Transitioning to Inverter-based Resources: Stability of AC Power Systems considering Current Constraints, Small-signal Sensitivities, and Large Loads

Trager Joswig-Jones

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

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Baosen Zhang, Chair

Daniel Kirschen

Jing Yu

Program Authorized to Offer Degree:
Electrical and Computer Engineering

University of Washington

Abstract

Transitioning to Inverter-based Resources:
Stability of AC Power Systems considering Current Constraints,
Small-signal Sensitivities, and Large Loads

Trager Joswig-Jones

Chair of the Supervisory Committee:
Keith & Nancy Rattie Endowed Professor Baosen Zhang
Electrical and Computer Engineering

The large-scale integration of inverter-based resources (IBRs) has fundamentally altered the dynamics of alternating-current (AC) power systems. Traditionally founded on the physical dynamics of synchronous machines, the grid is transitioning toward a power-electronics-dominated system driven by renewable generation and the proliferation of inverter-based loads from data centers and electrification efforts. While digitally controlled IBRs provide programmable flexibility, they are bound by rigid hardware constraints that limit their capabilities. Ensuring grid stability now requires generalized approaches that move beyond synchronous-machine-centric models to account for the unique dynamics of power electronics. This dissertation investigates IBR control and stability across three critical domains: individual device constraints, system-wide network stability, and large load regulation.

First, we address the nonlinear control challenges posed by inverter current magnitude limits, which are essential for protecting semiconductor devices, but inhibit IBRs ability to provide a strong voltage source response during transients. We characterize the geometry of the feasible output of current-limited inverters and derive Lyapunov stability conditions for devices with direct current limiters. We then propose a safety filter approach that minimally modifies desired control actions to maintain safety, and finally propose a safe trajectory gradient flow controller that iteratively optimizes the future trajectory to remain within a defined safe set.

Second, we examine the impact of IBRs on system-wide small-signal stability. Leveraging equivalent admittance modeling of systems, we propose a sensitivity-based metric that quantifies how perturbations to power flow parameters, such as bus power injections or network admittances, impact system eigenvalues. We demonstrate the insights these metrics provide for identifying and mitigating stability challenges in high-IBR networks.

Lastly, we investigate leveraging IBRs within data centers to mitigate load volatility. We demonstrate how reactive power control and energy storage dispatch can be optimized for voltage regulation, reducing fluctuations and violations caused by the load variations typical of LLM training tasks.

Collectively, this work provides an overview of a range of stability-related challenges that must be considered when integrating IBRs into modern, large-scale power systems.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I wish to express my sincere appreciation to Professor Baosen Zhang. I'm grateful to have been taken on as his student and to have had the opportunity to learn how to formulate and solve challenging and interesting research problems with his support. I will always remember the advice he has given me throughout the years: figure out what assumptions we can make to simplify the problem enough to get an insightful analytical solution, think about what result I expect to see before trying something, and before trying to prove something, to try it a few thousand times with random inputs – then if it still works we can try to prove it. Lastly, I am thankful for the space and freedom I was provided to learn to work independently on research problems, and manage myself.

I am also incredibly thankful to Professor Daniel Kirschen, the first professor I had for a power systems course, as well as for many following courses. He has been a part of my academic studies from nearly the start of my undergraduate degree and has provided me opportunities and support along every step of the way through to my dissertation defense. I'll always remember these two takeaways from his lectures: Reactive power is like the foam on top of beer, and power system operators in the UK have to prepare the grid for the surge of electric kettle usage during commercial breaks of popular shows.

I extend my sincere appreciation to my two other committee members, Jing Yu, and Chaoyue Zhao. I would also like to thank some of the many other UW faculty members who positively impacted my education during my long term here at UW through their teaching and feedback on my work: Rich Christie, Samuel Burden, June Lukuyu, Jungwong Choi, Cody Schlenker, Alex Mamishev, and Brian Johnson.

Next, I would like to thank the other mentors that I have had throughout the years: Marcos Donolo, Samuel Bhushan, Zafar (ZAK) Khan, Gord Stephen, Ahmed Zamzam, Kyri Baker, Shuan Dong, and Jin Tan. I learned so much through my internships and research projects that I would not have had the opportunity to pursue without all of these wonderful people. I would like to emphasize my gratitude to Gord Stephen; it was through my work with him that I really discovered my passion for power systems research which led me to apply to graduate school.

I am very grateful for all of the wonderful colleagues I've had throughout the years. My labmates – Wenqi Cui, Daniel Tabas, Jiayi Li, Ling Zhang, Matt Motoki, Lauren Streitmatter, Jiayi Yao, Lila Perkins, Margaret Capetz, Yan Jiang, Zixiao Ma, Jerry Anunrojwong, Weiqian Cai, Soham Dutta, Rahul Mallik, Pranav Chandran, Minghui Lu, Nimesh Vamanan – fellow PhD students – Ahana Mukherjee, Kwame Donkar, Eliane Nirere, Aya Alayli, Miquilina Anagbah, Hongyu Yi, Jesse Wang – our visiting students – Jiarui Long, Yangze Zhou, Jochen Stiasny – the Crags setting crew – Cal, Syd, Oliver, Van, Eva, Lila, Hailey, Troy, Nic, Austin, Ian, and James – and lab partner – Ishaan Bhimani. I have been fortunate to work alongside a group of great people in every facet of my academic career.

I am grateful for all the staff who have supported me and the programs I have participated in throughout my journey at UW: Stephanie Swanson (my first student advisor and primary source of candy left out front of her office), Danica Hendricks (who made outreach through CEI a wonderful experience), Jeremy Cox (an excellent manager and overseer of the Crags Climbing Gym), Jessi Allen-Navarre, Ary Prasetyowati, Tyler Pippins, Anthony Pumilia, Mack Carter, Jennifer Huberman, and all the other staff that keep UW running.

I am wholly indebted to my family and friends who have patiently supported me and have made Washington State so much more than just a place to live. Washington State has been my home for the past 27 years, and I am grateful to have been able to appreciate the beautiful varied landscape this state resides upon and formally complete my education here before graduating into the real world.

Lastly, a special thank you to the jumping spider who lived on my computer monitors in my home office and kept the fruit flies at bay during the long work days of the 5th year of my Ph.D.

DEDICATION

to my parents, Carolyn and Jeff, and maja kachanaja, Angelina

Chapter 1
INTRODUCTION

The development of high-efficiency semiconductor switches and microprocessor-based digital controllers made power electronic converters viable for wide-spread deployment in modern electrical systems. This combined with advances in power electronic dependent energy technologies, such as solar photovoltaics, wind turbines, and battery energy storage systems, has led to the proliferation of grid-interfacing direct current (DC) to alternating current (AC) inverters in AC power systems. The adoption of these Inverter-Based Resources (IBRs) has caused a renaissance in power system control and stability, which had previously stagnated due to the fact that many challenges related to conventional synchronous generation had been considered solved. Of particular interest at the moment is the control and stability of power systems with high levels of IBRs. These devices are fast and flexible in how we can control them; They are simply a set of switches controlled with a microcontroller. However, they are inflexible in other ways due to the physical limitations of the hardware used in these devices and uncertainty in power availability from the resource behind the inverter. This flexibility in control and rigidity in constraints make integrating IBRs into the grid an interesting optimization and control problem.

1.1 Motivation

Electrical power grids worldwide were built around the generation and transmission of AC electric power; A decision shaped by the capabilities of AC synchronous generators and the ability to increase or decrease AC voltages efficiently using transformers[1] [1]. High-voltage AC transmission enabled long-distance power delivery with lower losses, while three-phase AC systems reduced conductor material and improved power quality. These AC system networks provide power at nominal voltage values to many loads that are connected in parallel across the network. Around the world, AC wide-area synchronous interconnections continue to enable the economical delivery of electric power to every corner of our modern lives.

At the heart of these wide-area interconnections is the physical reliance on the rotating mass of synchronous generators. The inherent inertia of synchronous generator's rotors provides a critical link between power balance and frequency; any mismatch between generation and load causes the rotors to accelerate or decelerate, directly manifesting as frequency deviations. This inertia is derived from the kinetic energy stored within the rotating mass, which acts as a physical buffer against power imbalances. A larger rotating mass stores more kinetic energy, effectively slowing the rate of change of frequency following a disturbance and providing critical time for system regulation. This relationship allows frequency to serve as a global, decentralized measurement of the grid's power balance, facilitating the stabilization and regulation of the system.

Today, however, the technological landscape of the electric power grid is shifting. Many increasingly prevalent energy technologies, such as solar photovoltaics, battery-energy storage, electric vehicles, and large computational loads, either produce, store, or operate on DC power or require power electronics to efficiently interface with an AC grid. The conversion between DC and AC power that enables these devices to interface with the grid is handled by power electronics devices called inverters[2]. As more inverters are interfaced with

---

[1] In the 1890s, no practical technology (e.g. power electronics) existed for stepping the voltage up and down in a DC system making the interconnection of systems that were separated by long distances infeasible.

[2] The conversion of AC to DC power is performed by rectifiers. However, inverter devices can also be controlled to rectify AC to DC, as well as convert DC to AC (i.e. consume or produce AC power).

the grid, the underlying physics of AC systems, that used to be dominated by synchronous generators, are changing.

1.1.1 Inverter-Based Resources

AC power systems were founded on conventional synchronous generators that generally combust fossil fuels, such as oil, coal, and natural gas to produce the power required to generate electricity. Recently, there has been a push to transition to renewable energy based resources, such as solar, wind, and waves, to reduce the cost of generating electricity as well as reduce the emission of greenhouse gases released when combusting fossil fuels[3]. Indeed, this transition is already well under way with one third of electrical energy world-wide being produced from renewable energy resources in 2024 [3]. Fig. 1.1 shows the rapid rate at which renewable energy technologies, particularly solar photovoltaics and wind, have increased and are projected to continue increasing their contribution of total electricity generation globally. Looking at the North American Western Interconnection, in 2024, entities retired 1.4 GW of generation capacity, most of which was coal (920 MW) and natural gas (413 MW), while 75% of the new generation (24 GW) added was IBRs [4].

While renewable energy resources can provide low cost power production, this power production is variable and uncertain based on the availability from the resource (e.g. wind turbines only produce power when the wind is strong enough.). To mitigate the challenge of uncertainty and variability in renewable power production, battery-energy storage has been seeing increased adoption ranging from home battery storage systems to grid-scale storage projects. Global battery storage additions reached 77 GW, in 2024 alone [3]; This is higher than in any previous year as shown in Fig 1.2 and demonstrates that these yearly additions have been growing in terms of total battery power capacity.

This growth in renewable generation and battery energy storage has been fueled by the decreasing costs of building wind and solar generation and battery energy storage systems over the past decade. When evaluated without subsidies, utility-scale solar and onshore wind are the most cost-effective forms of new-build energy technologies in the United States. As of

[3]The combustion of fossil fuels produces greenhouse gases, which contribute to anthropogenic climate change [2].

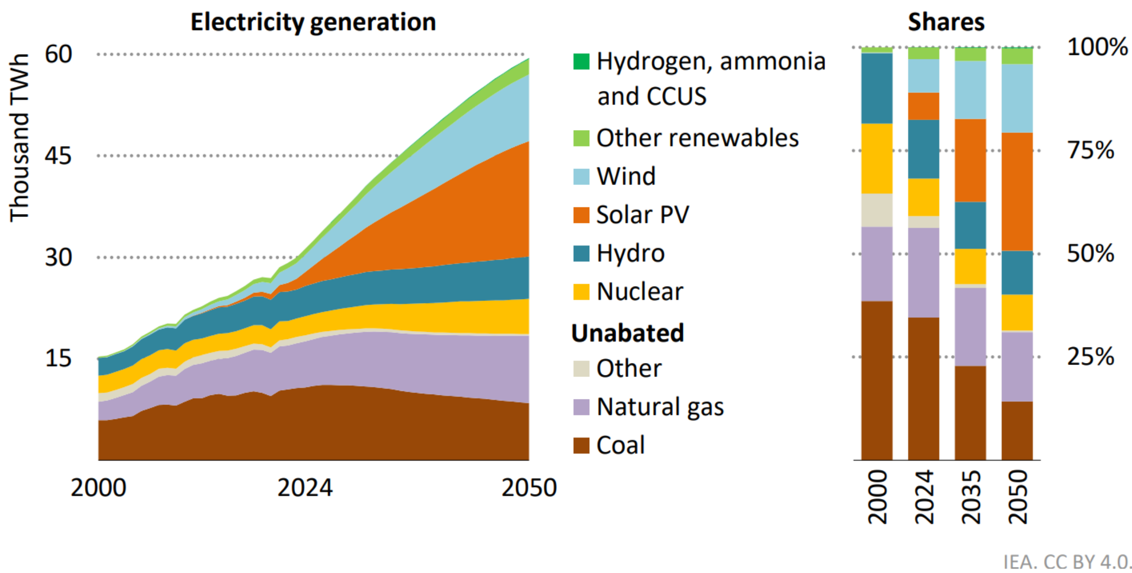

Figure 1.1: Global electricity generation by source in the Current Policies Scenario (CPS) to 2050 (Figure 3.18 from [3]).

2025, utility-scale solar has a levelized cost of energy (LCOE) of \$58 per MWh, while onshore wind has an LCOE of \$61 per MWh [6]. Utility-scale standalone battery energy storage systems with 4-hour capacities have reached a levelized cost of storage of approximately \$200 per MWh [6]. Meanwhile, Gas-fired generation have reached 10-year high LCOE at \$78 per MWh with construction costs expected to continue rising for combined cycle gas turbine generating plants due to turbine shortages and long delivery times [6].

With the growth of IBRs outpacing that of synchronous generators, systems are becoming dominated by IBRs. This has brought with it a range of challenges relating to the control and stability of power systems.

1.1.2 Challenges with Grid-interfacing Inverters

The conventional synchronous generators, that AC systems were founded upon, inherently have inertia due to the large spinning mass that makes up the rotor of these devices, and provide voltage support through field excitation of their rotor windings. This inertia has long provided a stable foundation for AC power systems and allowed these systems to

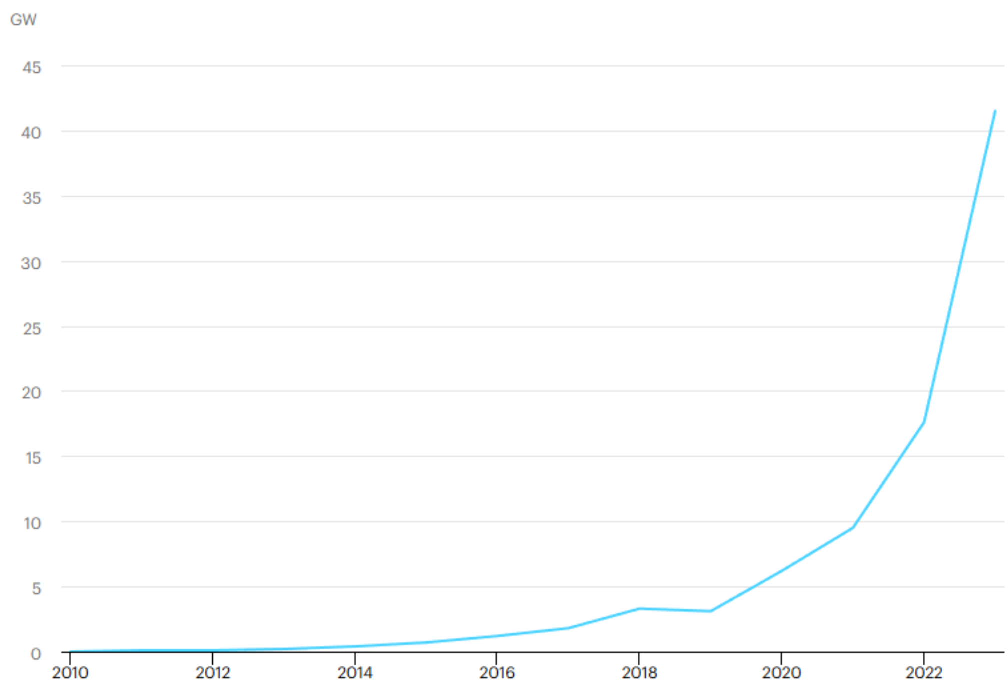

Figure 1.2: Global battery storage capacity additions, 2010-2023 [5].

maintain stability following disturbances. IBRs do not come with this inherent physical inertia or voltage regulation, and must be intentionally controlled in ways that support AC power systems.

