground truth. The resulting price forecasts are representative of forecasts that a state-of-the-art price forecasting method could produce and are treated as equally likely. Pseudocode for this approach is provided in Algorithm 4.3.

Algorithm 4.3 Day-ahead price scenario development

Require: Baseline price data set ("DATA"), accuracy metric computation function ("METRIC"), target accuracy metric value ("TARGET"), error reduction factor ("FACTOR")
1: scenarios ← price profiles of comparative dates from DATA
2: TRUTH ← price profiles from DATA, organized to correspond with scenarios
3: **while** METRIC(scenarios, TRUTH)>TARGET **do**
4: *Improve scenarios:* scenarios ← TRUTH +(1−FACTOR) × (scenarios−TRUTH)
5: **end while**
6: **return** scenarios

This study selected the mean value of weekly-weighted mean absolute errors (WMAE) as the metric and 5% as its target value, informed by the results presented in [118] and [124]. To utilize WMAE, scenarios for at least seven days must be computed at once. Scenarios were initialized to be the past ten weekdays, for each weekday, and the past six weekend days, for each Saturday and Sunday. 0.05 was used as the error reduction factor. Figure 4.3.3 provides an example of the final day-ahead price scenarios (those returned on line 6 of Algorithm 4.3) compared with the initial, naive scenarios.

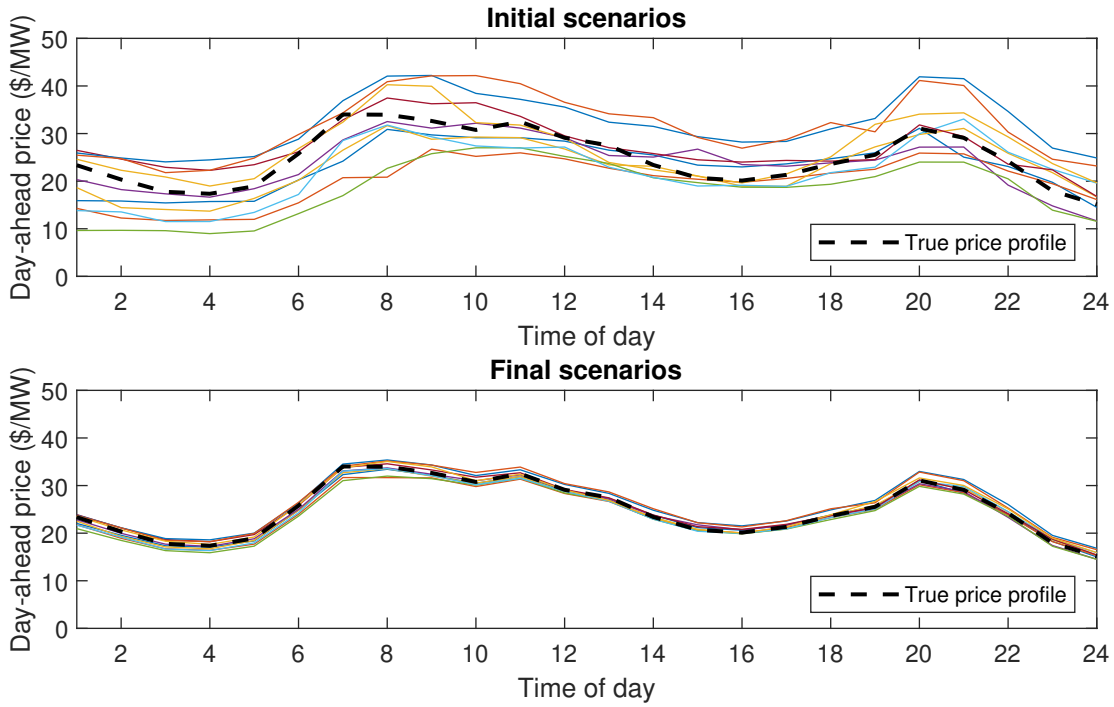

Figure 4.3.3: Day-ahead price scenario development: The top graph shows how the initial scenario set (ten most recent weekdays) for a given weekday compares to the true prices. For the same date, the bottom graph shows the final price scenarios which, along with scenarios for each day of a year, correspond to a mean WMAE of 4.9%.

Scenarios for the difference between intraday and day-ahead market prices (ρ^{IN})

In this project a real-time market was not simulated, so Λ was set to zero and $\rho^{IN}_{t,s}$ also took a value of zero, for all $t \in \mathcal{T}$ and $s \in \mathcal{S}$. An approach similar to that proposed for generation scenarios could be utilized, with the following substitutions: F =day-ahead price, E_1 =price difference between the intraday and day-ahead markets, E_2 =price difference between the intraday and day-ahead markets in the preceding period. Alternative methods can be found in the literature, including [41, 18, 69].

Scenarios for the deviation price ratios ($\eta^{+/-}$)

As stated in Lemma 4.1, deviations will never be priced advantageously to the day-ahead market if $\eta^{+}_{t,s} < 1$ and $\eta^{-}_{t,s} > 1$ for all $t \in \mathcal{T}$ and $s \in \mathcal{S}$. In this project, the $\eta^{+/-}$ values were designed to generally limit differences between the hybrid's actual and scheduled output, as opposed to reflecting deviation charges in a specific market structure. The choices $\eta^{+}_{t,s} = 0.5$ and $\eta^{-}_{t,s} = 1.5$ for all $t \in \mathcal{T}$ and $s \in \mathcal{S}$ accomplished this objective.

4.3.3 Combined Scenarios

A symmetric scenario tree was created to pair each generation scenario with each day-ahead price scenario. Since the day-ahead price scenarios are all equally likely, $\{\pi_s\}_{s\in\mathcal{S}}$ are proportional to the generation scenario probabilities.

4.3.4 Hyperparameters

This section will discuss the selection of design parameters which the researcher must provide to the bidding model in section 4.2. The parameters specific to a plant's technology and configuration, including $poi, \mu, O, G^{max}, E^{cha/dis,max}$ and $C^{min/max}$, are not covered.

Risk parameters

Risk in the bidding model is managed through the parameters ζ, the weight applied to the CVaR in the objective function, and σ, which regulates the scenarios CVaR represents. We chose to fully consider financial risk by setting $\zeta = 1$ and have CVaR be the expected profits of the lowest 5% of scenarios by setting $\sigma = 0.95$.

Time horizon

The choice of N_T depends on the market structure and storage duration. If the day-ahead market covers 24 hours, clearly N_T must be at least 24. Considering a longer time horizon when making bidding decisions, even if the bids produced for $t = 25, 26, \ldots, N_T$ are not submitted to the market, is beneficial because the bids are based on a more complete view of the opportunity costs of charging and discharging the ESR. For example, if $N_T = 24$ and $\rho^{DA}_{24,s} > 0$ then $C^{SC}_{24,s}$ will generally be zero or near zero since there is no opportunity in the model for energy in storage at $t = 24$ to ever be sold. In practice, there will be another day and automatically emptying the battery each night without considering it is not a wise strategy. Requirement 7 can be satisfied by selecting N_T to be sufficiently greater than the number of binding time periods. The hybrids considered in this study had a storage duration of four hours, so a one-day look-ahead period ($N_T = 48$) was appropriate. Hybrids with longer duration storage would benefit from modeling longer time horizons. Bidding model simplifications may be used during the look-ahead period in order to satisfy computational or memory limits, for example using self-scheduled bids or identical scenarios during the look-ahead hours.