Inverter Control

Most commercial IBRs currently utilize grid-following (GFL) control, which employs a phase-locked loop (PLL) to synchronize to the grid voltage, allowing the inverter to adjust current injections to meet power targets. While this approach is effective in strong grids, it relies on the assumption that the point of interconnection voltage remains stable and provides a reliable phase reference–an assumption that is increasingly challenged as IBR penetration rises. This reliance on a grid reference creates several stability challenges. Conventional power systems have historically relied on synchronous generator inertia to limit the rate of change of frequency (RoCoF) following a disturbance [7]. As GFL IBRs displace synchronous generation, system inertia decreases, leading to higher RoCoF and increased frequency instability. Although GFL inverters can be augmented with grid-supporting features, such as fast frequency response [8] or reactive power-voltage droop controls, these solutions are often constrained by the inverter's primary operating mode. GFL inverters are typically controlled to inject maximum active power at near-unity power factor, which leaves limited headroom for the reactive power adjustments necessary to provide the voltage support typically offered by synchronous generators.

To overcome these fundamental limitations of grid-following control, inverters can instead use grid-forming (GFM) control. With GFM control an inverter synthesizes its own angle reference and can effectively provide *synthetic inertia*. Many GFM control methods have been proposed in the literature, but the three most common are virtual synchronous machine (VSM) control [9], virtual oscillator control (VOC) [10], and power angle droop control [11, 12]. VSM enables the inverter to imitate the physical behavior of synchronous machines. VOC generates its internal reference from a digitally simulated oscillator. Droop control adjusts its voltage magnitude (V) and frequency (ω) from nominal values based on deviations in the measured active (P) and reactive (Q) power output values from desired reference values (P^*, Q^*) with the control law

$$V = V_{\text{nom}} - m_q(Q - Q^*), \tag{1.1}$$

$$\dot{\theta} = \omega = \omega_{\text{nom}} - m_p(P - P^*). \tag{1.2}$$

All of these GFM control methods can be considered as instances of a generic model, and exhibit similar droop-like relations in steady-state [13]. Each method fundamentally commands the inverter voltage to operate as an oscillator that is coupled to the grid through the device's measured power outputs, thereby enforcing synchronization and regulation through active and reactive power balance.

Another interpretation of GFM control is that it causes the inverter to act as a stiff voltage source behind an impedance [14]. By controlling the inverter to act as a stiff voltage source, grid-forming control can imitate the physical properties of synchronous generators, which resist changes in the voltage frequency due to the inertia of their rotors, inherently supporting the system during transients. If there is a perturbation on the grid-side voltage, then the inverter will temporarily hold its voltage output based on its internal angle reference and inject more power into the grid resisting the change. This can be seen in (1.2), where the voltage angle of the droop control takes time to respond to sudden changes in its active power output during a transient where the grid frequency deviates from the nominal value. These GFM control methods have been shown to be able to help stabilize IBR-based systems [15, 16] and are widely consider an integral part of achieving high levels of IBR penetration in large-scale power systems [17].

Inverter Current Limits

While GFM control methods appear to be promising in supporting the development of stable power systems with high IBR penetration, there are still challenges with deploying these methods in practice. A significant challenge with GFM control is how to handle fault-ride through when maintaining the desire-able stiff voltage reference from the GFM control would require current output levels beyond the current magnitude limits of the inverter device [18, 19]. When a fault occurs on the system the current output from GFM devices increases rapidly as the grid-voltage magnitude or frequency deviates from its nominal values[4]. To maintain the voltage source behavior of the GFM requires a significant increase in the current output of the device which may not be feasible due to the device's current limits. These current limits are important to prevent damaging current levels from conducting through the switches that compose an inverter [20]. Therefore, fault-ride through strategies are required for GFM controlled IBRs. At the moment, it is unclear what the best approach for controlling a GFM IBR when its current limits are reached as it becomes a trade-off between maintaining the voltage source behavior of the GFM control mode and maintaining the physical current limits of the device.

System Stability

As IBRs are increasingly influencing system dynamics, stability analysis approaches relevant to systems dominated by synchronous generators are becoming less applicable. Voltage

[4]This is typical of low voltage faults and phase jumps

stability in conventional power grids is typically related to slower timescale dynamics, while frequency stability relates to faster timescales, allowing for the decoupling of the stability analysis of these parameters. With increasing IBR penetration, traditional boundaries in timescales used to decouple stability analysis of voltage and frequency are becoming less distinct and both mechanisms are increasingly influenced by faster dynamics [21]. This has made it more challenging to evaluate the stability of these systems.

The ability of power systems to maintain stable voltage and frequency under disturbances can broadly be referred to as system strength [22]. The large-scale integration of IBRs over the last decade has led to a decrease in the strength of the grid [23], which has resulted in challenges with maintaining stability. Weak grids and GFL IBRs are associated with stability challenges such as voltage instability [24] and subsynchronous oscillations [25]. These instabilities occur due to control interactions between IBR devices and the rest of the system or other IBR devices. GFM IBRs can help support system stability, but are prone to becoming unstable in strong grids and can also induce oscillations due to interactions between their controllers and the rest of the system.

Ensuring systems with high IBR contributions remain stable requires careful consideration of how the inverters in the system are controlled, and rigorous assessments of the system's performance. Electromagnetic transient (EMT) simulations can be used to assess grid stability, but these simulations are computationally and time intensive for large systems and can be performed only for a limited number of possible transients and system states. Metrics that reflect the strength of the grid can be used to screen for weak system areas and identify potential instability issues to focus these studies. However, traditional indices like the short-circuit ratio (SCR) that rely on short-circuit analysis cannot accurately quantify system strength for systems with high IBR penetration. Consequently, they can fail to identify areas prone to instabilities caused by control interactions between IBR devices. Therefore, there is a demand for new system strength metrics that can both be used to identify weak areas of the grid to focus EMT studies and serve as online indicators of grid strength.

Inverter-based Loads

The impact of power electronics on AC power system stability is not limited to generation, but is increasingly influenced by the growth of inverter-based loads. Historically, power system loads were passive, characterized by constant impedance properties, and were largely treated as stochastic, aggregate variables. However, modern computational loads act as fast changing constant power loads that break this load passivity assumption.

This challenge is exacerbated by the recent growth in data center demand, particularly for the training of large language models (LLMs). LLM training requires thousands of GPUs to operate in temporal synchronization, creating a highly volatile bimodal load profile where aggregate demand steps from idle to maximum capacity in milliseconds. These rapid transitions induce recurring voltage sags across distribution feeders, complicating local voltage regulation. Furthermore, modern data centers are highly sensitive to transient voltage dips. If voltage sags at the point of interconnection trigger sensitive undervoltage protection, facilities may instantaneously disconnect from the grid to transition to backup power. This sudden loss of load can induce significant system-level power imbalances and frequency excursions.

As computational loads scale, the assumption that they function as passive grid participants is no longer valid. To maintain system stability, it is imperative that these large loads be subjected to regulation. Focus must shift from supply-side generation regulation toward demand-side mechanisms that align load behavior with system stability requirements. Consequently, the adoption of active, integrated control is required for these loads to mitigate their impact on the system.

1.2 Outline

In Chapter 2, we introduce the challenge of controlling IBRs considering current limits and our approaches to this problem. We first identify stability conditions for the control of an IBR with a current limiter, next we propose a safety filter approach for IBR control to ensure current limits are maintained, then we show that the geometry of the feasible output of an IBR constrained by current limits is convex, and lastly we propose a safe trajectory optimizing controller and apply it to a grid-interfacing inverter.

In Chapter 3, we focus on the small-signal stability of power systems containing IBRs. We review the topics of grid strength and system strength as they relate to small-signal stability, and evaluate the applicability of existing metrics to systems with high levels of IBRs. Next, we propose a new system strength metric that is based on the sensitivity of a system's eigenvalues to perturbations in power flow parameters.

In Chapter 4, we turn our attention towards the application of IBRs in voltage regulation problems. In particular, we consider how IBRs can be used to mitigate voltage fluctuations due to fast changing data center loads. First, we consider the use of switching reference reactive power control scheme that regulates the grid voltage. Then we consider how uninterruptible power supply energy storage headroom can be used to regulate the active power load of data centers.

The dissertation concludes by identify future directions for research that builds upon this work.

Chapter 2

OPTIMAL & SAFE CONTROL OF GRID-INTERFACING INVERTERS WITH CURRENT MAGNITUDE LIMITS

In this chapter, we investigate the control of grid-interfacing inverters with current magnitude limits. This chapter is split up into four sections on this topic: First, we introduce the problem and a standard inverter system model that we base our investigation on. Next, Section 2.2 covers conditions for the stability of IBRs with current limiters. Then, Section 2.3 proposes a safety filter approach for ensuring current limits are not violated. Section 2.4 discusses the geometry of the feasible output space of IBRs with current magnitude limits. Lastly, Section 2.5 proposed a control strategy based on safe gradient flows for trajectory optimization and applies it to the IBR current limiting control problem.

2.1 Introduction, Background & Modeling

A number of control approaches have been proposed with the aim of supporting the stability of the grid, such as droop control [11], virtual synchronous machines [9], and virtual oscillator control [10]. However, these approaches assume the inverter model to be ideal (i.e. the device can achieve arbitrary voltage and current outputs) and do not model physical limitations on the devices that are important in practice. In reality, controlling an inverter connected to the grid requires a trade-off between system stability and self-protection [26]. One of these limitations needed for self-protection is the amount of current that can pass through the switches of inverter devices before they are damaged. To compensate for neglecting these limitations when designing controllers auxiliary modules such as current limiters and virtual impedance methods are used, but these can introduce undesirable side effects, such as altered voltage dynamics or poor stability [27]. While many current limiting methods for grid-interfacing inverters have been proposed, there are still open issues related to how to limit the current when controlling an inverter as a voltage source (as is the case with grid-forming control) [27].

2.1.1 Literature on the Control of IBRs with Current Limits

A current limiter saturates the reference of a closed-loop current controller to avoid exceeding the allowed current magnitude, as seen in Fig. 2.1. This approach, although simple, changes the behavior of an inverter, and makes its stability and voltage support capability depend on the exact operating conditions [28, 29]. To avoid these qualitative behavioral changes, voltage limiters have been proposed to directly limit the voltage difference across the inverter filter, such that the current would not exceed its magnitude limits [30, 31]. Virtual impedance methods can also be used to avoid changes in voltage leading to large changes in current that exceed the current magnitude constraint [32, 33]. However, these methods require careful design based on detailed knowledge of parameters or system measurements. In addition, they invariably involve a tradeoff between current limiting, and the performance and stability of the inverter.

This type of saturation is nontrivial to analyze. In particular, currents are vectors

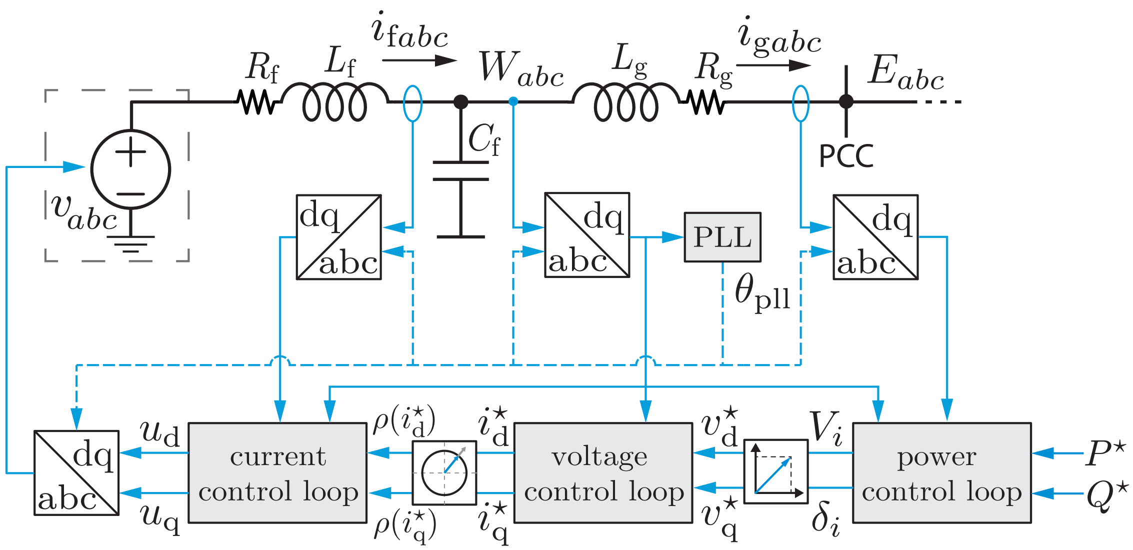

Figure 2.1: Full-order inverter system model with inner control loops and current limiter. This diagram shows the angle reference being generated using a phase-lock loop tracking the LCL capacitor voltage angle, but this angle reference for the dq reference frame can also be generated from other sources (e.g. the power control loop) for use with this same general control scheme.

in $\mathbb{R}^2$ (in the dq frame) and the saturation is on their magnitudes. Hence the tools for sector-constraints developed for Lur'e systems [34] do not readily apply. Currently, current saturation is either ignored or analyzed in linearized [35] and reduced model systems [36].

To optimize the tradeoff between satisfying current limits and performance, a number of techniques have been proposed. A common approach is to design a "good" controller by looking at a system without consideration for the current constraint, then modify its control action as little as possible when current limits are encountered. Some research has been presented for virtual impedance methods indicating that the magnitude of the change of the inverter terminal voltage is minimized when the phase angle of the virtual impedance equals that of the passive impedance composed by filter impedance and grid impedance [37]. This provides a method for selecting the virtual impedance parameters, but requires knowledge of the passive impedance which can change under different operating scenarios and be challenging to obtain exactly. Another relevant approach is the safety filter framework presented in [38], which is demonstrated with a model of an inverter connected to an infinite bus. This framework can design safety filters (functions that modify control actions) by solving a sum of squares optimization problem and implement the resulting safety filter as a quadratically constrained quadratic program. However, finding these guarantees is computationally difficult and would have to be computed separately for differing inverter systems. Next, projected-droop control, a control method that explicitly considers the current limits of an inverter, is proposed in [39]. This approach is designed specifically for the dynamics of a droop controller and is not generally applicable to other control methods. [40] presents a similar approach that uses primal-dual gradient dynamics to develop a feedback controller for a generic grid-forming control problem and considers a range of constraints. However, this approach does not guarantee constraint satisfaction for all times as the model does not capture the circuit dynamics of the system.

2.1.2 Current-limited Inverter Dynamics & Modeling

We base our approaches on the dynamics and general control structure of the following model of an inverter system connected to the grid via an LCL filter [41]. This inverter system model employs inner voltage and current control loops [36], a phase-locked loop (PLL) [42] and a current limiter. The dynamics for this full-order model can be found in Appendix A.1.

Simplified Inverter Model

To simplify the full-order model shown in Fig. 2.1, we assume the inner voltage and current control loops act fast enough, such that we can ignore the dynamics of the capacitor. This results in a dynamical system model of a voltage-source inverter connected to the grid through an RL filter, which is

$$\begin{bmatrix} \dot{I}_{\mathrm{d}} \\ \dot{I}_{\mathrm{q}} \end{bmatrix} = \begin{bmatrix} -\frac{R}{L} & \omega \\ -\omega & -\frac{R}{L} \end{bmatrix} \begin{bmatrix} I_{\mathrm{d}} \\ I_{\mathrm{q}} \end{bmatrix} + \frac{1}{L} \left(\begin{bmatrix} V\cos(\delta) \\ V\sin(\delta) \end{bmatrix} - \mathbf{E}_{\mathrm{dq}} \right), \tag{2.1}$$

where R and L are the resistance and impedance of the RL filter, respectively, $\mathbf{E}_{\mathrm{dq}}$ is the grid-side voltage with reference to some angle θ, ω is the rotational frequency of the reference frame, V is the inverter voltage magnitude, and δ is the angle difference between the inverter voltage and θ. A diagram representing this simplified inverter model is shown in Fig. 2.2.

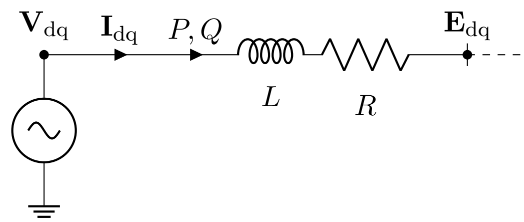

Figure 2.2: A simplified model of an inverter as a controllable voltage source in the dq-frame connected via an RL filter to an infinite bus.

Making a small-angle assumption, such that $\cos(\delta) \approx 1$ and $\sin(\delta) \approx \delta$, assuming the reference frame rotates near to a nominal value ω_{nom}, and assuming $\mathbf{E}_{\mathrm{dq}} = (E, 0)$, where E is the grid-side voltage magnitude, this reduces to the affine system [43, 44]

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} - \mathbf{E}_{\mathrm{dq}}, \tag{2.2}$$

with $\mathbf{A} = \begin{bmatrix} -\frac{R}{L} & \omega_{\mathrm{nom}} \\ -\omega_{\mathrm{nom}} & -\frac{R}{L} \end{bmatrix}$ and $\mathbf{B} = \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix}$, where our states are $\mathbf{x} = (I_{\mathrm{d}}, I_{\mathrm{q}})$ and our inputs are $\mathbf{u} = (V, \delta)$. Note that we can alternatively take the inputs to be $\mathbf{u} = (V_{\mathrm{d}}, V_{\mathrm{q}})$, with $\mathbf{B} = \frac{1}{L}$ and $V_{\mathrm{q}} = E \cdot \delta$. Taking the inputs to be centered around the grid voltage and centering the states to be around an equilibrium value results in a linear system model.

Inverter Current Limits

To protect the inverter from damaging current levels, the magnitude of the current must be limited [45, 46][1] With the full-order model with LCL filter, introduced in Section 2.1.2, we consider the current limits as a constraint on the magnitude of the inverter-side current $\mathbf{I}_{i\mathrm{dq}}$. We denote this limit as $I_{\max}$, and we define the safe (or feasible) current operating

[1]Inverter current limits arise from two aspects, thermal limits and inductor saturation [47]. Thermal limits are "soft constraints," in that they can be violated for short periods of time. Inductor saturation is a hard constraint and has values typically a few times larger than thermal constraints [48].

region as

$$\mathcal{I} := \{\mathbf{I}_{i\mathrm{dq}} \mid \|\mathbf{I}_{i\mathrm{dq}}\|_2^2 \leq I_{\max}^2\}. \tag{2.3}$$

Fig. 2.3 visualizes this circular magnitude limit. Note that we can find the current reference

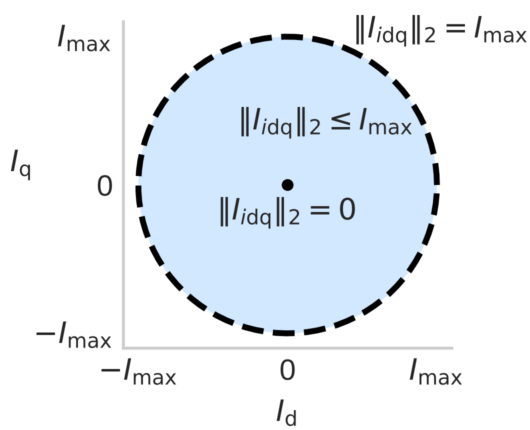

Figure 2.3: Current magnitude constraint imposed on inverter output current.

from desired power reference values as

$$\begin{bmatrix} I_{\mathrm{d}}^* \\ I_{\mathrm{q}}^* \end{bmatrix} = \begin{bmatrix} V_{\mathrm{d}}^* & V_{\mathrm{q}}^* \\ V_{\mathrm{q}}^* & -V_{\mathrm{d}}^* \end{bmatrix} \begin{bmatrix} P^* \\ Q^* \end{bmatrix}, \tag{2.4}$$

if we solve for $(V_{\mathrm{d}}^*, V_{\mathrm{q}}^*) = (V^* \cos\delta^*, V^* \sin\delta^*)$ from steady-state power flow equations with the stiff grid voltage and given power references [49].

2.2 Optimal Control of IBRs with Current Limiters[2]

One common approach to current limiting is to simply saturate a controller that is designed for an unconstrained system. This can be done with a direct current limiter, which is an element that can address over-currents that may appear during faults and voltage fluctuations, and damage sensitive semiconductor devices in inverters. These limiters achieve this by directly saturating the current reference provided to the fast inner current control loop, such that the current cannot exceed the device's limits [50]. This approach can lead to stability issues or limit the control actions to be too conservative. This type of saturation is nontrivial to analyze. In particular, currents are vectors in $\mathbb{R}^2$ (in the dq frame) and the saturation is on their magnitudes. Hence, the tools for sector-constraints developed for Lur'e systems [34] do not readily apply. Currently, current saturation is either ignored or analyzed in linearized [35] and reduced model systems [36].

In this section, we directly focus on a nonlinear system that explicitly accounts for the saturation of the current magnitude. We use a Lyapunov stability approach to determine a stability condition for the system, guaranteeing that a class of controllers would be stabilizing if they satisfy a simple semidefinite programming condition. With this condition we fit a linear-feedback controller by sampling the output of (offline) model predictive control problems. This learned controller has improved performances compared to existing designs.

2.2.1 System Model

We use the simplified inverter system model described in Section 2.1.2. With this model, our goal is to design a feedback controller such that the system tracks a constant power setpoint, (P^*, Q^*), at the grid side. The dynamics of the system comes from the RL filter element in the circuit, and we think of the current $\mathbf{I}_{\mathrm{dq}} = (I_{\mathrm{d}}, I_{\mathrm{q}})$ as the state of the system. We model the inverter itself as a voltage source converter, and treat its output, $\mathbf{V}_{\mathrm{dq}} = (V_{\mathrm{d}}, V_{\mathrm{q}})$ as the actuator output. We take the magnitude of the inverter voltage, V, and the angle difference between the inverter and grid voltages, δ, to be the controller outputs and choose the d-axis

[2] Adapted from T. Joswig-Jones and B. Zhang, "Optimal control of grid-interfacing inverters with current magnitude limits," in 2024 IEEE 63rd Conference on Decision and Control (CDC). IEEE, 2024, pp. 4623–4628 [29]

to align with the inverter voltage such that $V_d = V$ and $V_q = 0$[3]. For the simplified model, we assume that the angle of the grid-voltage is known, however in the full-order model this is a measured value using a phase-lock loop. We further assume the grid side voltage E has a constant magnitude such that $\mathbf{E}_{\mathrm{DQ}} = (\sqrt{2}V_{\mathrm{nom}}, 0)$, where V_{nom} is the nominal phase-neutral voltage of the grid[4]. In the inverter reference frame, this is $\mathbf{E}_{\mathrm{dq}} = (E_{\mathrm{d}}, E_{\mathrm{q}}) = R(\delta) \cdot \mathbf{E}_{\mathrm{DQ}}$, where $R(\cdot)$ is a counter-clockwise rotation matrix.

With the above notations and assumptions, the active and reactive powers at the grid side are given by

$$P = \frac{3}{\sqrt{2}} V_{\mathrm{nom}} \cdot I_{\mathrm{d}}, \;\; Q = -\frac{3}{\sqrt{2}} V_{\mathrm{nom}} \cdot I_{\mathrm{q}}.$$

Therefore, tracking a given P^*, Q^* is equivalent to tracking some I_{d}^* and I_{q}^*.

The reason we make these assumptions is to isolate the challenge introduced by the nonlinear saturation block, which we describe next.

Current Saturation

To protect internal devices, the inverter's output current cannot exceed a preset limit. Different from typical saturation limits in the literature, which limits each component of control input or the state, the limit in our case is on the magnitude of the current vector. More precisely, we define

$$\mathrm{sat}(\mathbf{z}) = \begin{cases} \mathbf{z} & \text{if } \|\mathbf{z}\|_2 \le 1 \\ \frac{1}{\|\mathbf{z}\|_2} \cdot \mathbf{z} & \text{if } \|\mathbf{z}\|_2 > 1 \end{cases} = \frac{1}{\max\left(\|\mathbf{z}\|_2, 1\right)} \cdot \mathbf{z}. \tag{2.5}$$

This modeling approach assumes the direct current limiter instantaneously tracks the given saturated current reference values projecting the current state to the circular magnitude limit. This saturation on the magnitude[5] of the currents is not a sector inequality, therefore some of the results from Lur'e problems do not directly apply. The main theoretical result

[3]With a small-angle assumption this is equivalent to taking V_{dq} to be the controller outputs.

[4]The assumption that $E_{\mathrm{q}} = 0$ can be made without loss of generality, but the assumption E_{d} is constant is more material. An important future question is to allow time variation in E_{d}.

[5]Without loss of generality, we can take the value of the saturated magnitude to be 1 when developing these theorems.

of this paper is that a simple geometric argument allows us to characterize the stability of a linear system with magnitude saturation.

Inverter Model

We obtain the discrete time dynamical system model by discretizing the differential equations (2.2) governing the current through the inductor using the forward Euler method, with a saturation limit on the current:

$$\mathbf{x}_{t+1} = \mathrm{sat}\{\hat{\mathbf{A}}\mathbf{x}_t + \hat{\mathbf{B}}\mathbf{u}_t\}, \tag{2.6}$$

where

$$\hat{\mathbf{A}} = \mathbf{I} - \Delta t \begin{bmatrix} -\frac{R}{L} & \omega_{\mathrm{nom}} \\ -\omega_{\mathrm{nom}} & -\frac{R}{L} \end{bmatrix}, \quad \hat{\mathbf{B}} = \Delta t \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix},$$

where $\mathbf{I}$ is the identity matrix and Δt is a sufficiently small discretization timestep. Fig. 2.4 provides a diagram of this system model.

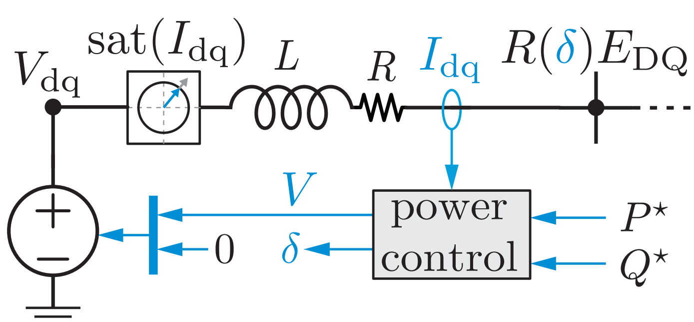

Figure 2.4: Simplified inverter system model under study.

Our goal is to design an optimal linear feedback controller such that the system in (2.6) is asymptotically stable around a setpoint $\mathbf{x}^*$.

2.2.2 Lyapunov Stability

We assume that the tracking problem is feasible, that is, the reference $\mathbf{x}^*$ satisfies $||\mathbf{x}^*||_2 < 1$. Let $\Delta\mathbf{x} = \mathbf{x} - \mathbf{x}^*$ denote the shifted state with the reference at the origin. Note that the saturation is on $||\Delta\mathbf{x}+\mathbf{x}^*||_2$. In this paper, we study the class of linear feedback controllers with the gain matrix K, where $\mathbf{u}=\mathbf{u}^*-K(\mathbf{x}-\mathbf{x}^*)$. In this section we find a constraint on K, such that the saturated-state system is asymptotically stable and has a unique equilibrium.

The system with shifted state $\Delta\mathbf{x}$ is

$$\Delta\mathbf{x}_{t+1} = \mathbf{A}(\hat{\mathbf{x}}_t - \mathbf{x}^*) - BK(\hat{\mathbf{x}}_t - \mathbf{x}^*), \tag{2.7a}$$

$$\hat{\mathbf{x}}_t = \mathrm{sat}(\Delta\mathbf{x}_t + \mathbf{x}^*), \tag{2.7b}$$

where we introduce $\hat{\mathbf{x}}_t$ as a variable that represents the saturated state at time t. We have the following result.

Theorem 1. *Consider the system in* (2.7)*. Suppose the reference and the initial starting point satisfy* $||\mathbf{x}^*||_2 < 1$ *and* $||\mathbf{x}_0||_2 < 1$*, respectively. The system is asymptotically stable around* $\mathbf{x}^*$ *if* K *satisfies*

$$(\mathbf{A} - BK)^\top(\mathbf{A} - BK) - I \prec 0. \tag{2.8}$$

We will first develop some intuitions as to the issues that including the saturation bound on the states introduces. The condition in (2.8) is the standard Lyapunov stability condition for linear systems where the P matrix is chosen to be identity [51]. This means that the unsaturated system needs to be stable with a Lyapunov function that has circular level sets. The intuition behind this result is that if the Lyapunov function's level sets and the saturation function have the same shape, we can use triangle inequality to show that the trajectory of the system with saturation never gets "stuck". It is possible to find systems that are stable but do not have Lyapunov functions with circular level sets, where the trajectory of the saturated system can remain stationary at some point along the magnitude bound and not converge to the reference (see Section 2.2.4). Fig. 2.5a shows how without a circular level set the state dynamics can align with the normal of the circular saturation function, resulting in the state getting "stuck" in this position. Fig. 2.5b visualizes how with a circular level set the normal of the saturation function circle and the dynamics step in the states cannot align. The proof of Theorem 1, can be found in Appendix A.2

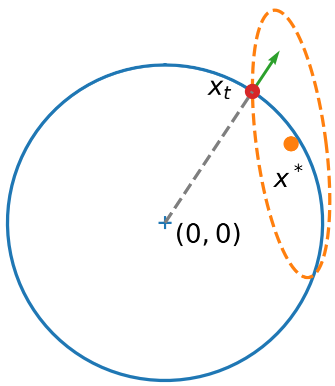

(a) A case where $\mathbf{x}_t$ is "stuck" and will remain stationary; never reaching $\mathbf{x}^*$.