Number of scenarios (N_S)

The aim is to create bids which are not overly sensitive to small changes in the model inputs, chief among which are the scenarios. An example of a small change is including one additional scenario created in the same manner as the rest. The price scenario methodology introduced in 4.3.2 produces ten price scenarios for a weekday and six price scenarios for a

weekend day. The options for changing the number of price scenarios are more limited than for generation scenarios and have a smaller impact on the resulting bids due to their lower variability (compare Figures 4.3.2 and 4.3.3). So, this section will focus on how to determine the number of generation scenarios.

There are two numbers of generation scenarios which need to be determined: (1) the number of equally-likely scenarios created using Algorithm 4.2 and (2) the number, N_{FFS}, to be selected by FFS and ultimately used as bidding model inputs. Selecting a large number, say 500 or 1000, for (1) produces a more complete distribution and is less computationally taxing than increasing the dimension of the bidding problem, especially since the scenarios can be generated in parallel. Thus the more challenging problem is determining a good value for N_{FFS}. The trade-off is between bidding problem computation (increasing as N_{FFS} increase) and the resulting bid robustness to generation uncertainty (improving as N_{FFS} increases). As the specific terms of this trade-off will vary based on the setting and the resources available, we share an approach to evaluating how many scenarios are enough:

> Repeatedly solve the hybrid bidding problem in section 4.2 with increasing numbers of generation scenarios (values ranging from 1 to 200 are used in Figure 4.3.4), and the same price scenarios, for each of a small selection of dates (six are used in Figure 4.3.4). As you grow the scenario set, be sure to only add scenarios (e.g., all scenarios in the trial with 10 generation scenarios are used in the trial with 20 generation scenarios) and note the computation times. By analyzing the rate at which bids converge to the bids based on the greatest number of scenarios ($x_t^{(max)}$), search for a value of N_{FSS} that is computationally tractable for the project and will produce bids which are acceptably close to those based on much more information and very similar to those based on a few more or a few less scenarios.

Figure 4.3.4 provides an example of this analysis, where the difference between two sets of N_T bid curves, one based on $N_{FSS} = k$ scenarios and one based on $N_{FSS} = max$ scenarios, is measured as follows:

$$\text{Average hourly bid difference} := \frac{\sum_{t=1}^{N_T} \|x_t^{(k)} - x_t^{(max)}\|_2}{\sum_{t=1}^{N_T} \|x_t^{(max)}\|_2}. \tag{4.42}$$

In this example, if you are comfortable with bids that are approximately 10-20% different on average than the bids based on 200 generation scenarios, then you only need to use 10 generation scenarios moving forward. However, if you want bids that are within 5% of the 200-generation-scenario bids on average, Figure 4.3.4 suggests that close to 100 generation scenarios are needed. Ultimately, we decided to use 20 generation scenarios after analyzing Figure 4.3.4 and our computational resources.

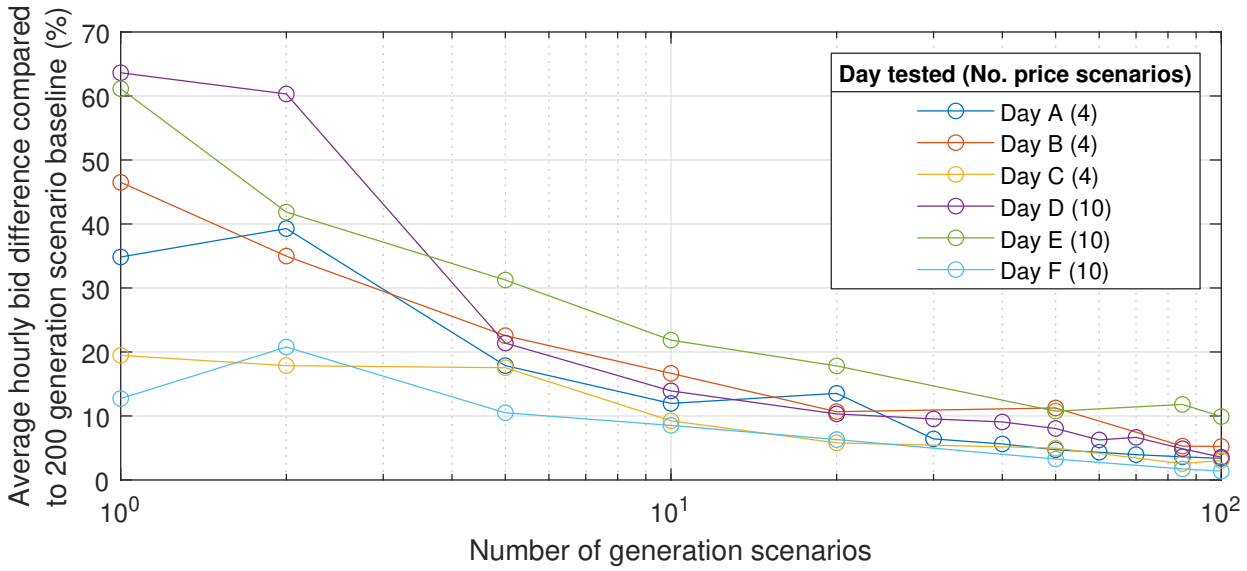

Figure 4.3.4: Experimental results for 6 days comparing bid curves based on 200 generation scenarios to those based on 1-100 generation scenarios.

Initial state of charge

Ideally, the initial state of charge parameter C_{start} is known precisely or can be estimated well. In practice this is usually the case. Bids are typically submitted 12 hours in advance, so the state of charge can be estimated based on the current state of charge, current day's bids and 12 hours of generation forecasts, as well as managed using the intraday market. Current production cost models do not allow for the co-simulation of battery state of charge and hybrid participant decision-making. A result of this software limitation for multi-day studies is that bids for the entire study horizon (weeks, or even years worth of bids) must be submitted to the "market" in advance. In this case, what value of C_{start} should be used for day 42, for example, given the uncertainty in generation and dispatch for days 1-41?

As Figure 4.3.6 demonstrates, this choice has a large influence on the bids for the first several hours of the day. Bids resulting from five different

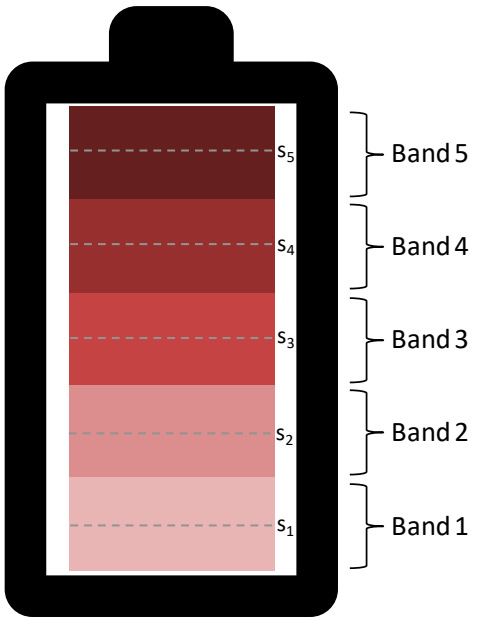

Figure 4.3.5: Illustration of solution to initial state of charge estimation challenge

C_{start} values, and identical values for all other parameters, are compared and show that a large C_{start} allows the solar hybrid to deliver power in hour 1 and then remain idle until the morning peak price hours (i.e., hour 7). Conversely, hybrids with a low C_{start} charge from the grid via negative quantity bids during the early morning hours.

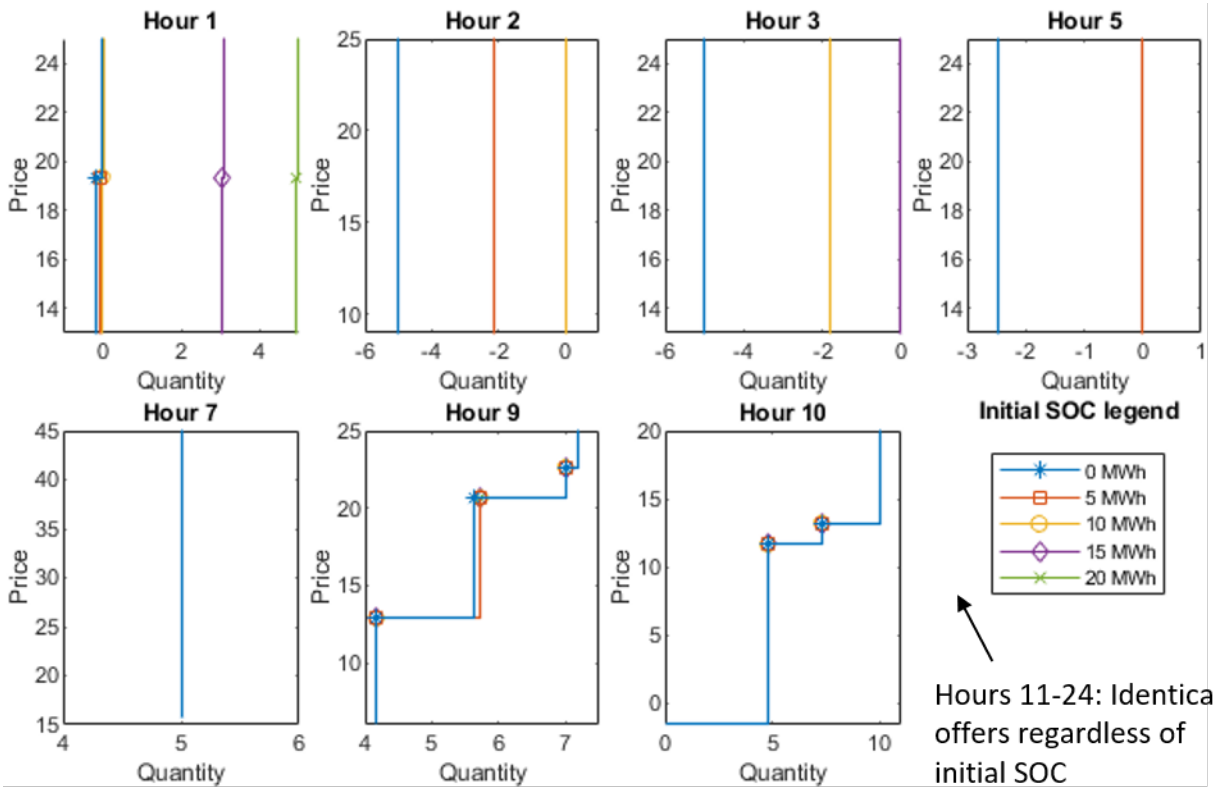

Figure 4.3.6: Sensitivity analysis of the initial state-of-charge parameter C_{start}.

The solution implemented in this project was to create several sets of bidding curves for each day, each corresponding to a different initial state of charge condition. As illustrated in Figure 4.3.5, the approach begins by solving the bidding problem in section 4.2 five times for $C_{start} = s_1, s_2, \ldots, s_5$ and all other parameters held constant. Then, for each simulated day-ahead market, the production cost model utilizes only the single set of bidding curves corresponding to the band of the observed state of charge. For example, if the state of charge at the end of day 41 in the simulations falls within band 2, then the bids for day 42 are based on $C_{start} = s_2$. The number and placement of the state of charge levels and their corresponding bands is up to the individual researcher. A larger number of narrow bands will improve the bids, but increase the computation and memory requirements. This study used the values listed in Table 4.3.2; narrower bands near the extreme values were used to guard against infeasible schedules early in the day.

C_{start}	2.5	12.5	30	50	70	87.5	97.5
Band	[0, 5]	(5, 20]	(20, 40]	(40, 60]	(60, 80]	(80, 95]	(95, 100]

Table 4.3.2: Values used to define and select state-of-charge dependent bidding curve sets (percentage of $C_{1,s}^{max}$).

4.3.5 Implementation

A MATLAB implementation of the hybrid power plant bidding model in section 4.2 with the parameters discussed in section 4.3.4 (notably, $N_T = 48$ and $N_S = 200$ for a weekday) was solved in, on average, 45 seconds using MATLAB's mixed-integer linear programming solver from its "optimization toolbox" and 24 seconds using Gurobi on a personal computer, which satisfied the tractability demands of requirement 9.

4.3.6 Heuristic to Add Elasticity for High- and Low-price Events

The advantage of bidding curves over self-scheduled bids is the ability to be dispatched differently if prices are higher or lower than your expectation. The day-ahead price scenarios developed in section 4.3.2 allow for elasticity over the range of probable prices. However, prices may spike outside of the range of day-ahead price scenarios in ways that are difficult to predict, and the hybrid operator wants to capitalize on unexpectedly high and low prices. This section presents a heuristic approach to extending the bids created by the section 4.2.4 model to address exceptionally high and low prices.

In the event of a high price, a hybrid operator wants to be dispatched at its maximum possible power output. The challenges are to determine what constitutes a high-price event and to ensure the resulting curve is logical with respect to the marginal cost of generation. Figure 4.3.7 illustrates the proposed heuristic for high price events in both notable cases. The details are found in Algorithm 4.4, which relies on two hyperparameters: $\underline{y}^{high}$ sets an absolute floor on what price is considered "high" and α establishes a high price as a multiple of the highest hourly price among all scenarios over the time horizon. In this study, we used $\alpha = 2$ and set $\underline{y}^{high}$ equal to the expected 80$^{\text{th}}$ percentile day-ahead price in the given month.

In the event of a low price, a hybrid operator wants to charge the ESR as much as possible. Similar to high-price events, the challenges are to determine what prices are treated as "low" and to ensure the final bidding curve is logical with respect to the marginal cost of generation and potential grid charging restrictions. Figure 4.3.8 illustrates the proposed heuristic, which is detailed in Algorithm 4.5, for two common cases. The user-defined hyperparameters are $\bar{y}^{low}$, which sets an absolute ceiling on what price is considered "low," and α, which again establishes a low price as a multiple of the highest hourly price among all scenarios over the time horizon. In this study, we used $\alpha = 2$ and set $\bar{y}^{low}$ equal to the expected 20$^{\text{th}}$ percentile day-ahead price in the given month.