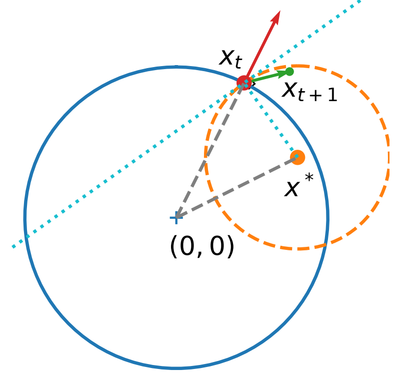

(b) A case where the Lyapunov level set is circular and $\mathbf{x}_{t+1} - \mathbf{x}_t$ can not align with the saturation function normal $\mathbf{x}_t$.

Figure 2.5: Illustrative figures showing the impact the shape of the Lyapunov level sets has on asymptotic stability. The saturation function boundary and the Lyapunov level sets are represented as ellipses with solid and dashed lines, respectively.

Stability with unknown line parameters

The resistance and inductance parameters—R and L defining the matrices $\mathbf{A}$ and B—are the sum of the inverter side and grid side parameters. In practice, the inverter side R and L are known since they are part of the design parameters, but the grid side parameter depends on the exact connection setup and is often unknown. Therefore, it is useful to design K to be robust to this uncertainty in the grid parameters. The next theorem provides a way to do this assuming that the grid-side line is inductive ($L_g > 0$).

Theorem 2. *Suppose the gain matrix K satisfies* (2.8) *for some given* $\mathbf{A}$ *and* B *matrices. Then K satisfies* (2.8) *if the value of R is increased, as long as $L > 0$ and the discretization steps Δt are sufficiently small.*

Therefore, if the parameters of the grid-side line inductance and resistance are unknown, then it suffices to design the controller assuming the grid side resistance is 0 (the value of the inductance does not enter into stability). This leads to a suboptimal, but robust controller. The proof of Theorem 2 is in Appendix A.3.

2.2.3 Controller Development

In the last section, Theorems 1 and 2 provided stability conditions on the gain matrix $\mathbf{K}$, but not optimality conditions for how to select a stabilizing $\mathbf{K}$. A number of objective functions are used in practice to optimize the controller gain, including settling times, rate of change of current, overshoot, control efforts, and a combination of these. In general, there are no simple rules on how to find an optimal $\mathbf{K}$. In this paper, we take a data-driven approach, where we first simulate a set of MPC controlled trajectories and then fit a $\mathbf{K}$ that best approximates the MPC solutions.

Model Predictive Controller

We formulate the optimal control problem for the discrete time system with saturation as the constrained optimization problem

$$\begin{aligned}
\mathbf{u}_{\text{mpc}}^{\mathbf{x}_{\text{init}},\mathbf{x}^*} := \underset{\mathbf{u}_1,\ldots,\mathbf{u}_T}{\arg\min} \sum_{i=1}^{T} \quad & f(\mathbf{x}_i, \mathbf{u}_i, \mathbf{x}^*) && \text{(2.9a)}\\
\text{s.t.} \quad & \Delta\mathbf{x}_{t+1} = \mathbf{A}(\mathbf{x}_t - \mathbf{x}^*) + \mathbf{B}\mathbf{u}_t && \text{(2.9b)}\\
& \mathbf{x}_t = \text{sat}(\Delta\mathbf{x}_t + \mathbf{x}^*) && \text{(2.9c)}\\
& \mathbf{x}_0 = \mathbf{x}_{\text{init}} && \text{(2.9d)}
\end{aligned}$$

where $\mathbf{x}_{\text{init}}$ and $\mathbf{x}^*$ are the initial value of the states and the reference value for the controller, respectively, and $f(\mathbf{x}_i, \mathbf{u}_i, \mathbf{x}^*)$ is a user selected objective function. We use this problem formulation to run a nonlinear model predictive controller.

Fitting a Static Controller

Using the stability constraint from Theorem 1 we formulate the linear regression problem:

$$\begin{aligned}
\underset{\mathbf{K}}{\arg\min} \quad & \sum_{i\in U} \left(\sum_{t=0}^{T_i-1} \| -\mathbf{K}(\Delta\mathbf{x}_t^i) - \mathbf{u}_t^i \|_2 \right) && \text{(2.10a)}\\
\text{s.t.} \quad & (\mathbf{A} - \mathbf{B}\mathbf{K})^\top (\mathbf{A} - \mathbf{B}\mathbf{K}) - \mathbf{I} \prec 0. && \text{(2.10b)}
\end{aligned}$$

This problem is convex, since we can express (2.10b) as a linear matrix inequality using Schur complement as shown in [52]:

$$\underset{\mathbf{K}}{\arg\min} \quad \sum_{i\in U}\left(\sum_{t=0}^{T_i-1}\|-\mathbf{K}(\Delta\mathbf{x}_t^i)-\mathbf{u}_t^i\|_2^2\right) \tag{2.11a}$$

$$\text{s.t.} \quad \begin{bmatrix} \mathbf{I} & (\mathbf{A}-\mathbf{B}\mathbf{K})^\top \\ (\mathbf{A}-\mathbf{B}\mathbf{K}) & \mathbf{I} \end{bmatrix} \succ 0. \tag{2.11b}$$

Here U is a dataset containing sets of state values, $\Delta\mathbf{x}^i \in \mathbb{R}^{n\times T_i}$ and associated MPC optimal input values, $\mathbf{u}^i \in \mathbb{R}^{m\times T_i-1}$, where n is the number of states, m is the number of inputs, and T_i is the length of simulation i.

2.2.4 Simulation Results

To test our approach, we fit a controller to a series of MPC responses and we simulate the response of the system with the MPC-fit controller and a base-line controller for comparison. We select the base-line controller to be the linear–quadratic regulator (LQR) with

$$\mathbf{K}_{\text{base}} = \begin{bmatrix} 1.206 & 0.0957 \\ 0.096 & 0.0671 \end{bmatrix}, \tag{2.12}$$

found with $\mathbf{Q} = \begin{bmatrix} 1 & 0 \\ 0 & 0.1 \end{bmatrix}$ and $\mathbf{R} = 5{\cdot}\mathbf{B}$. In the following we use the function $\text{linspace}(a, b, n) \coloneqq a + \frac{n-1}{i-1}(b-a)$ to define a list of evenly spaced numbers over a specified interval, and the notation $(\cdot \times \cdot)$ to represent a Cartesian product of two sets.

We generate a set of MPC responses, $U \coloneqq \{\mathbf{x}_{\text{mpc}}^{\mathbf{x}_{\text{init}},\mathbf{x}^*}, \mathbf{u}_{\text{mpc}}^{\mathbf{x}_{\text{init}},\mathbf{x}^*} \mid (\mathbf{x}_{\text{init}}, \mathbf{x}^*) \in X_{\text{init}} \times X_{\text{ref}}\}$. U contains the state trajectories and optimal MPC inputs for each $(\mathbf{x}_{\text{init}}, \mathbf{x}^*)$ pair from a mesh of values in the polar grids, $X \coloneqq \{r\cos\theta + \frac{\pi}{4} \mid (r,\theta) \in R \times \Theta\}$, where $R \coloneqq \text{linspace}(0, I_{\max}, 3)$ and $\Theta \coloneqq \text{linspace}(0, 2\pi - \frac{2\pi}{4}, 4)$. $I_{\max}$ is the current magnitude limit of the inverter.

We use GurobiPy [53] to solve a rolling time horizon version of (2.9a) with an LQR objective $\sum_{i=1}^{T} \ (\mathbf{x}_i - \mathbf{x}^*)^\top \mathbf{Q}(\mathbf{x}_i - \mathbf{x}^*) + \mathbf{u}_i^\top \mathbf{R}\mathbf{u}_i$ where $\mathbf{Q}$ and $\mathbf{R}$ are the same matrices used to design the LQR base-line controller. We choose a discretization step size of $\Delta t = 10$ μs

Table 2.1

INVERTER & FILTER SIMPLIFIED SYSTEM PARAMETERS

Parameter	V_{nom}	S_{nom}	I_{nom}	I_{max}	E	ω_{nom}	R	L
Value	120 V	1.5 kV A	4.167 A	4.167 A	120 V	$2\pi 60\ \mathrm{rad\,s^{-1}}$	1.3 Ω	3.5 mH

and use parameter values that are modified from those used in [36] for the inverter and *LCL* filter system. The parameter values used in these simulations are listed in Table 2.1. The optimization horizon used is 5 periods and each simulation runs until the states reach a steady-state value; $\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2 < 10^{-5}$ for 10 consecutive periods.

We then fit a linear feedback controller $\mathbf{K}_{\text{fit}}$ by solving (2.11) with our dataset U using CVXPY [54]. We find the optimal fit matrix to be

$$\mathbf{K}_{\text{fit}} = \begin{bmatrix} 0.608 & 0.027 \\ 0.012 & 0.026 \end{bmatrix}. \tag{2.13}$$

We then test this MPC-fit controller and base-line controller on the simplified system by running simulations with all $(\mathbf{x}_{\text{init}}, \mathbf{x}^*)$ in the set X. The results of these tests are detailed in the next section. The code used to generate these results is available at https://github.com/TragerJoswig-Jones/Current-Magnitude-Limited-Inverter-Control.

Simplified Model Test Results

The average cost of control for the MPC-fit and base-line controllers were 59.2 and 115.1, respectively. While the fit controller outperformed the base-line controller in terms of the average cost, the base-line controller showed lower costs for a majority of the test cases. However, for some test cases (22 to be exact) the base-line controller got stuck along the boundary and did not converge to the reference value.

Fig. 2.6a shows the trajectory of the absolute value of the state errors for the MPC-fit controlled and base-line controlled systems for a selected test case. In this case the base-line controller gets stuck at the circular magnitude bound and does not converge to the

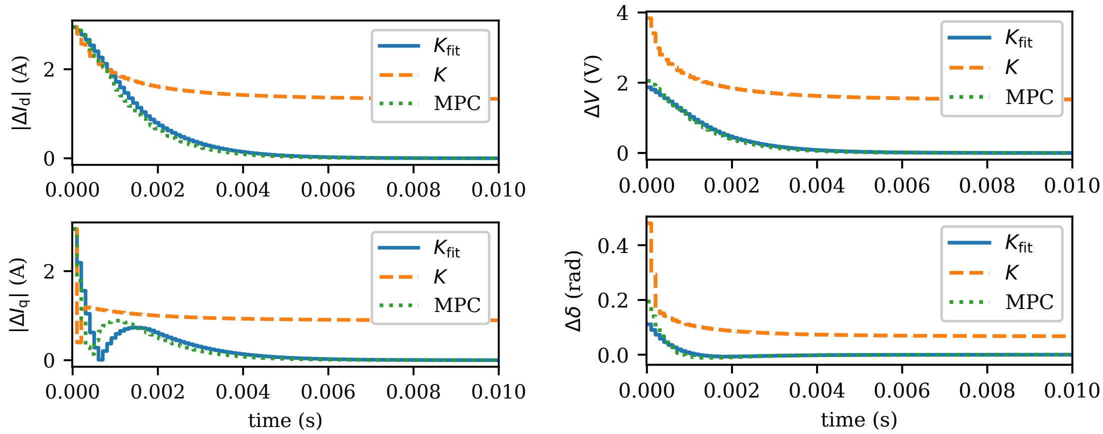

(a) $|\Delta \mathbf{I}_{\mathrm{dq}}|$ trajectories.

(b) $\Delta V, \Delta\delta$ input trajectory.

Figure 2.6: State and input trajectories for $\mathbf{K}, \mathbf{K}_{\text{fit}}$, and MPC controllers with $\mathbf{x}_0 = (0.00 \text{ A}, 0.00 \text{ A})$, $\mathbf{x}^* = (2.95 \text{ A}, 2.95 \text{ A})$. The $\mathbf{K}_{\text{fit}}$, and MPC controllers converge to the reference value, while the base K controller gets stuck at a point along the current limit.

reference value. We note that for this selection of $\mathbf{Q}$ and $\mathbf{R}$ the LQR controller, $\mathbf{K}_{\text{base}}$, does not satisfy Theorem 1. In this case, the MPC-fit controller converges to the reference value along a similar trajectory to the MPC controlled system. The control inputs of the MPC-fit controlled system and the MPC controlled system align well, as seen in Fig. 2.6b, further demonstrating that the static linear feedback controller can adequately reproduce the optimal inputs from a rolling horizon MPC controller for this system and choice of LQR objective function.

Comparison of Simplified and Full-order model

We compare simulation response of the full-order system described by (A.1), (A.2), (A.3) with the simplified system model used in this paper and described by (2.6). We control both of these systems with the MPC-fit controller and simulate their response to steps in P^*, Q^*. I_d^*, I_q^* are found from P^*, Q^* using (2.4). The full-order system parameter values are selected according to the methods used in [42, 36] are similar to those in Table 2.1 with $C = 1$ µF. We compare the grid-side currents $\mathbf{I}_{g\mathrm{dq}}$ of the full-order system to the currents

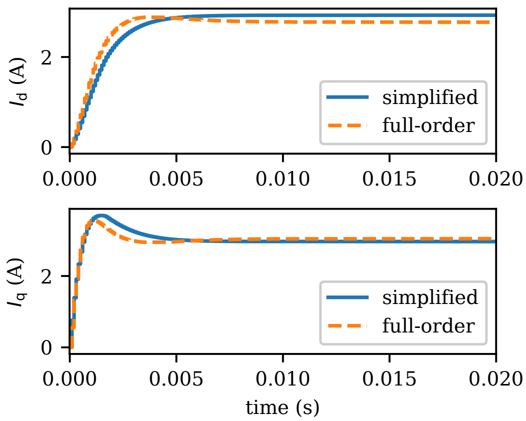

(a) Trajectories of the simplified system $\mathbf{I}_{\mathrm{dq}}$ and full-order system $\mathbf{I}_{g\mathrm{dq}}$ state values.

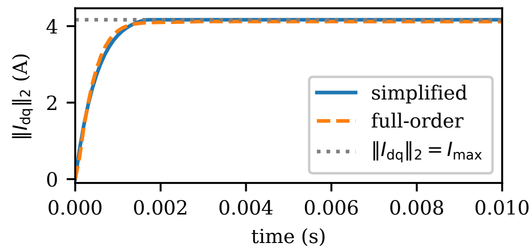

(b) Magnitude of the current for the simplified and full-order system.

Figure 2.7: dq current values and magnitude of the simplified and full-order systems for a step in (P^*, Q^*) from $(0\text{ W}, 0\text{ VAr})$ to $(775\text{ W}, -775\text{ VAr})$.

of the simplified model. Fig. 2.7a shows that both systems display similar trajectories and settling times. Discrepancies in the trajectories are likely due to the dynamics and steady-state response of the capacitor.

2.3 Safe Control of IBRs with Current Limits[6]

In this section, we demonstrate a control barrier function based safety filter method for current limiting of inverter-based resources controlled as a voltage source. The safety filter optimization problem is formulated with a control barrier function constraint that encodes the current magnitude limit. The feasibility of this optimization problem is guaranteed by the existence of a safe linear feedback controller for a given feasible reference. To guarantee the existance of a safe linear controller we use the dynamic properties of an inverter connected to a voltage source through an RL branch. This method aims to minimally alter the nominal control action to try to preserve the performance of a nominal controller, while safely limiting the current output at all times. We provide a simple closed-form solution to the safety filter problem that can be applied to a nominal control action resulting in a safe controller. Simulations are performed comparing an unsafe nominal controller to the safety filter controller and a safe linear feedback controller.