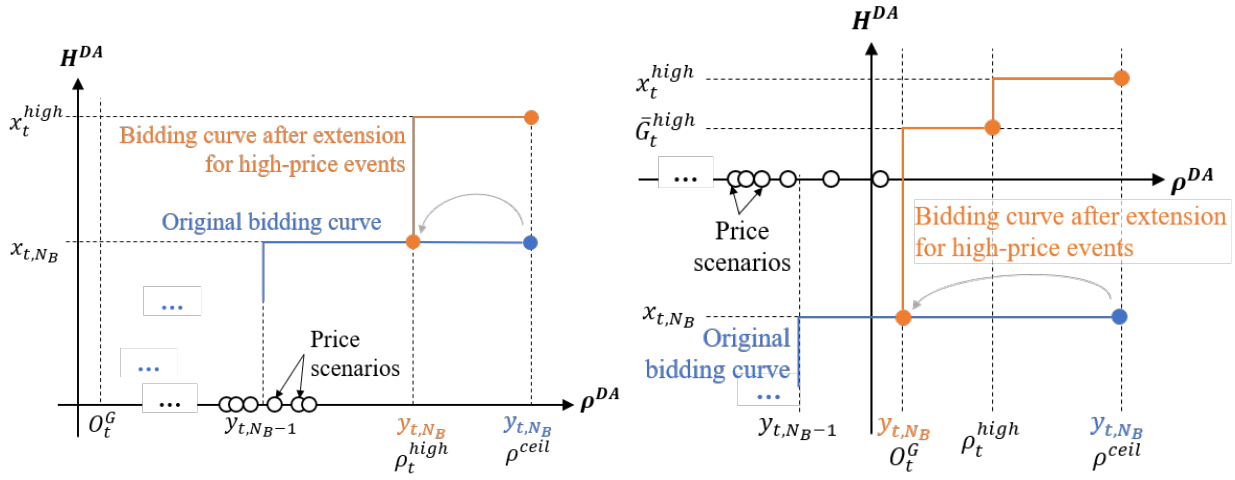

(a) Case 1 (b) Case 2

Figure 4.3.7: Illustration of the high-price event heuristic (Algorithm 4.4) in two typical cases.

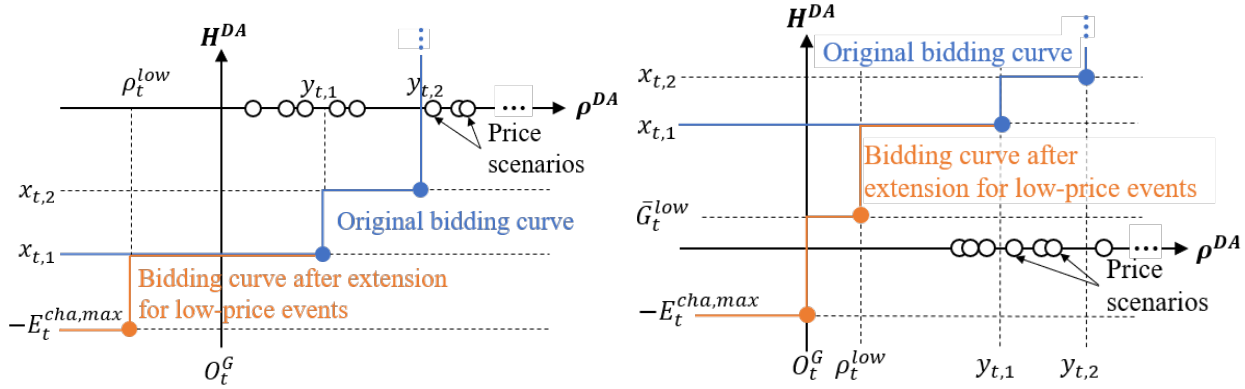

(a) Case 1 (b) Case 2

Figure 4.3.8: Illustration of the low-price event heuristic (Algorithm 4.5) applied to two common cases in which grid charging is allowed ($\psi = 1$).

4.4 Conclusion

Hybrid power plants are expected to contribute significant capacity to power systems in the near future. To analyze and prepare for this future, researchers require the ability to represent hybrids' market actions in production cost models. This paper addresses this need by modeling a hybrid operator's day-ahead bidding decisions as a stochastic optimization problem and then detailing methods for deploying the proposed model in a simulation environment. The model builds upon those in the literature first by allowing for greater flexibility of the generator and ESR individually, without lessening constraints on the combined hybrid resource, in accordance with a single, self-managed resource participation model. Second,

the model respects limits on a bid's complexity through its number of breakpoints, without restricting the number and distribution of price scenarios. The unique challenges of developing price and generation scenarios for a hypothetical system are addressed through algorithms and examples. Finally, the bid curves are adapted to provide favorable dispatch in the face of low-probability prices events.

Algorithm 4.4 Bid curve extension for high-price events

Require: Model inputs and results: $\{\{\pi_{t,s}, gen_{t,s}, \rho^{DA}_{t,s}, \bar{G}_{t,s}\}_{s\in\mathcal{S}}, \{x_{t,b}, y_{t,b}\}_{b\in\mathcal{B}}, E^{dis,max}_t, O^G_t\}_{t\in\mathcal{T}}$
Require: Market price ceiling: ρ^{ceil}
Require: Minimum price considered a high-price event: $\underline{y}^{high}$
Require: Factor defining a high-price event relative to the time horizon: $\alpha > 1$
for $t \in \mathcal{T}$ **do**
  Calculate the expected maximum possible power output:

$$x^{high}_t = \left(\sum_{s\in\mathcal{S}} \pi_s\right)^{-1} \left(\sum_{s\in\mathcal{S}} \pi_{t,s} \cdot \min\{E^{dis,max}_t + gen_{t,s},\ poi\}\right) \tag{4.43}$$

  if $x^{high}_t > x_{t,N_B}$ **then**
    Insert $(\rho^{ceil},\ x^{high}_t)$ into the hour t bidding curve
    Calculate the high-price event threshold:

$$\rho^{high}_t = \max\left\{\underline{y}^{high},\ \alpha \cdot \left(\max_{s\in\mathcal{S},t\in\mathcal{T}} \rho^{DA}_{t,s}\right),\ \frac{1}{\alpha} \cdot \left(\max_{s\in\mathcal{S},t\in\mathcal{T}} \rho^{DA}_{t,s}\right)\right\} \tag{4.44}$$

    Calculate the expected curtailment of scenarios on the highest-price segment:

$$\bar{G}^{high}_t = \frac{\sum_{s\in\mathcal{S}} \pi_{t,s}\bar{G}_{t,s}\mathbb{I}\left[y_{t,N_B} \le \rho^{DA}_{t,s} \le y_{t,N_B+1}\right]}{\sum_{s\in\mathcal{S}} \pi_{t,s}\mathbb{I}\left[y_{t,N_B} \le \rho^{DA}_{t,s} \le y_{t,N_B+1}\right]} \tag{4.45}$$

    if $\max_{s\in\mathcal{S}} \rho^{DA}_{t,s} < O^G_t$ **and** $\rho^{high}_t > O^G_t$ **and** $\bar{G}^{high}_t > 0$ **then** $\{Case\ 2\}$
      Update $y_{N_B} = O^G_t$
      Insert $(\rho^{high}_t, \min\{\bar{G}^{high}_t + [x_{t,N_B}]_+,\ poi\})$ into the hour t bidding curve
    else $\{Case\ 1\}$
      Update $y_{N_B} = \rho^{high}_t$
    end if
  end if
end for

Algorithm 4.5 Bid curve extension for low-price events

Require: Model inputs and results: $\{\{\pi_{t,s}, gen_{t,s}, \rho^{DA}_{t,s}, C^{SC}_{t,s}, E^{DA}_{t,s}\}_{s\in\mathcal{S}}, \{x_{t,b}, y_{t,b}\}_{b\in\mathcal{B}}, E^{cha,max}_t, O^G_t\}_{t\in\mathcal{T}}$
Require: Maximum price considered a low-price event: $\bar{y}^{low}$
Require: Factor defining a high-price event relative to the time horizon: $\alpha > 1$

for $t \in \mathcal{T}$ **do**
  if $x_{t,1} > -\psi E^{cha,max}_t$ **then**
    Calculate the low-price event threshold:

$$\rho^{low}_t = \min\left\{\bar{y}^{low}, \alpha \cdot \left(\min_{s\in\mathcal{S},t\in\mathcal{T}} \rho^{DA}_{t,s}\right), \frac{1}{\alpha} \cdot \left(\min_{s\in\mathcal{S},t\in\mathcal{T}} \rho^{DA}_{t,s}\right)\right\} \tag{4.46}$$

    if $\min_{s\in\mathcal{S}} \rho^{DA}_{t,s} \geq O^G_t$ **and** $x_{t,1} \geq 0$ **then**
      if $\rho^{low}_t \geq O^G_t$ **then**
        Insert $(O^G_t,\ -\psi E^{cha,max}_t)$ into the hour t bidding curve
        Calculate the expected curtailment if day-ahead dispatch is less than 0:

$$\bar{G}^{low}_t = \frac{\sum_{s\in\mathcal{S}} \pi_{t,s} \cdot \left[gen_{t,s} - \min\{E^{cha,max}_t, C^{max}_{t,s} - C^{SC}_{t,s}\}\right]_+}{\sum_{s\in\mathcal{S}} \pi_{t,s}} \tag{4.47}$$

        if $0 < \bar{G}^{low}_t < x_{t,1}$ **then** $\{Case\ 2\}$
          Insert $(\rho^{low}_t,\ \bar{G}^{low}_t)$ into the hour t bidding curve
        else if $\bar{G}^{low}_t = 0$ **then**
          Calculate the expected rate of grid charging if ESR charges maximally:

$$x^{low}_t = \left(\sum_{s\in\mathcal{S}} \pi_{t,s}\right)^{-1} \left(\sum_{s\in\mathcal{S}} \pi_{t,s} \cdot \left[gen_{t,s} - E^{cha,max}_t\right]_+\right) \tag{4.48}$$

          Insert $(\rho^{low}_t,\ \psi \cdot x^{low}_t)$ into the hour t bidding curve
        end if
      else
        Calculate expected ESR output for scenarios on the lowest-price segment:

$$E^{low}_t = \frac{\sum_{s\in\mathcal{S}} \cdot \pi_{t,s} \cdot E^{DA}_{t,s} \cdot \mathbb{I}\left[y_{t,1} \leq \rho^{DA}_{t,s} \leq y_{t,2}\right]}{\sum_{s\in\mathcal{S}} \pi_{t,s} \cdot \mathbb{I}\left[y_{t,1} \leq \rho^{DA}_{t,s} \leq y_{t,2}\right]} \tag{4.49}$$

        Insert $(O^G_t, [\min\{x_{t,1}, E^{low}_t\}]_+)$, $(\rho^{low}_t, -\psi E^{cha,max}_t)$ into the hour t bidding curve
      end if
    else $\{Case\ 1\}$
      Insert $(\rho^{low}_t,\ -\psi E^{cha,max}_t)$ into the hour t bidding curve
    end if
  end if
end for

Bibliography

[1] Khalid Abdulla et al. "Optimal Operation of Energy Storage Systems Considering Forecasts and Battery Degradation". In: *IEEE Transactions on Smart Grid* 9.3 (May 2018), pp. 2086–2096.

[2] Xin Ai et al. "Robust operation strategy enabling a combined wind/battery power plant for providing energy and frequency ancillary services". In: *International Journal of Electrical Power & Energy Systems* 118 (June 2020), p. 105736.

[3] Alireza Akbari-Dibavar, Kazem Zare, and Sayyad Nojavan. "A hybrid stochastic-robust optimization approach for energy storage arbitrage in day-ahead and real-time markets". In: *Sustainable Cities and Society* 49 (Aug. 2019), p. 101600.

[4] Hossein Akhavan-Hejazi and Hamed Mohsenian-Rad. "Optimal Operation of Independent Storage Systems in Energy and Reserve Markets With High Wind Penetration". In: *IEEE Transactions on Smart Grid* 5.2 (Mar. 2014), pp. 1088–1097.

[5] M. H. Amini, J. Mohammadi, and S. Kar. "Distributed Holistic Framework for Smart City Infrastructures: Tale of Interdependent Electrified Transportation Network and Power Grid". In: *IEEE Access* 7 (2019), pp. 157535–157554.

[6] M Salman Asif and Justin Romberg. "Sparse recovery of streaming signals using l1-homotopy". In: *IEEE Trans. on Signal Processing* 62.16 (2014), pp. 4209–4223.

[7] Ahmad Attarha, Nima Amjady, and Shahab Dehghan. "Affinely Adjustable Robust Bidding Strategy for a Solar Plant Paired With a Battery Storage". In: *IEEE Transactions on Smart Grid* 10.3 (May 2019), pp. 2629–2640.

[8] Ross Baldick. *Applied optimization: formulation and algorithms for engineering systems*. Cambridge University Press, 2006.

[9] Dimitri Bertsekas. *Nonlinear Programming*. Third. Belmont, Mass.: Athena Scientific, 2016, pp. 35–355, 770.

[10] Omar Besbes, Yonatan Gur, and Assaf Zeevi. "Non-stationary stochastic optimization". In: *Operations Research* 63.5 (2015), pp. 1227–1244.

[11] Mark Bolinger et al. *Hybrid Power Plants: Status of Operating and Proposed Plants*. URL: `https://emp.lbl.gov/hybrid`.

[12] Saverio Bolognani, Guido Cavraro, and Sandro Zampieri. "A Distributed Feedback Control Approach to the Optimal Reactive Power Flow Problem". In: *Control of Cyber-Physical Systems: Workshop held at Johns Hopkins University, March 2013*. Ed. by Danielle C. Tarraf. Heidelberg: Springer International Publishing, 2013, pp. 259–277.

[13] W. A. Bukhsh et al. "Local Solutions of the Optimal Power Flow Problem". In: *IEEE Transactions on Power Systems* 28.4 (Nov. 2013), pp. 4780–4788.

[14] Raymond H. Byrne et al. "Opportunities for Energy Storage in CAISO: Day-Ahead and Real-Time Market Arbitrage". In: *2018 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*. June 2018, pp. 63–68.

[15] *California ISO OASIS*. URL: `http://oasis.caiso.com/`.

[16] E. J. Candes and Y. Plan. "Matrix completion with noise". In: *Proc. of the IEEE* 98.6 (2010), pp. 925–936.

[17] E. J. Candes and Ben Recht. "Exact matrix completion via convex optimization". In: *Foundations of Computational mathematics* 9.6 (2009), pp. 717–772.

[18] Thomas Carriere et al. "Strategies for combined operation of PV/storage systems integrated into electricity markets". In: *IET Renewable Power Generation* 14.1 (May 2019), pp. 71–79.

[19] Ting-Jui Chang and Shahin Shahrampour. "On Online Optimization: Dynamic Regret Analysis of Strongly Convex and Smooth Problems". In: *AAAI Conference on Artificial Intelligence (AAAI)*. Vol. 35. 2021.

[20] Jie Chen and Ronny Luss. "Stochastic gradient descent with biased but consistent gradient estimators". In: *arXiv preprint arXiv:1807.11880* (2018).

[21] Bolong Cheng and Warren B. Powell. "Co-Optimizing Battery Storage for the Frequency Regulation and Energy Arbitrage Using Multi-Scale Dynamic Programming". In: *IEEE Transactions on Smart Grid* 9.3 (May 2018), pp. 1997–2005.

[22] Carleton Coffrin, Hassan L Hijazi, and Pascal Van Hentenryck. "The QC relaxation: A theoretical and computational study on optimal power flow". In: *IEEE Transactions on Power Systems* 31.4 (2015), pp. 3008–3018.

[23] A.L. Costa and A. Simões Costa. "Energy and ancillary service dispatch through dynamic optimal power flow". In: *Electric Power Systems Research* 77.8 (2007), pp. 1047–1055.