2.3.1 System Model & Control Objective

In this paper, we consider a three-phase inverter connected to an infinite bus via an RL branch. The model assumes we have balanced three-phase and can use a direct-quadrature (dq) reference frame with reference to some rotating angle to describe all the rotating physical quantities [56]. This model was introduced in Section 2.1.2, and is commonly used when studying the transient stability of inverters controlled as a voltage source.

For this work, we set $V = E$ and take $\mathbf{u} = \delta$ to be the only control input to the system. The states are still $\mathbf{x} = (I_\mathrm{d}, I_\mathrm{q})$. The active and reactive power outputs from the inverter are given by

$$
\begin{aligned}
P &= \frac{3}{2}\left(V\cos(\delta)I_\mathrm{d} + V\sin(\delta)I_\mathrm{q}\right), \\
Q &= \frac{3}{2}\left(V\sin(\delta)I_\mathrm{d} - V\cos(\delta)I_\mathrm{q}\right).
\end{aligned}
$$

[6] Adapted from T. Joswig-Jones and B. Zhang, "Safe control of grid-Interfacing inverters with current magnitude limits," in 2025 Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences, 2025 [55].

With the above notation and assumptions, tracking a given P^*, Q^* is equivalent to tracking some I_d^* and I_q^*.

Safety and Stability

To protect internal switching devices, the inverter's output current cannot exceed a given magnitude limit 2.3. To satisfy this bound we must always have $\|\mathbf{x}\|_2 \leq I_\text{max}$. We define a set of safe states in terms of the current magnitude bounds as $\mathcal{S} \coloneqq \{\mathbf{x} \mid h(\mathbf{x}) \geq 0\}$, where

$$h(\mathbf{x}) = I_\text{max}^2 - \|\mathbf{x}\|_2^2,$$

and the set of states at the boundary of this set as $\partial\mathcal{S} \coloneqq \{\mathbf{x} \mid h(\mathbf{x}) = 0\}$. We want the inverter to be able to operate near or at this bound to maximize the use of the inverter's capabilities especially during faults or transients when the grid requires support.

A control action is safe if it pushes the states such that they do not leave the safe set anytime they are at the boundary of the safe set. Mathematically, the set of safe control inputs for some $\mathbf{x} \in \mathcal{S}$, with dynamics $\dot{\mathbf{x}} = f(\mathbf{x}) + g(\mathbf{x}) \cdot \mathbf{u}$, can be defined as the values of u that satisfy the control barrier function (CBF) constraint

$$\dot{h}(\mathbf{x}) = \nabla h(\mathbf{x})^\top (f(\mathbf{x}) + g(\mathbf{x}) \cdot \mathbf{u}) \geq -\alpha(h(\mathbf{x})) \;\; \forall \;\; \mathbf{x} \in \mathcal{S}, \tag{2.14}$$

where $\nabla h(\mathbf{x})$ is the gradient of $h(\mathbf{x})$ with respect to $\mathbf{x}$, and α is a class $\mathcal{K}$ function (strictly increasing and $\alpha(0) = 0$) [57]. The inclusion of the term $\alpha(h(\mathbf{x}))$ allows the safety condition to be applied to the entire safe region instead of just at its boundary. We note that for our system $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ and $g(\mathbf{x}) = \mathbf{B}$.

To control the power output of the inverter to a desired value, we need the system to be able to track to a given set point, $\mathbf{x}^*$. Formally, we require the control to be asymptotically stable with respect to some control Lyapunov function (CLF) [58, 59]. Given a CLF, $V(\mathbf{x})$, the set of stabilizing control inputs for some $\mathbf{x} \in \mathcal{S}$ can be defined as the values of $\mathbf{u}$ satisfying the constraint

$$\dot{V}(\mathbf{x}) = \nabla V(\mathbf{x})^\top (f(\mathbf{x}) + g(\mathbf{x}) \cdot \mathbf{u}) \leq 0. \tag{2.15}$$

We want to control our system such that $\mathcal{S}$ is an invariant set and the control is stabilizing (i.e. given an initial condition $\mathbf{x}_0 \in \mathcal{S}$ the system will converge to a feasible $\mathbf{x}^*$ without the states leaving $\mathcal{S}$ at any point during their trajectory). Typically, a nominal voltage controller $\mathbf{u}_{\text{nom}}(x)$ that is stabilizing will not satisfy the safety constraint for references near the current magnitude boundary and will produce unsafe control actions.

2.3.2 Current Limiting Safety Filter

We assume that the reference point is feasible: $\|\mathbf{x}^*\|_2 \leq I_{\max}$, and $\mathbf{u}^*, \mathbf{x}^*$ satisfy

$$\mathbf{A}\mathbf{x}^* + \mathbf{B}\mathbf{u}^* = 0. \tag{2.16}$$

Our goal is to design a feedback controller such that the system in (2.2) is stable with respect to a given setpoint $\mathbf{x}^*$ ($\mathbf{x}(t) \to \mathbf{x}^*$ as $t \to \infty$), and is safe with respect to the current magnitude limit.

Safety Filter Formulation

A nonlinear control method often used to achieve safety guarantees with a well performing controller is a *safety filter*. A safety filter is a function that is applied to some control action that can detect unsafe control inputs that may lead to constraint violations and minimally modifies them to ensure safety [60]. In this section we introduce how such a safety filter can be applied to our system and provide guarantees on the feasibility of the safety filter problem constraints.

We can formulate this safety filter as a quadratic program (QP) as follows [61].

$$\bar{\mathbf{u}} = \arg\min_{\mathbf{u}} \|\mathbf{u} - \mathbf{u}_{\text{nom}}(\mathbf{x})\|_2^2 \tag{QP}$$

$$\text{s.t.} \quad \nabla h(\mathbf{x})^\top (f(\mathbf{x}) + g(\mathbf{x}) \cdot \mathbf{u}) \geq -\alpha h(\mathbf{x}) \tag{CBF}$$

$$\nabla V(\mathbf{x})^\top (f(\mathbf{x}) + g(\mathbf{x}) \cdot \mathbf{u}) \leq 0, \tag{CLF}$$

where $\bar{\mathbf{u}}$ is the filtered voltage control action, and we select α to be a positive constant. For this paper we use the CLF $V(\mathbf{x}) = (\mathbf{x} - \mathbf{x}^*)^\top (\mathbf{x} - \mathbf{x}^*)$.

Although QPs are considered to be simple optimization problems, using iterative solvers in real-time on inverters themselves is not trivial, because of microprocessor limitations and sampling/input delays [62, 63]. However, this form of QP has a closed-form solution [57].

In our case, because the optimization is over a scalar variable u, there is actually a simple closed-form solution (see Algorithm 1 of [55]). This algorithm comes from projecting a point into an interval, and can be implemented on existing inverter microcontrollers.

When using an optimization-based controller it is important to ensure that there exists at least one feasible solution at any possible state of the system. This is especially relevant to our problem where the safety and stability constraints often can seem to be working towards opposite goals (we want a controller that can track current references near the current magnitude limit, but never exceeds this current magnitude limit). In general, there is no guarantee that there is a feasible input to this safety filter problem for a generic linear system, but given the properties of our system and safety constraint we can guarantee the existence of a feasible solution. In the following section we prove that there always exists a feasible control action that satisfies (CBF) and (CLF) for (2.2).

Safety Filter Feasibility

The theorem below states the main result of the paper.

Theorem 3. *Given the linear system in* (2.2) *with input* $u = \delta$ *and assume that the setpoints* $\mathbf{x}^*$ *and* $\mathbf{u}^*$ *are feasible. Consider the magnitude barrier function* $h(\mathbf{x}) = I_{\max}^2 - \|\mathbf{x}\|_2^2$*. If* $\mathbf{A} + \mathbf{A}^\top \prec 0$*, and* $\mathbf{A}^{-1}\mathbf{B} \neq 0$*, then a safe and stable control actions always exists such that if* $h(\mathbf{x}_0) \geq 0$ *then* $h(\mathbf{x}_t) \geq 0$ *and* $\mathbf{x}_t \to \mathbf{x}^*$*.*

Our approach to proving that there always exists a feasible solution to the safety filter problem is to prove the existence of a safe and stable linear feedback controller, $\mathbf{K}$, for a given $\mathbf{x}^*, \mathbf{u}^*$. Because this controller exists, there is at least one feasible action $\mathbf{u}$ that satisfies (CBF) and (CLF), namely, the linear control action, $\mathbf{u} = \mathbf{u}^* - \mathbf{K}(\mathbf{x} - \mathbf{x}^*)$. The main condition on the system is that the A matrix is "stable" enough, namely $\mathbf{A} + \mathbf{A}^T \prec 0$ (this can be seen as the Lyapunov stability condition for linear systems with a Lyapunov function of $||\mathbf{x}||_2^2$). This result is intuitive in the following sense. Because the level sets of the Lyapunov function and the barrier function have the same shape—they are both circular—the tension between stability and safety can be resolved. The formal proof of this theorem is nontrivial, and the details are given in Appendix A.4.

Even though a safe and stable linear controller exists, there may be more optimal control actions that are safe and stabilizing. Imposing a safe linear controller over the entire safe region is overly conservative and can result in a suboptimal performance. Further, for any linear system and reference value there is a limited set of safe linear feedback controllers that satisfy the criteria required by **Theorem** 3. For some systems this set is small and can be restrictive when designing a controller to meet certain performance criteria. Therefore, we are motivated to look at the nonlinear safety filter controller when the control objective includes safety constraints and performance specifications. The ability of this safety filter to provide more optimal safe control actions compared to a safe linear controller can be seen in Section 2.3.3.

2.3.3 Simulation Results

To test this approach we simulate the response of the inverter system for a range of initial conditions and power references. We test the system with an LQR controller, a safe linear K controller, and the LQR controller with CBF safety filter. The parameters used for the system in these simulations are the same as those in Table 2.1 other than the current limit which is $|I|_{\max} = 5$ A. The code used to generate these results is available at `https://github.com/TragerJoswig-Jones/Safe-Current-Magnitude-Limit-Inverter-Control`.

The nominal LQR controller is computed with $\mathbf{Q} = \mathbf{I}$, $R = V_{\text{nom}}/(10 \cdot L)$ and is found to be $\mathbf{K}_{\text{LQR}} = \begin{bmatrix} 0.0009 & 0.0099 \end{bmatrix}$. To design a safe linear feedback controller we use the conditions of **Lemma** 2 to formulate the optimization problem

$$\begin{aligned} \min_{\mathbf{K},\lambda} \quad & \|\mathbf{K}\|_2 \\ \text{s.t.} \quad & (\mathbf{x}^*)^\top (\mathbf{A} - \mathbf{B}\mathbf{K}) = (\mathbf{x}^*)^\top \lambda, \\ & \lambda_{\max}((\mathbf{A} - \mathbf{B}\mathbf{K}) + (\mathbf{A} - \mathbf{B}\mathbf{K})^\top) \leq \lambda, \\ & (\mathbf{A} - \mathbf{B}\mathbf{K}) + (\mathbf{A} - \mathbf{B}\mathbf{K})^\top \prec 0, \end{aligned} \tag{2.18}$$

where we choose to minimize the spectral norm of $\mathbf{K}$ as the objective. Note that it is difficult to select an objective that designs a controller for performance with this form of static controller optimization problem. We also note that these constraints are convex in $\mathbf{K}$ and λ. We solve (2.18) using `CVXPY` [54, 64] to get $\mathbf{K} = \begin{bmatrix} -0.0111 & 0.0111 \end{bmatrix}$. The CBF

Table 2.2

AVERAGE COST WITH $\mathbf{x}^* = (3.56\text{ A}, 3.51\text{ A})$ AND $\mathbf{x}_0 \in X_0$

controller	$\mathbf{K}$	CBF	LQR
cost	82.22	59.16	58.57

safety filter, (QP), is ran with the nominal control action, $u_{\text{nom}} = u^* - \mathbf{K}_{\text{LQR}}(\mathbf{x} - \mathbf{x}^*)$, and $\alpha = 1000$, and is solved using the closed-form solution. Simulations are performed using the odeint solver from `SciPy` with a sampling time of $\Delta t = 10$ µs and simulation length of $t_{\text{end}} = 50$ ms. The costs for these tests are calculated as

$$\text{cost}_i = 1000 \cdot \sum_{t=0}^{t_{\text{end}}} \Delta t \left(\tilde{\mathbf{x}}_{i,t}^\top \mathbf{Q} \tilde{\mathbf{x}}_{i,t} + \tilde{u}_{i,t}^\top R \tilde{u}_{i,t} \right),$$

where $\tilde{\mathbf{x}}_{i,t} = \mathbf{x}_{i,t} - \mathbf{x}_{i,t}^*$ and $\tilde{u}_{i,t} = u_{i,t} - u_{i,t}^*$ are the state and input errors at time t for test i.

In the following we use the function $\text{lin}(a, b, n) \coloneqq a + \frac{n-1}{i-1}(b-a)$ to define a list of evenly spaced numbers over a specified interval.

$\mathbf{x}_0 \in \partial\mathcal{S}$ *tests*

We test the three controllers for a range of initial conditions from $\partial\mathcal{S}$, specifically $X_0 \coloneqq \{(I_{\max} \sin(\phi), I_{\max} \cos(\phi])) \mid \phi \in \text{lin}(0, 2\pi - \frac{2\pi}{100}, 100)\}$. We select a single $\mathbf{x}^*$ reference value for these tests, that is $\mathbf{x}^* = (3.56\text{ A}, 3.51\text{ A})$, by selecting a feasible $\mathbf{x}^*$ value from (2.16) with magnitude $I_{\max}$.

The average cost for each controller can be found in Table 2.2. The nominal LQR controller always performs optimally, but also always has a trajectory that is unsafe in these tests. The safe $\mathbf{K}$ controller is always safe, but performs poorly for the given LQR costs. Lastly, the CBF safety filter controller has a similar, yet slightly worse, performance to the LQR controller, but the trajectory always remains safe due to the safety filter. A selected example trajectory can be seen in Fig 2.8a. Notably, the CBF trajectory aligns with the LQR trajectory up until the safety constraint would be violated and the safety filter

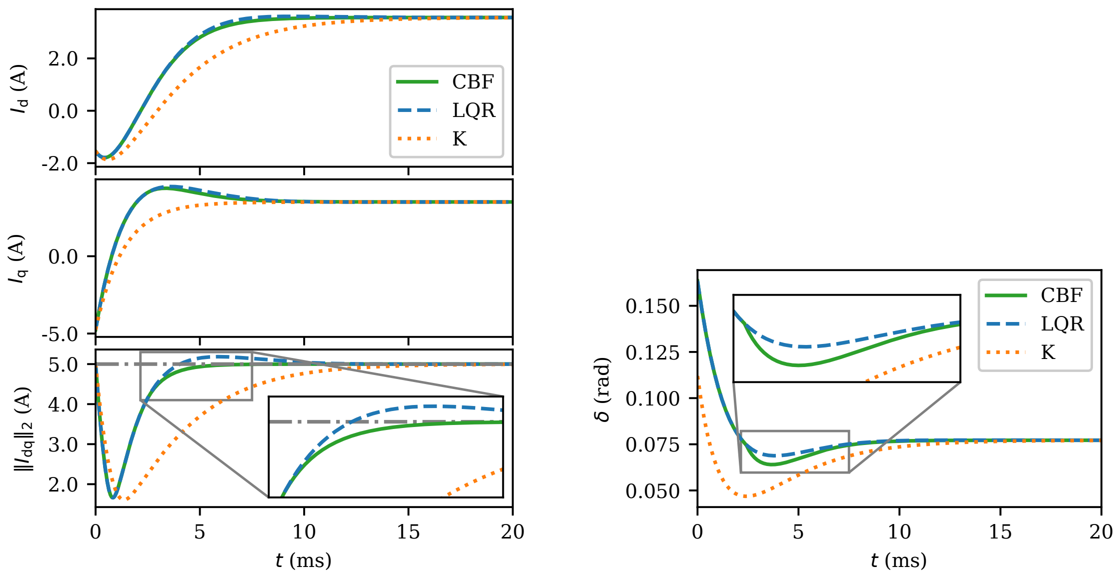

(a) $\mathbf{I}_{\mathrm{dq}}$ state values and magnitude trajectory.