[24] Claude Crampes and Jean-Michel Trochet. "Economics of stationary electricity storage with various charge and discharge durations". In: *Journal of Energy Storage* 24 (Aug. 2019), p. 100746.

[25] James R. Cruise, Lisa Flatley, and Stan Zachary. "Impact of storage competition on energy markets". In: *European Journal of Operational Research* 269.3 (Sept. 2018), pp. 998–1012.

[26] E. Dall'Anese and A. Simonetto. "Optimal Power Flow Pursuit". In: *IEEE Transactions on Smart Grid* 9.2 (Mar. 2018), pp. 942–952.

[27] Kaushik Das et al. "Optimal battery operation for revenue maximization of wind-storage hybrid power plant". In: *Electric Power Systems Research* 189 (Dec. 2020), p. 106631.

[28] Maria Dicorato et al. "Planning and Operating Combined Wind-Storage System in Electricity Market". In: *IEEE Transactions on Sustainable Energy* 3.2 (Apr. 2012), pp. 209–217.

[29] Huajie Ding et al. "Optimal Offering and Operating Strategy for a Large Wind-Storage System as a Price Maker". In: *IEEE Transactions on Power Systems* 32.6 (Nov. 2017), pp. 4904–4913.

[30] Yuhao Ding, Javad Lavaei, and Murat Arcak. "Escaping spurious local minimum trajectories in online time-varying nonconvex optimization". In: *American Control Conference* (2021).

[31] C. B. Do, Q. V. Le, and C. S. Foo. "Proximal regularization for online and batch learning". In: *ICML* (2009).

[32] Asen L Dontchev and R Tyrrell Rockafellar. *Implicit functions and solution mappings*. Vol. 543. Springer, 2009.

[33] Yury Dvorkin et al. "Comparison of scenario reduction techniques for the stochastic unit commitment". In: *2014 IEEE PES General Meeting*. 2014, pp. 1–5.

[34] S Fattahi et al. *Absence of spurious local trajectories in time-varying optimization*. 2019. URL: `https://lavaei.ieor.berkeley.edu/Time_Varing_2019_1.pdf`.

[35] Salar Fattahi et al. "Absence of Spurious Local Trajectories in Time-Varying Optimization: A Control-Theoretic Perspective". In: (2020). IEEE Conference on Control Technology and Applications (CCTA).

[36] Maryam Fazel et al. "Global convergence of policy gradient methods for the linear quadratic regulator". In: *International Conference on Machine Learning*. 2018.

[37] *FERC Order No. 841*. Dec. 2020. URL: `https://www.ferc.gov/media/order-no-841`.

[38] L. Gan and S. H. Low. "An Online Gradient Algorithm for Optimal Power Flow on Radial Networks". In: *IEEE Journal on Selected Areas in Communications* 34.3 (Mar. 2016), pp. 625–638.

[39] Xiand Gao, Xiaobo Li, and Shuzhong Zhang. "Online learning with non-convex losses and non-stationary regret". In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2018, pp. 235–243.

[40] Javier Garcia-Gonzalez et al. "Stochastic Joint Optimization of Wind Generation and Pumped-Storage Units in an Electricity Market". In: *IEEE Transactions on Power Systems* 23.2 (May 2008), pp. 460–468.

[41] Sahand Ghavidel et al. "Risk-Constrained Bidding Strategy for a Joint Operation of Wind Power and CAES Aggregators". In: *IEEE Transactions on Sustainable Energy* 11.1 (Jan. 2020), pp. 457–466.

[42] S. Gill, I. Kockar, and G. W. Ault. "Dynamic Optimal Power Flow for Active Distribution Networks". In: *IEEE Transactions on Power Systems* 29.1 (Jan. 2014), pp. 121–131.

[43] I. L. R. Gomes et al. "Stochastic coordination of joint wind and photovoltaic systems with energy storage in day-ahead market". In: *Energy* 124 (Apr. 2017), pp. 310–320.

[44] Amaia Gonzalez-Garrido et al. "Annual Optimized Bidding and Operation Strategy in Energy and Secondary Reserve Markets for Solar Plants With Storage Systems". In: *IEEE Transactions on Power Systems* 34.6 (Nov. 2019), pp. 5115–5124.

[45] A. Gopalakrishnan et al. "Global optimization of multi-period optimal power flow". In: *2013 American Control Conference*. June 2013, pp. 1157–1164.

[46] Will Gorman et al. "Motivations and options for deploying hybrid generator-plus-battery projects within the bulk power system". In: *The Electricity Journal* 33.5 (2020), p. 106739.

[47] N. Growe-Kuska, H. Heitsch, and W. Romisch. "Scenario reduction and scenario tree construction for power management problems". In: *2003 IEEE Bologna Power Tech Conference Proceedings,* vol. 3. June 2003, 7 pp. Vol.3.

[48] Eric C. Hall and Rebecca M. Willett. "Online Convex Optimization in Dynamic Environments". In: *IEEE Jrnl. of Selected Topics in Signal Processing* 9.4 (2015), pp. 647–662.

[49] Adrian Hauswirth, Saverio Bolognani, and Florian Dörfler. "Projected dynamical systems on irregular, non-Euclidean domains for nonlinear optimization". In: *SIAM Journal on Control and Optimization* 59.1 (2021).

[50] Elad Hazan. *Introduction to Online Convex Optimization*. Vol. 2. 4. Foundations and Trends in Optimization, 2015, pp. 157–325.

[51] Elad Hazan and Satyen Kale. "Online Submodular Minimization." In: *Journal of Machine Learning Research* 13.10 (2012).

[52] Guannan He et al. "Optimal Bidding Strategy of Battery Storage in Power Markets Considering Performance-Based Regulation and Battery Cycle Life". In: *IEEE Transactions on Smart Grid* 7.5 (Sept. 2016), pp. 2359–2367.

[53] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Convex analysis and minimization algorithms I: Fundamentals*. Vol. 305. Springer science & business media, 2013.

[54] *Hybrid Resources White Paper*. May 2021. URL: https://www.ferc.gov/media/hybrid-resources-white-paper.

[55] California ISO. *Final Flexible Capacity Needs Assessment for 2023*. May 2022. URL: http://www.caiso.com/InitiativeDocuments/Final2023FlexibleCapacityNeedsAssessment.pdf.

[56] Midwest ISO. *Production Cost Model Fundamentals*. URL: https://www.engineering.iastate.edu/~jdm/ee590-Old/ProductionCostModleFundamentals_EE590.pdf.

[57] Ali Jadbabaie, Shahin Shahrampour Alexander Rakhlin, and Karthik Sridharan. "Online Optimization: Competing with Dynamic Comparators". In: *Artificial Intelligence and Statistics, Proc. of Machine Learning Research* 38 (2015), pp. 398–406.

[58] Ali Jamali et al. "Self-scheduling approach to coordinating wind power producers with energy storage and demand response". In: *IEEE Transactions on Sustainable Energy* 11.3 (2019), pp. 1210–1219.

[59] George F. Jenks. *Optimal data classification for choropleth maps*. University of Kansas, 1977.

[60] Daniel R. Jiang and Warren B. Powell. "Optimal Hour-Ahead Bidding in the Real-Time Electricity Market with Battery Storage Using Approximate Dynamic Programming". In: *INFORMS Journal on Computing* 27.3 (Aug. 2015), pp. 525–543.

[61] Cédric Josz et al. "Application of the moment-SOS approach to global optimization of the OPF problem". In: *IEEE Transactions on Power Systems* 30.1 (2014), pp. 463–470.

[62] Hamed Karimi, Julie Nutini, and Mark Schmidt. "Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition". In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2016, pp. 795–811.

[63] Jalal Kazempour and Benjamin F Hobbs. "Value of flexible resources, virtual bidding, and self-scheduling in two-settlement electricity markets with wind generation—part I: principles and competitive model". In: *IEEE Transactions on Power Systems* 33.1 (2017), pp. 749–759.

[64] Muhammad Khalid et al. "On maximizing profit of wind-battery supported power station based on wind power and energy price forecasting". In: *Applied Energy* 211 (Feb. 2018), pp. 764–773.

[65] Robert Kleinberg, Yuanzhi Li, and Yang Yuan. "An alternative view: When does SGD escape local minima?" In: *arXiv preprint arXiv:1802.06175* (2018).

[66] Burak Kocuk, Santanu S Dey, and X Andy Sun. "Strong SOCP relaxations for the optimal power flow problem". In: *Operations Research* 64.6 (2016), pp. 1177–1196.

[67] Magnus Korpaas, Arne T. Holen, and Ragne Hildrum. "Operation and sizing of energy storage for wind power plants in a market system". In: *International Journal of Electrical Power & Energy Systems*. 14th Power Systems Computation Conference 25.8 (Oct. 2003), pp. 599–606.

[68] D. Kourounis, A. Fuchs, and O. Schenk. "Toward the Next Generation of Multiperiod Optimal Power Flow Solvers". In: *IEEE Transactions on Power Systems* 33.4 (2018), pp. 4005–4014.

[69] Dheepak Krishnamurthy et al. "Energy Storage Arbitrage Under Day-Ahead and Real-Time Price Uncertainty". In: *IEEE Transactions on Power Systems* 33.1 (Jan. 2018), pp. 84–93.

[70] Alexander Y Kruger et al. "Hölder error bounds and Hölder calmness with applications to convex semi-infinite optimization". In: *Set-Valued and Variational Analysis* 27.4 (2019), pp. 995–1023.

[71] Ranjeet Kumar et al. "Benchmarking stochastic and deterministic MPC: A case study in stationary battery systems". In: *AIChE Journal* 65.7 (2019), e16551.

[72] Bernd Kummer. "Inclusions in general spaces: Hoelder stability, solution schemes and Ekeland's principle". In: *Journal of mathematical analysis and applications* 358.2 (2009), pp. 327–344.

[73] Javad Lavaei and Steven H Low. "Zero duality gap in optimal power flow problem". In: *IEEE Transactions on Power Systems* 27.1 (2012), pp. 92–107.

[74] Luke Lavin, Ningkun Zheng, and Jay Apt. *Market Power Challenges and Solutions for Electric Power Storage Resources*. Working Paper CEIC-21-02. Pittsburg, PA: Carnegie Mellon Electricity Industry Center, 2021. URL: `https://www.cmu.edu/ceic/research-publications/ceic_21_02-esr-policy.pdf`.

[75] A. Lesage-Landry, J. A. Taylor, and I. Shames. "Second-order Online Nonconvex Optimization". In: *IEEE Trans. on Automatic Control* (2020).

[76] H. Li et al. "Visualizing the loss landscape of neural nets". In: *Advances in Neural Information Processing Systems* (2018), pp. 6389–6399.

[77] Jinghua Li, Fei Lan, and Hua Wei. "A scenario optimal reduction method for wind power time series". In: *IEEE Transactions on Power Systems* 31.2 (2015), pp. 1657–1658.

[78] Nan Li et al. "Flexible Operation of Batteries in Power System Scheduling With Renewable Energy". In: *IEEE Transactions on Sustainable Energy* 7.2 (Apr. 2016), pp. 685–696.

[79] Xiuxian Li, Xinlei Yi, and Lihua Xie. "Distributed Online Convex Optimization with an Aggregative Variable". In: *IEEE Transactions on Control of Network Systems* (2021), pp. 1–1.

[80] Guodong Liu, Yan Xu, and Kevin Tomsovic. "Bidding Strategy for Microgrid in Day-Ahead Market Based on Hybrid Stochastic/Robust Optimization". In: *IEEE Transactions on Smart Grid* 7.1 (Jan. 2016), pp. 227–237.

[81] Nils Löhndorf, David Wozabal, and Stefan Minner. "Optimizing Trading Decisions for Hydro Storage Systems Using Approximate Dual Dynamic Programming". In: *Operations Research* 61.4 (Aug. 2013), pp. 810–823.

[82] Fengji Luo et al. "Coordinated Operational Planning for Wind Farm With Battery Energy Storage System". In: *IEEE Transactions on Sustainable Energy* 6.1 (Jan. 2015), pp. 253–262.

[83] Odalric-Ambrym Maillard and Rémi Munos. "Online learning in adversarial lipschitz environments". In: *Joint european conference on machine learning and knowledge discovery in databases*. Springer. 2010, pp. 305–320.

[84] Robert Małkowski, Marcin Jaskólski, and Wojciech Pawlicki. "Operation of the Hybrid Photovoltaic-Battery System on the Electricity Market—Simulation, Real-Time Tests and Cost Analysis". In: *Energies* 13.6 (Jan. 2020), p. 1402.

[85] Vincenzo Marano, Gianfranco Rizzo, and Francesco Antonio Tiano. "Application of dynamic programming to the optimal management of a hybrid power plant with wind turbines, photovoltaic panels and compressed air energy storage". In: *Applied Energy* 97 (2012), pp. 849–859.

[86] Olivier Massicot and Jakub Marecek. "On-line Non-Convex Constrained Optimization". In: *arXiv preprint arXiv:1909.07492* (2019).

[87] Madeleine McPherson et al. "Impacts of storage dispatch on revenue in electricity markets". In: *Journal of Energy Storage* 31 (Oct. 2020), p. 101573.

[88] Dennis van der Meer, Guang Chao Wang, and Joakim Munkhammar. "An alternative optimal strategy for stochastic model predictive control of a residential battery energy management system with solar photovoltaic". In: *Applied Energy* 283 (Feb. 2021), p. 116289.

[89] Hamed Mohsenian-Rad. "Optimal Bidding, Scheduling, and Deployment of Battery Systems in California Day-Ahead Energy Market". In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016). Conference Name: IEEE Transactions on Power Systems, pp. 442–453.

[90] Aryan Mokhtari et al. "Online optimization in dynamic environments: Improved regret rates for strongly convex problems". In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE. 2016, pp. 7195–7201.

[91] James A. Momoh. *Electric Power System Applications of Optimization*. Boca Raton: CRC Press, 2017.

[92] Anna Nagurney and Ding Zhang. *Projected dynamical systems and variational inequalities with applications*. Vol. 2. Springer Science & Business Media, 2012.

[93] Ehsan Nasrolahpour et al. "A bilevel model for participation of a storage system in energy and reserve markets". In: *IEEE Transactions on Sustainable Energy* 9.2 (2017), pp. 582–598.

[94] Amparo Núñez-Reyes et al. "Optimal scheduling of grid-connected PV plants with energy storage for integration in the electricity market". In: *Solar Energy* 144 (Mar. 2017), pp. 502–516.

[95] Anthony Papavasiliou and Shmuel S. Oren. "Multiarea Stochastic Unit Commitment for High Wind Penetration in a Transmission Constrained Network". In: *Operations Research* 61.3 (2013), pp. 578–592.

[96] SangWoo Park et al. "Diminishing Regret for Online Nonconvex Optimization". In: *2021 American Control Conference (ACC)*. 2021, pp. 978–985.