(b) δ input values.

Figure 2.8: State and input trajectories for CBF, LQR, and $\mathbf{K}_{\mathrm{safe}}$ controllers with $\mathbf{x}_0 = (-1.55\ \mathrm{A}, -4.76\ \mathrm{A}), \mathbf{x}^* = (3.56\ \mathrm{A}, 3.51\ \mathrm{A})$. The LQR control is seen to exceed the safety bound, the safe $\mathbf{K}$ control has costly performance, and the CBF control performs well and remains safe.

intervenes. The alteration of the control action that causes this divergence in trajectories can be seen in Fig 2.8b.

Sampled $\mathbf{x}_0, \mathbf{x}^*$ *tests*

We then test these controllers with randomly sampled $\mathbf{x}_0$ and $\mathbf{x}^*$ values. We sample $\mathbf{x}^*$ reference values by selecting a feasible $\mathbf{x}^*$ value from (2.16) and scaling the magnitude to a uniformly sampled value from the range $[-I_{\max}, I_{\max}]$. $\mathbf{x}_0$ values are calculated as $\mathbf{x}_0 = (r_0 \cos(\phi_0), r_0 \sin(\phi_0))$ with r_0, ϕ_0 uniformly sampled from the ranges $[0, I_{\max}]$ and $[0, 2\pi]$, respectively.

Testing with 1,000 randomly sampled $\mathbf{x}_0$ and $\mathbf{x}^*$ values we see a similar trend in the average costs with $\mathbf{K}$ performing the worst, LQR control performing the best, and the CBF control performing slightly worse than the LQR controller. The distribution of costs from

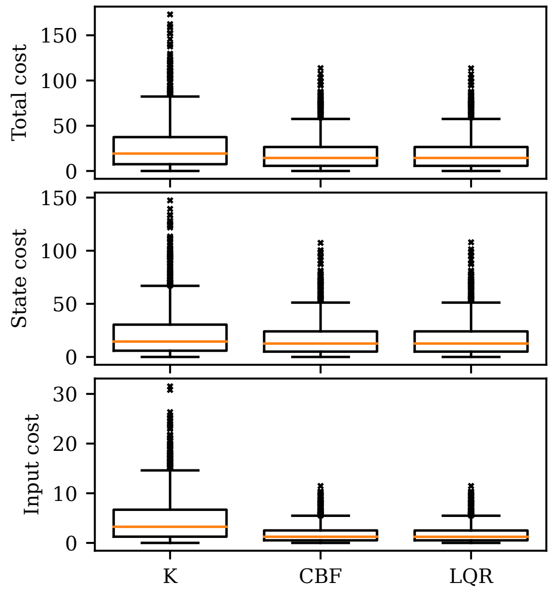

Figure 2.9: Costs of CBF, LQR, and $\mathbf{K}_{\text{safe}}$ controllers for 1,000 randomly sampled $\mathbf{x}_0$ and $\mathbf{x}^*$ values.

these tests can be seen in Fig 2.9. Of these tests 24 of them had unsafe trajectories with LQR control, while all tests were safe with the other two controllers. The difference in the LQR and CBF controller costs are small, but for all tests the CBF controller cost is equal or slightly higher than the LQR controller cost due to the safety filter intervening.

Small-angle assumption

Lastly, we test the validity of the small-angle assumption with the CBF control. In this section we compare the linear system used for analysis to the nonlinear system without the small-angle assumption (2.1). For the nonlinear system we use the CBF control formulated for the linear system. We subject the two systems to the same range of initial conditions as in Section 2.3.3, but take $\mathbf{x}^*$ to be (3.42 A,3.64 A), a feasible reference value for the nonlinear system (2.1).

An exemplary case can be seen in Fig 2.10a. The trajectories of the two systems do not differ significantly for any of the test cases, demonstrating that the small-angle assumption holds well with the CBF control. However, the trajectories do differ and the nonlinear

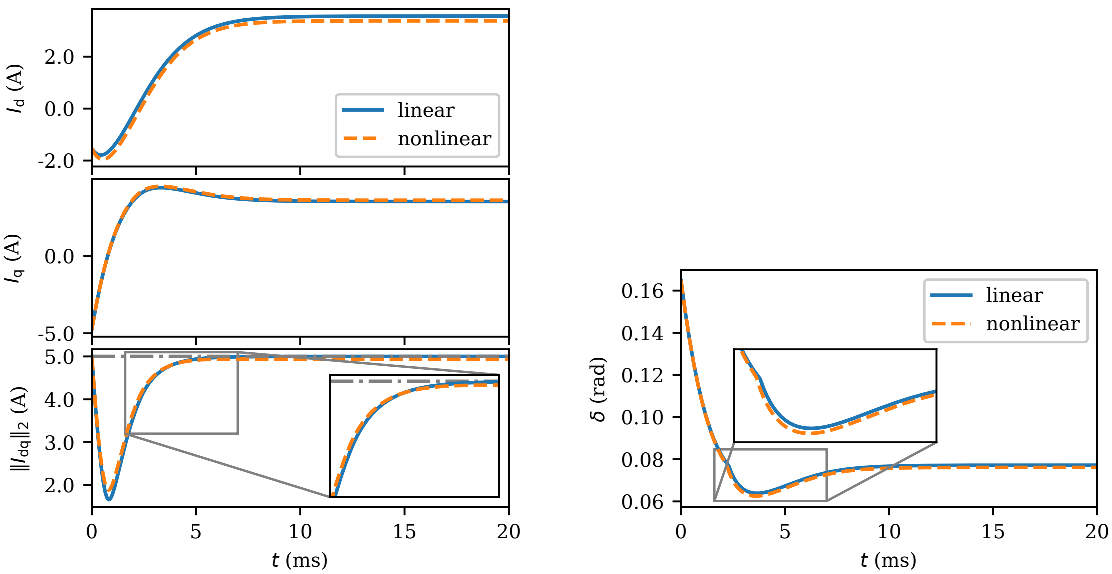

(a) I_{dq} trajectories.

(b) δ trajectories.

Figure 2.10: State and input trajectories for for the linear and nonlinear systems with CBF control and $\mathbf{x}_0 = (-1.55\text{ A}, -4.76\text{ A}), \mathbf{x}^* = (3.42\text{ A}, 3.64\text{ A})$. The linear and nonlinear systems have a similar trajectories.

system slightly exceeds the current magnitude bound by up to 0.5% in some of the tests. Further the nonlinear system is seen to not fully converge to the reference value due to the safety filter. Fig 2.10b shows the input trajectories are nearly aligned, but when the safety filter activates at 2 ms it introduces an offset between the inputs which prevents the states of the nonlinear system from completely approaching the magnitude bound.

These issues seen when applying the linear CBF control to the nonlinear system are small, but may not be tolerable. The boundary violation issue can be mitigated by setting the boundaries to be slightly tighter than the actual current limit values. A solution to resolve both issues would be to formulate the CBF control for the nonlinear system. However, this would introduce nonlinearity into the constraints of the safety filter problem making it more computationally expensive to solve in real time.

2.4 Geometry of the Feasible Output Regions of IBRs with Current Limits[7]

As seen in the previous two sections, current limits complicate the control of inverter-based resources. One topic that is not addressed in these works is the geometry of the feasible output region defined by current-magnitude limits. Understanding this geometry is important to the control problem, particularly when trying to determine how to control the inverter device to optimal steady-state outputs.

In this section, we study the geometry of the feasible output region of an inverter with current magnitude limits. That is, given a current limit and the inverter parameters, what is the region of all achievable outputs? Answering this question is the natural first step towards ensuring the safe operations of inverters and designing controllers that optimize their performances. We show that, under the simplified RL branch model, the feasible region is convex when two of active power, reactive power and square of the voltage magnitudes are considered. We provide an explicit characterization of this region using linear matrix inequalities. This analytical definition of the region can be used to efficiently find optimal operating points of the inverter.

Because power and voltage magnitude depend quadratically on the current, and the current is a vector in $\mathbb{R}^2$ (in the dq frame), the feasible region is a quadratic map of a disk (current magnitude is upper bounded). Unlike affine transformations, a quadratic map of a convex set is not, in general, convex. It turns out that the particular structure of the inverter circuit preserves convexity, which we will explain in detail in this paper. We demonstrate how knowing the feasible set and its convexity allows us to improve upon existing grid-forming inverters such that their steady-state currents always remain within the current magnitude limit, whereas standard grid-forming controllers can lead to instabilities and violations.

---

[7] This section is adapted from L. Streitmatter, T. Joswig-Jones, and B. Zhang, "Geometry of the feasible output regions of grid-interfacing inverters with current limits," IEEE Control Systems Letters, 2025 [65].

2.4.1 Model and Problem Formulation

System Model

This study uses a similar simplified model of a three-phase inverter connected to a fixed grid voltage represented by an infinite bus through an RL branch as introduced in Section 2.1.2 with affine dynamics 2.2. Note that for this work we take our state to be $\mathbf{I}_{\text{dq}}$, the inverter current output, with input $\mathbf{V}_{\text{dq}}$, the inverter voltage, both in the dq-reference frame of the grid [41]. This choice is mainly for notational convenience. We show in Section 2.4.2 that our overall approach and main results do not require the knowledge of the grid voltage angle and in fact we can take any angle to be the reference. The matrix $\mathbf{A}$ being skew-symmetric with equal diagonal elements is the key structural property of the inverter system that allows us to derive the main results.

Inverter Output

Inverters in the grid are usually provided set points which they try to track with their output, typically a combination of active power, reactive power, and voltage magnitude. The natural question becomes how well we can optimize these outputs given the constraint that current needs to stay in the safe region (or feasible set) $\mathcal{I}$, (2.3).

The active power, reactive power, and the squared voltage magnitude at the inverter terminals can all be expressed as quadratic functions of the current and the voltage:

$$P = \frac{3}{2}\mathbf{I}_{\text{dq}}^\top \mathbf{V}_{\text{dq}} \tag{2.19a}$$

$$Q = \frac{3}{2}\mathbf{I}_{\text{dq}}^\top \mathbf{J}\mathbf{V}_{\text{dq}} \tag{2.19b}$$

$$V^2 = \mathbf{V}_{\text{dq}}^\top \mathbf{V}_{\text{dq}}, \tag{2.19c}$$

where $\mathbf{J} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$. We note that here we consider the squared magnitude of the voltage. This is equivalent to considering the magnitude from the point of view of achieving a desired setpoint, but the squared form in (2.19c) is easier to work with.

The quantities in (2.19) are defined for all time t since $\mathbf{I}_{\text{dq}}$ and $\mathbf{V}_{\text{dq}}$ are functions of t, but we are usually interested in optimizing the equilibrium values resulting from the dynamics in (2.2). Assuming that (2.2) is asymptotically stable (we show this in Section 2.4.2), we denote the equilibrium of current and voltage as $\overline{\mathbf{I}}_{\mathrm{dq}}$ and $\overline{\mathbf{V}}_{\mathrm{dq}}$, respectively. Substituting these into (2.19), we get the equilibrium values of $\overline{P}$, $\overline{Q}$ and $\overline{V}^2$. We are interested in two related questions: 1) finding the optimal values of $\overline{P}$, $\overline{Q}$ and $\overline{V}^2$ and 2) using these values to improve the performance of existing grid-forming controllers.

We first look at the question of optimizing the equilibrium values. Let $S_1, S_2 \in (\overline{P}, \overline{Q}, \overline{V}^2)$ denote some choice of two out of the three quantities. We are interested in solving the following problem:

$$\begin{aligned} \text{minimize} \quad & f(S_1, S_2) \\ \text{subject to} \quad & \overline{\mathbf{I}}_{\mathrm{dq}} \in \mathcal{I}, \end{aligned} \tag{2.20}$$

where $f : \mathbb{R}^2 \to \mathbb{R}$ is some objective function. For example, suppose that we are interested in tracking some setpoints S_1^*, S_2^* as closely as possible. The objective would be $f(S_1, S_2) = \frac{1}{2}(S_1 - S_1^*)^2 + \gamma \frac{1}{2}(S_2 - S_2^*)^2$ with γ being some tradeoff parameter.

It is not immediately clear whether the problem (2.20) is easy to solve. Even if f is a convex function, the quadratic forms in (2.19) make the overall problem not convex in $\overline{\mathbf{I}}_{\mathrm{dq}}$.

Feasible Operating Regions

We can rewrite (2.20) as an optimization over S_1 and S_2:

$$\begin{aligned} \text{minimize} \quad & f(S_1, S_2) \\ \text{subject to} \quad & (S_1, S_2) \in \mathcal{S} \end{aligned} \tag{2.21}$$

where $\mathcal{S}$ is the set of all feasible points achieved by currents $\overline{\mathbf{I}}_{\mathrm{dq}} \in \mathcal{I}$. In the next section, we show that $\mathcal{S}$ is convex, and consequently, (2.21) is a convex optimization problem if f is convex. Using the tracking example again, (2.21) with $f(S_1, S_2) = \frac{1}{2}(S_1 - S_1^*)^2 + \gamma \frac{1}{2}(S_2 - S_2^*)^2$ can be solved efficiently at real-time, allowing us to achieve perfect tracking if possible, and the closest setpoints if not.

2.4.2 Geometry of the Feasible Region

In this section, we study the geometry of the feasible set $\mathcal{S}$. The main result is given by Theorem 4.

Theorem 4. *Let (S_1, S_2) be a pair of points formed by choosing any two of the three quantities $\overline{P}, \overline{Q}, \overline{V}^2$. Let $\mathcal{S} \in \mathbb{R}^2$ be the set of all achievable points (S_1, S_2) by $\bar{\mathbf{I}}_{\mathrm{dq}} \in \mathcal{I}$. Then $\mathcal{S}$ is convex.*

The proof of this theorem is given in [65], with an alternate proof provided in Appendix A.5. Fig. 2.11 plots what the regions look like for the combinations of (P, Q), (P, V^2), and (Q, V^2). The convexity of the regions do not depend on the values of R and L (as long as they are positive) or $\mathbf{E}_{\mathrm{dq}}$. The exact shape depends on R, L and $||\mathbf{E}_{\mathrm{dq}}||_2$, but not the phase angle of $\mathbf{E}_{\mathrm{dq}}$. Therefore, we do not need a PLL to sense the grid voltage angle. We do assume that inverters know their own RL filter specifications or can estimate them, as well as $||E_{\mathrm{dq}}||$ is measurable [66, 67].