[97] B.T. Polyak. "Gradient methods for minimizing functionals". In: *Zh. Vychisl. Mat. Mat. Fiz.* 3.4 (1963), pp. 643–653.

[98] H. M. I. Pousinho et al. "Robust optimisation for self-scheduling and bidding strategies of hybrid CSP–fossil power plants". In: *International Journal of Electrical Power & Energy Systems* 67 (May 2015), pp. 639–650.

[99] Morteza Rahimiyan and Luis Baringo. "Strategic Bidding for a Virtual Power Plant in the Day-Ahead and Real-Time Markets: A Price-Taker Robust Optimization Approach". In: *IEEE Transactions on Power Systems* 31.4 (July 2016), pp. 2676–2687.

[100] Joseph Rand et al. *Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2021*. URL: `https://emp.lbl.gov/queues`.

[101] Christopher V. Rao, James B. Rawlings, and David Q. Mayne. "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations". In: *IEEE Trans. on Automatic Control* 48.2 (2003).

[102] R Tyrrell Rockafellar, Stanislav Uryasev, et al. "Optimization of conditional value-at-risk". In: *Journal of risk* 2 (2000), pp. 21–42.

[103] Ernest K Ryu and Stephen Boyd. "Stochastic proximal iteration: a non-asymptotic improvement upon stochastic gradient descent". In: (2017). URL: `http://stanford.edu/~boyd/papers/spi.html`.

[104] Agustín A. Sánchez de la Nieta, Javier Contreras, and José Ignacio Muñoz. "Optimal coordinated wind-hydro bidding strategies in day-ahead markets". In: *IEEE Transactions on Power Systems* 28.2 (May 2013), pp. 798–809.

[105] Soroush Shafiee et al. "Risk-constrained bidding and offering strategy for a merchant compressed air energy storage plant". In: *IEEE Transactions on Power Systems* 32.2 (2016), pp. 946–957.

[106] Alexander Shapiro. "Monte Carlo Sampling Methods". In: *Stochastic Programming*. Vol. 10. Handbooks in Operations Research and Management Science. Elsevier, 2003, pp. 353–425.

[107] Aditi Sheshadri et al. "Midlatitude error growth in atmospheric GCMs: the role of eddy growth rate". In: *Geophysical Research Letters* 48.23 (2021), e2021GL096126.

[108] Somayeh Sojoudi and Javad Lavaei. "Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure". In: *SIAM Journal on Optimization* 24.4 (2014), pp. 1746–1778.

[109] Gregory Steeger, Luiz Augusto Barroso, and Steffen Rebennack. "Optimal Bidding Strategies for Hydro-Electric Producers: A Literature Survey". In: *IEEE Transactions on Power Systems* 29.4 (July 2014), pp. 1758–1766.

[110] Y. Tang, K. Dvijotham, and S. Low. "Real-Time Optimal Power Flow". In: *IEEE Transactions on Smart Grid* 8.6 (2017), pp. 2963–2973.

[111] Yujie Tang. "Time-Varying Optimization and Its Application to Power System Operation". PhD thesis. California Institute of Technology, 2019.

[112] Yujie Tang, Krishnamurthy Dvijotham, and Steven Low. "Real-Time Optimal Power Flow". In: *IEEE Trans. on Smart Grid* 8.6 (2017), pp. 2963–2973.

[113] Anupam A. Thatte et al. "Risk Measure Based Robust Bidding Strategy for Arbitrage Using a Wind Farm and Energy Storage". In: *IEEE Transactions on Smart Grid* 4.4 (Dec. 2013), pp. 2191–2199.

[114] Sadegh Vejdan and Santiago Grijalva. "Maximizing the Revenue of Energy Storage Participants in Day-Ahead and Real-Time Markets". In: *2018 Clemson University Power Systems Conference (PSC)*. Sept. 2018, pp. 1–6.

[115] Yishen Wang, Yuzong Liu, and Daniel S. Kirschen. "Scenario Reduction With Submodular Optimization". In: *IEEE Transactions on Power Systems* 32.3 (2017), pp. 2479–2480.

[116] Yishen Wang et al. "Look-Ahead Bidding Strategy for Energy Storage". In: *IEEE Transactions on Sustainable Energy* 8.3 (July 2017), pp. 1106–1117.

[117] Ming Wei and Jin Zhong. "Optimal bidding strategy for demand response aggregator in day-ahead markets via stochastic programming and robust optimization". In: *2015 12th International Conference on the European Energy Market (EEM)*. May 2015, pp. 1–5.

[118] Rafał Weron. "Electricity price forecasting: A review of the state-of-the-art with a look into the future". In: *International Journal of Forecasting* 30.4 (Oct. 2014), pp. 1030–1081.

[119] Yunyun Xie et al. "Robust MPC-based bidding strategy for wind storage systems in real-time energy and regulation markets". In: *International Journal of Electrical Power & Energy Systems* 124 (Jan. 2021), p. 106361.

[120] Bolun Xu, Magnus Korpas, and Audun Botterud. "Operational Valuation of Energy Storage under Multi-stage Price Uncertainties". In: *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, Dec. 2020, pp. 55–60.

[121] Xiao Xu et al. "Scheduling of wind-battery hybrid system in the electricity market using distributionally robust optimization". In: *Renewable Energy* 156 (Aug. 2020), pp. 47–56.

[122] L. Yang et al. "An Optimal Algorithm for Online Non-Convex Learning". In: *Proc. of the ACM on Measurement and Analysis of Computing Syst.* 25 (2018).

[123] Tianbao Yang et al. "Tracking Slowly Moving Clairvoyant: Optimal Dynamic Regret of Online Learning with True and Noisy Gradient". In: *Proceedings of The 33rd International Conference on Machine Learning*. 2016.

[124] Zhang Yang, Li Ce, and Li Lian. "Electricity price forecasting by a hybrid model, combining wavelet transform, ARMA and kernel-based extreme learning machine methods". In: *Applied Energy* 190 (2017), pp. 291–305.

[125] Sen Zhan et al. "Co-optimized trading of hybrid wind power plant with retired EV batteries in energy and reserve markets under uncertainties". In: *International Journal of Electrical Power & Energy Systems* 117 (May 2020), p. 105631.

[126] Lijun Zhang et al. "Improved Dynamic Regret for Non-degenerate Functions". In: *Advances in Neural Information Processing Systems*. Vol. 30. 2017.

[127] Lijun Zhang et al. "Online bandit learning for a special class of non-convex losses". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 29. 1. 2015.

[128] Richard Y Zhang, Javad Lavaei, and Ross Baldick. "Spurious local minima in power system state estimation". In: *IEEE Transactions on Control of Network Systems* 6.3 (2019), pp. 1086–1096.

[129] Richard Y Zhang, Somayeh Sojoudi, and Javad Lavaei. "Sharp restricted isometry bounds for the inexistence of spurious local minima in nonconvex matrix recovery". In: *Journal of Machine Learning research* (2019).

[130] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas. "MATPOWER: Steady-State Operations, Planning, and Analysis Tools for Power Systems Research and Education". In: *IEEE Transactions on Power Systems* 26.1 (Feb. 2011), pp. 12–19.

[131] Martin Zinkevich. "Online convex programming and generalized infinitesimal gradient ascent". In: *Proc. of the Int. Conf. on Machine Learning* (2003), pp. 928–936.

[132] Fariba Zohrizadeh et al. "A survey on conic relaxations of optimal power flow problem". In: *European Journal of Operational Research* 287.2 (2020), pp. 391–409.