Theorem 4 shows that for any two combinations of P, Q, V^2, the feasible region is convex. But it does not immediately provide a way to solve (2.21), since it does not directly give an algebraic description of the set. However, as the next theorem shows, we can describe the set using linear matrix inequalities.

Lemma 1. *Let $\mathbf{W}$ be a 3 by 3 matrix whose $(i, j)^{th}$ element is denoted by W_{ij} and define $\hat{C}$ as*

$$\hat{C} = \{(\mathrm{Tr}(\mathbf{M}_1\mathbf{W}), \mathrm{Tr}(\mathbf{M}_2\mathbf{W})| W_{11} + W_{22} \leq 1, W_{33} = 1, \\ \mathbf{W} \succeq 0\},$$

where ζ_1, ζ_2 are arbitrary constants and $\mathbf{I}_2$ is the 2 by 2 identity matrix and

$$\mathbf{M}_1 = \begin{bmatrix} \mathbf{I}_2 & \frac{1}{2}\mathbf{a} \\ \frac{1}{2}\mathbf{a}^\top & \zeta_1 \end{bmatrix} \text{ and } \mathbf{M}_2 = \begin{bmatrix} \mathbf{I}_2 & \frac{1}{2}\mathbf{b} \\ \frac{1}{2}\mathbf{b}^\top & \zeta_2 \end{bmatrix}, \tag{2.22}$$

Then with $\mathcal{C}$ defined as $\mathcal{C} = \left\{ \left(\alpha\mathbf{x}^\top\mathbf{x} + \mathbf{a}^\top\mathbf{x}, \beta\mathbf{x}^\top\mathbf{x} + \mathbf{b}^\top\mathbf{x} \right) \mid \mathbf{x}^\top\mathbf{x} \leq 1 \right\}$, we have $\mathcal{C} = \hat{\mathcal{C}}$.

Lemma 1 is proven in [65] and shows that we can find optimal steady-state current values by solving

$$\min f(\mathrm{Tr}(\mathbf{M}_1\mathbf{W},)\, \mathrm{Tr}(\mathbf{M}_2\mathbf{W})) \tag{2.23a}$$

$$\text{s.t. } W_{11} + W_{22} \leq I_{\max}^2 \tag{2.23b}$$

$$W_{33} = 1 \tag{2.23c}$$

$$\mathbf{W} \succeq 0, \tag{2.23d}$$

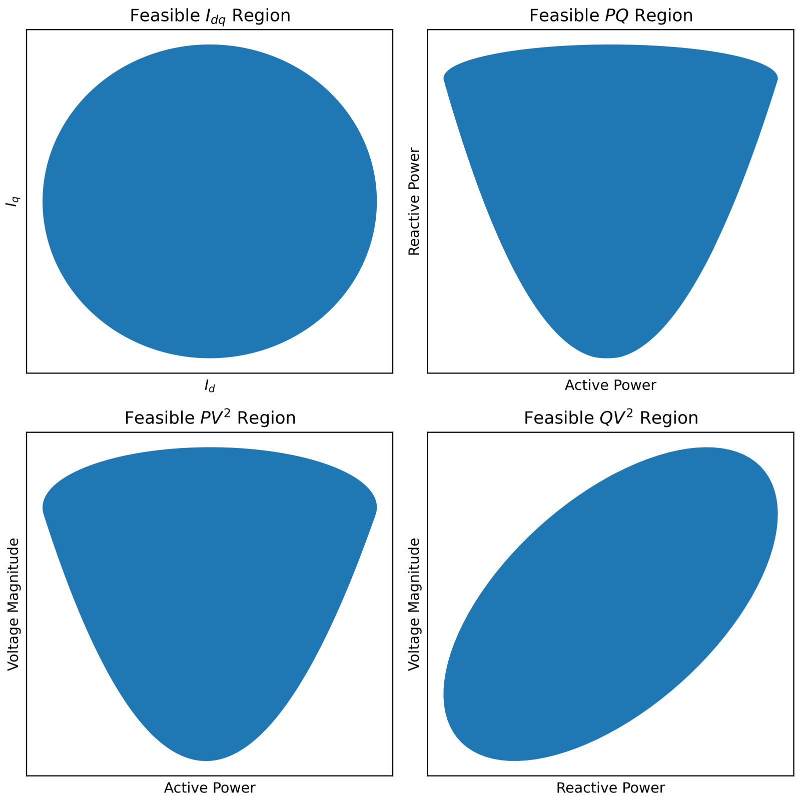

Figure 2.11: The set of feasible inverter currents ($\mathcal{I}$) (top left) always forms a circle of radius $I_{\max}$ due to the magnitude constraint on $\mathbf{I}_{\mathrm{dq}}$. The shapes of feasible (P, Q)(top right), (P, V^2)(bottom left), and (Q, V^2)(bottom right) regions depend on specific network parameters, namely RL filter impedance values and grid-side voltage magnitude $||\mathbf{E}_{\mathrm{dq}}||_2$.

where $\mathbf{M}_1, \mathbf{M}_2$ can be chosen to be $\begin{bmatrix} \frac{3}{2}R\mathbf{I}_2 & \frac{3}{4}\mathbf{E}_{\mathrm{dq}} \\ \frac{3}{4}\mathbf{E}_{\mathrm{dq}}^\top & 0 \end{bmatrix}$ for P, $\begin{bmatrix} \frac{3}{2}\omega L\mathbf{I}_2 & \frac{3}{4}\mathbf{J}^\top\mathbf{E}_{\mathrm{dq}} \\ \frac{3}{4}\mathbf{E}_{\mathrm{dq}}^\top\mathbf{J} & 0 \end{bmatrix}$ for Q and $\begin{bmatrix} (R^2+\omega^2L^2)\mathbf{I}_2 & \mathbf{A}^\top\mathbf{E}_{\mathrm{dq}} \\ \mathbf{E}_{\mathrm{dq}}^\top\mathbf{A} & ||\mathbf{E}_{\mathrm{dq}}||_2^2 \end{bmatrix}$ for $\mathbf{V}_{\mathrm{dq}}^2$.

So far we have shown that given a feasible $\mathbf{W}$ to (2.23), we can find a feasible current $\bar{\mathbf{I}}_{\mathrm{dq}}$ that achieves the same objective value. However, the optimal solution to (2.23) may not be rank 1. We can easily solve a quadratic equation to find the corresponding $\bar{\mathbf{I}}_{\mathrm{dq}}$, but an interesting observation from simulations is that the solution to (2.23) is always rank 1 (possibly with a small regularization term on the nuclear norm of $\mathbf{W}$). An important part of our future work is to rigorously prove this observation.

We also remark that (2.23) can be solved with data that are easy to obtain in practice. The grid voltage magnitude $||E_{\mathrm{dq}}||_2$ is easy to measure and R, L values can be estimated with existing inverter hardware if not already known. Since the feasible region in (2.23) is invariant to the rotation of the angle in E_{eq}, an arbitrary phase angle (e.g., 0) can be chosen.

Inverter Control

Given the steady-state values $\overline{\mathbf{I}}_{\mathrm{dq}}$ for (2.2), we need to find a controller such that the current will reach this equilibrium. If we have full information, including the phase angle of the grid voltage, we could use a linear feedback control on $\mathbf{V}_{\mathrm{dq}}$:

$$\dot{\mathbf{V}}_{\mathrm{dq}} = -k_v(\mathbf{V}_{\mathrm{dq}} - \overline{\mathbf{V}}_{\mathrm{dq}}), \tag{2.24}$$

where k_v is a positive proportional gain constant. It is easy to show that the closed-loop system is always stable (e.g., by computing the eigenvalues), and the trajectories remain within $\mathcal{I}$ with a correctly chosen k_v [55].

If full information is not available, the following section demonstrates how our approach can be used to provide an optimal setpoint to existing grid-forming controllers that do not need to know the system parameters [68, 69]. Thus, our approach advances the grid-forming control technology by guaranteeing safe and optimal steady-state setpoints.

2.4.3 Case Study

To demonstrate our approach, we solve (2.23) for both feasible and infeasible reference setpoints, then use different controllers to achieve the optimal setpoints. We evaluate the performance of a standard grid-forming droop controller adapted from [70] given an infeasible setpoint against the same controller provided with an updated, optimal setpoint from (2.23). We then implement the linear controller in (2.24) to show that both transient and steady-state currents remain safe in the idealized case when full system information is available.

Table 2.3
P^*, Q^*, V^{2*} SETPOINTS USED IN SIMULATIONS

		P^*(W)	Q^*(VAr)	V^{2*}(V^2)
PQ	Pre-Disturbance	800	0	–
Tracking	Post-Disturbance	**1100**	0	–
PV^2	Pre-Disturbance	200	–	120^2
Tracking	Post-Disturbance	**850**	–	120^2
QV^2	Pre-Disturbance	–	0	120^2
Tracking	Post-Disturbance	–	$-\mathbf{500}$	120^2

To solve (2.23) for different $(P^*, Q^*), (P^*, V^{2*})$ and (Q^*, V^{2*}) setpoints, we select the appropriate matrices $\mathbf{M}_1$, $\mathbf{M}_2$ and use the `CLARABEL` solver in `CVXPY` [54]. From the optimal solution $\mathbf{W}^*$, we extract $\overline{\mathbf{I}}^*_{\text{dq}}$ and solve for the corresponding power and/or voltage equilibrium values that the controller needs to track.

All simulations are conducted using the `SciPy` odeint solver with a sampling time of $\Delta t = 100\mu s$ and simulation period of $t_{\text{end}} = 1s$. Reference setpoints used in the simulations are provided in Table 2.3 where the disturbance is modeled as an update to the setpoint values at time t_0 from pre- to post-disturbance values.

Grid-Forming Droop Control Comparison

The PQ droop controller first filters the inverter's active and reactive power (P,Q) through a low-pass filter with cut-off frequency ω_c to measure $\tilde{P}$ and $\tilde{Q}$. These filtered values have dynamics $\dot{\tilde{P}} = \omega_c(P - \tilde{P})$, $\dot{\tilde{Q}} = \omega_c(Q - \tilde{Q})$.

Then, voltage magnitude droop and frequency droop are implemented as $\dot{\mathbf{V}}_{\text{dq}} = m_q\omega_c(\tilde{Q} - Q)$ and $\Delta\omega_i = -m_p(\tilde{P} - P^*)$, where $\Delta\omega_i$ is the deviation of inverter frequency from grid frequency, and m_p and m_q are proportional gain parameters. These follow from $\Delta\mathbf{V}_{\text{dq}} = -m_q(\tilde{Q} - Q^*)$, where $\Delta\mathbf{V}_{\text{dq}}$ is the deviation in voltage magnitude from a desired $\mathbf{V}_{\text{dq,nom}}$ value.

The PV^2 droop controller has the same low-pass PQ filter and power-frequency dynamics as the PQ droop controller, but it replaces the reactive power-voltage dynamics with linear feedback on $\mathbf{V}_{\mathrm{dq}}^2$ as $\dot{\mathbf{V}}_{\mathrm{dq}}^2 = -m_{v^2}(\mathbf{V}_{\mathrm{dq}}^2 - \mathbf{V}_{\mathrm{dq}}^{2*})$. Parameters for the PQ and PV^2 droop controllers can be found in [65].

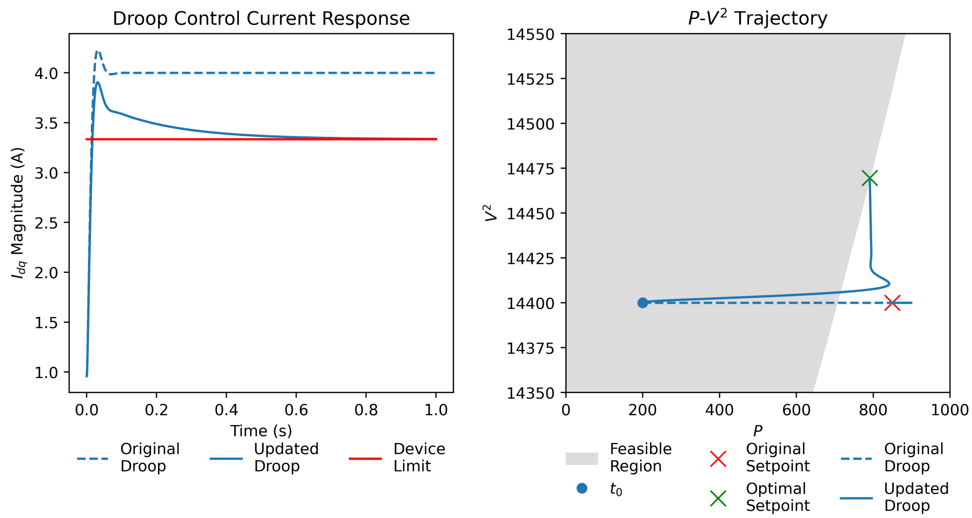

Figure 2.12: The two plots compare the current magnitude response (top) and (P, V^2) region trajectory (bottom) for the same droop controller tracking two different setpoints: the original infeasible setpoint (dashed blue line) and the best feasible setpoint output from the optimization problem (solid blue line).

Fig. 2.12 shows the unsafe transient *and* steady-state response of the original PV^2 droop controller to an infeasible setpoint of $P^* = 850W, V^{2*} = 120^2V$. In contrast, when we solve (2.23) to provide an updated, safe (P^*, V^{2*}) setpoint for the droop controller, we allow the grid-forming inverter to settle to the optimal feasible operating point. This prevents a consistent and unsafe violation of current limits in steady-state. The transient response violates the current constraints for a short time, which might be tolerable [48, 71] or can be mitigated with other techniques [50, 20].

Linear Feedback Optimal Controller (OC)

The linear feedback controller in (2.24) tracks $\overline{\mathbf{V}}_{\mathrm{dq}}^* = \mathbf{E}_{\mathrm{dq}} - L \cdot \mathbf{A}\, \overline{\mathbf{I}}_{\mathrm{dq}}^*$. We implement feedback on $\mathbf{V}_{\mathrm{dq}}$ with $k_v = 10$. We simulate two simple examples of the optimal controller tracking PQ and PV^2 setpoints. Fig. 2.13 and Fig. 2.13b show the inverter current response following

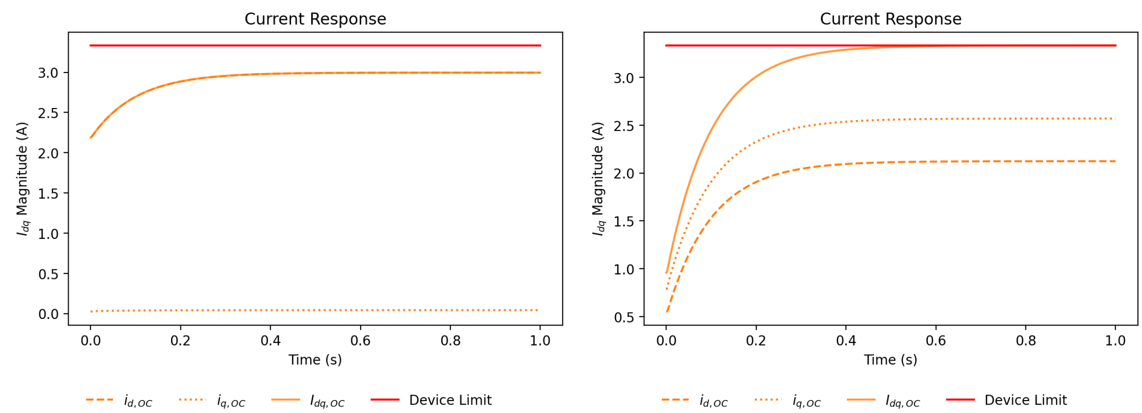

(a) The response of the optimal controller current magnitude to a feasible (P^*, Q^*) setpoint.

(b) The response of the optimal controller to an infeasible (P^*, V^{2*}) setpoint.

Figure 2.13: The response of the optimal controller current magnitude to a feasible and infeasible setpoints. In both simulations, the inverter current magnitude remains within the safe operating region. The infeasible setpoint in the bottom plot causes the inverter current to settle at the boundary of the safe operating region.

a feasible change in the PQ tracker's setpoint and an infeasible change in the PV^2 tracker's setpoint, respectively (see Table 2.3 for details). The linear feedback controller ensures both the transient and steady-state currents remain within the safe operating region of the inverter.

These advantages of the optimal controller motivate further investigation into whether we can control the inverter to achieve safety at all times without knowing the grid voltage magnitude or with different output filters. This is addressed in [72].

2.5 Safe Trajectory Gradient Flow Control of IBRs with Constraints[8]

Optimization-based control methods that steer a control system to a solution of an optimization problem have been studied in the context of constrained dynamical systems [74]. Recent works have proposed online feedback controllers that continuously regulate a system to the solution of an optimization problem with constraints on the state at steady-state [75, 76]. However, applying these methods to dynamic systems may lead to constraint violations during the transient. To solve this issue the approach in [77] uses high-order control barrier functions with safe gradient flows to ensure their feedback-optimization method enforces state and input constraints at all times.

In this section, we present a control framework that directly incorporates constraints into the control of a voltage-source inverter. We propose a safe trajectory gradient flow controller, which applies the safe gradient flow method to a rolling horizon trajectory optimization problem to ensure that the states remain within a safe set defined by the constraints while directing the trajectory towards an optimal equilibrium point of a nonlinear program. Simulation results demonstrate that our approach can drive the outputs of a simulated inverter system to optimal values and maintain state constraints, even when using a limited number of optimization steps per control cycle.

2.5.1 Notation

Given $f : \mathbb{R}^n \to \mathbb{R}^m$, let $\frac{\partial f}{\partial \mathbf{x}} \in \mathbb{R}^{m \times n}$ be the Jacobian of $f(\mathbf{x})$. When $m = 1$, $\nabla_{\mathbf{x}} f$ denotes the gradient. Consider a nonlinear program (NLP)

$$\min_{\mathbf{x}} \quad c(\mathbf{x}) \tag{2.25a}$$

$$\text{s.t.} \quad g(\mathbf{x}) \leq 0 \tag{2.25b}$$

$$h(\mathbf{x}) = 0 \tag{2.25c}$$

where $\mathbf{x} \in \mathbb{R}^n$, $g : \mathbb{R}^n \to \mathbb{R}^l$, and $h : \mathbb{R}^n \to \mathbb{R}^k$. Denote the feasible set as $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n | g(\mathbf{x}) \leq 0, h(\mathbf{x}) = 0\}$. Let $\mathbf{x}^*$ denote a locally optimal point that satisfies the KKT

[8] Adapted from T. Joswig-Jones and B. Zhang, "Safe Trajectory Gradient Flow Control of a Grid-Interfacing Inverter," IEEE Power & Energy Society General Meeting, 2026, in press [73].

conditions for (2.25) and X_{KKT} to be the set of all these points.

2.5.2 Safe Gradient Flow

One approach for solving (2.25) that is guaranteed to return a feasible solution regardless of when it is terminated, is referred to as *safe gradient flow* [78, 79]. This approach is shown to be asymptotically stable with equilibria in X_{KKT}, if the constraints satisfies the so-called Linear Independence Constraint Qualification Condition. Consider the control-affine system

$$\dot{\mathbf{x}} = -\nabla_{\mathbf{x}} c(\mathbf{x}) - \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}}^{\top} \mu - \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}}^{\top} \nu, \tag{2.26}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the state, and $\mu \in \mathbb{R}^l, \nu \in \mathbb{R}^k$ are inputs. Here the gradient of c pushes $\mathbf{x}$ in a direction that minimizes the cost function, while the additional inputs are used to ensure that the state remains within the safe set $\mathcal{C}$.

Applying CBF theory here, the admissible control set that ensures $\mathcal{C}$ is invariant is

$$\begin{aligned} K(\mathbf{x}) = \{(\mu, \nu) \ \mid \\ & -\frac{\partial g}{\partial \mathbf{x}} \frac{\partial g}{\partial \mathbf{x}}^{\top} \mu - \frac{\partial g}{\partial \mathbf{x}} \frac{\partial h}{\partial \mathbf{x}}^{\top} \nu \leq \frac{\partial g}{\partial \mathbf{x}} \nabla c(\mathbf{x}) - \alpha(g(\mathbf{x})), \\ & -\frac{\partial h}{\partial \mathbf{x}} \frac{\partial g}{\partial \mathbf{x}}^{\top} \mu - \frac{\partial h}{\partial \mathbf{x}} \frac{\partial h}{\partial \mathbf{x}}^{\top} \nu = \frac{\partial h}{\partial \mathbf{x}} \nabla c(\mathbf{x}) - \alpha(h(\mathbf{x}))\}. \end{aligned}$$

A feedback controller to find minimally altering values of μ, ν is synthesized as the quadratic program (QP)

$$(\mu(\mathbf{x}), \nu(\mathbf{x})) = \underset{\mu, \nu \in K(\mathbf{x})}{\operatorname{argmin}} \left\{ \left\| \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}}^{\top} \mu + \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}}^{\top} \nu \right\|^2 \right\} \tag{2.27}$$

This feedback controller finds the minimum alteration to the gradient descent term required to make the set $\mathcal{C}$ invariant while driving the state to a locally optimal point.

2.5.3 System Model

For this approach, we consider the simplified model of a three-phase inverter connected to an infinite bus via an RL branch as shown in Fig. 2.2 and first introduced in Section 2.1.2. The model assumes we have balanced three-phase on the AC-side and can use a direct-quadrature (dq) reference frame with reference to the inverter voltage which rotates with frequency ω [41]. Here we consider the system dynamics with reference to the inverter voltage angle, such that the dynamics are

$$f(\mathbf{x},\mathbf{u}) = \begin{bmatrix} \dot{I}_d \\ \dot{I}_q \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} -\frac{R}{L}I_\mathrm{d} + \omega I_\mathrm{q} + \frac{\sqrt{2}}{L}(V - E\cos(\delta)) \\ -\frac{R}{L}I_\mathrm{q} - \omega I_\mathrm{q} + \frac{\sqrt{2}}{L}(-E\sin(\delta)) \\ \omega_e - \omega \end{bmatrix}. \tag{2.28}$$

$\mathbf{I}_\mathrm{dq} = (I_\mathrm{d}, I_\mathrm{q})$ and $\mathbf{V}_\mathrm{dq} = (\sqrt{2}V, 0)$ are the inverter current and voltage, respectively, in the inverter reference frame. The grid voltage rotates with frequency ω_e and δ is the angle difference between the grid voltage and inverter voltage. $\mathbf{E}_\mathrm{dq} = (\sqrt{2}E\cos(\delta), \sqrt{2}E\sin(\delta))$ defines the grid voltage in the inverter reference frame. We take the states to be $\mathbf{x} = (I_\mathrm{d}, I_\mathrm{q}, \delta)$ and the control inputs to be $\mathbf{u} = (V, \omega)$. P, Q are the active and reactive powers injected into the grid, computed as $P = \frac{3}{2}(E\cos(\delta)I_\mathrm{d} + E\sin(\delta)I_\mathrm{q})$, $Q = \frac{3}{2}(E\sin(\delta)I_\mathrm{d} - E\cos(\delta)I_\mathrm{q})$. To prevent the semiconductor switches of the inverter from becoming damaged, the magnitude of the current must be limited according to (2.3):

$$g(\mathbf{x}) = \|\mathbf{I}_\mathrm{dq}\|_2^2 - I_\mathrm{max}^2 \leq 0. \tag{2.29}$$

We want to control the inverter such that we inject close to some given desired active and reactive power values P^*, Q^* without significantly deviating from the nominal voltage and frequency values $V_\mathrm{nom}, \omega_\mathrm{nom}$. We formulate our cost as

$$\begin{aligned} c(\mathbf{x},\mathbf{u}) =& \frac{m_p}{2}(P - P^*)^2 + \frac{m_q}{2\cdot\tau_v}(Q - Q^*)^2 \\ &+ \frac{1}{2\cdot\tau_v}(V - V_\mathrm{nom})^2 + \frac{1}{2}(\omega - \omega_\mathrm{nom})^2 \end{aligned} \tag{2.30}$$

where m_p, m_q, τ_v are the active power droop term, reactive power droop term, and voltage control speed term, respectively.

Our goal is to define a controller that drives the system to optimal equilibrium points, which is equivalent to solving:

$$\min_{\mathbf{x},\mathbf{u}} \quad c(\mathbf{x},\mathbf{u}) \tag{2.31a}$$

$$\text{s.t.} \quad g(\mathbf{x}) \leq 0 \tag{2.31b}$$

$$f(\mathbf{x},\mathbf{u}) = 0. \tag{2.31c}$$

We note that we cannot directly apply the SGF approach here as our dynamical system (2.28) does not directly align with (2.26) and we may not always be able to fully control $f(\mathbf{x}, \mathbf{u})$ such that it flows according to the gradient of the cost.

2.5.4 Safe Trajectory Gradient Flow

We design a controller using the *safe gradient flow* method by applying it to a rolling horizon trajectory optimization problem. We discretize (2.31) to obtain the following trajectory optimization problem that is related to (2.31). In this section we consider a given discrete time system

$$\mathbf{x}_{\tau+1} = f(\mathbf{x}_\tau, \mathbf{u}_\tau), \tag{2.32}$$

where $\mathbf{x}_\tau$ refers to the state at timestep τ. First, we formulate the trajectory optimization problem

$$\begin{aligned}
\min_{\mathbf{x},\mathbf{u}} \quad & c_f(\mathbf{x}_{t+T}) + \sum_{\tau=t}^{t+T} c(\mathbf{x}_\tau, \mathbf{u}_\tau) \\
\text{s.t.} \quad & g(\mathbf{x}_\tau) \leq 0 \ \ \forall \ \ \tau \in (t, t+T] \\
& h(\mathbf{x}_\tau) = 0 \ \ \forall \ \ \tau \in (t, t+T] \\
& \mathbf{x}_{\tau+1} = f(\mathbf{x}_\tau, \mathbf{u}_\tau,) \ \ \forall \ \ \tau \in (t, t+T-1]
\end{aligned} \tag{2.33}$$

where T is the horizon of the trajectory being optimized, and $\mathbf{x} \in \mathbb{R}^{n\times T}, \mathbf{u} \in \mathbb{R}^{m\times T}$ are the trajectories of the states and inputs, respectively. $c_f(\mathbf{x})$ is an altered cost function that only depends on the final state on the horizon T by removing the terms with the inputs V, ω. Next we denote

$$G(\mathbf{w}) = \begin{bmatrix} g(\mathbf{x}_t) \\ \vdots \\ g(\mathbf{x}_{t+T}) \end{bmatrix}, H(\mathbf{w}) = \begin{bmatrix} h(\mathbf{x}_t) \\ \vdots \\ h(\mathbf{x}_{t+T}) \\ \mathbf{x}_{t+1} - f(\mathbf{x}_t, \mathbf{u}_t) \\ \vdots \\ \mathbf{x}_{t+T} - f(\mathbf{x}_{t+T-1}, \mathbf{u}_{t+T-1}) \end{bmatrix},$$

and $C(\mathbf{w}) = c_f(\mathbf{x}_{t+T}) + \sum_{\tau=t}^{t+T} c(\mathbf{x}_\tau, \mathbf{u}_\tau)$, where $\mathbf{w} = (\mathbf{x}, \mathbf{u})$. With this notation we can fit this into the form of (2.25) as

$$\min_{\mathbf{w}} \quad C(\mathbf{w}) \tag{2.34a}$$

$$\text{s.t.} \quad G(\mathbf{w}) \leq 0 \tag{2.34b}$$

$$H(\mathbf{w}) = 0 \tag{2.34c}$$

Applying the safe gradient flow method to this trajectory optimization problem steps the trajectory $(\mathbf{x}, \mathbf{u})$ towards a locally optimal trajectory $(\mathbf{x}^*, \mathbf{u}^*)$. However, we want to implement a controller in real time and may only be able to update the trajectory for a limited number of times per step.

To solve this in real time and apply the control inputs to our system, we adopt a rolling horizon implementation inspired by model predictive control. At each time step, we formulate the above optimization problem using the current state as the initial condition, perform a limited number of updates, K, (i.e. gradient steps) toward minimizing the cost subject to constraints, apply the first control action from this solution, shift the horizon forward, and repeat the process. This process is detailed in Algorithm 1, which we call a *safe trajectory gradient flow* (STGF) controller.

This results in a feedback control policy that iteratively improves its trajectory while always maintaining feasibility (retaining the anytime property central to safe gradient flows), if the predicted state dynamics align with the true state dynamics [9] and a sufficiently small step size is used. Although the full optimal trajectory may not be reached at each step, the quality of the control inputs improves over time, and the feasible set remains invariant throughout.

2.5.5 Simulation Results

We test our STGF controller by apply the approach to the simplified inverter system with the per unitized parameters shown in Table 2.4 and compare it to model predictive control and a modified droop controller. To apply our STGF approach we discretize our dynamics

[9] Measuring exogenous values and estimating the system's true dynamics are important challenges with this approach, which we reserve for future work.

Algorithm 1 Safe Trajectory Gradient Flow

Input: Initial state $\mathbf{x}_0 \in \mathbb{R}^n$, Initial input trajectory $\mathbf{u} \in \mathbb{R}^{(T-1)\times m}$, horizon length T, termination time N, optimization updates per step K, dynamics $f(\mathbf{x}, \mathbf{u})$, constraints $(g(\mathbf{x}, \mathbf{u}), h(\mathbf{x}, \mathbf{u}))$, cost function $c(\mathbf{x}, \mathbf{u})$, time step Δt, optimization step-size ξ, CBF parameter α

Initialize trajectory: $\mathbf{x} \in \mathbb{R}^{T\times n} \leftarrow (\mathbf{x}_0, f(\mathbf{x}_0, \mathbf{u}_0), \ldots, f(\mathbf{x}_{T-1}, \mathbf{u}_{T-1}))$
Create C, G, H from applying c, g, h, f at each timestep in the horizon T
while $t \leq N$ **do**
    Measure current timestep state $\mathbf{x}_t$
    $\mathbf{x}_0 \leftarrow \mathbf{x}_t$
    Propagate state trajectory with predicted dynamics:
    for $\tau = 0$ to $T-1$ **do** $\mathbf{x}_{\tau+1} \leftarrow f(\mathbf{x}_\tau, \mathbf{u}_\tau)$
    end for
    for $j = 0$ to K **do**
        $\mathbf{w} \leftarrow (\mathbf{x}, \mathbf{u})$
        $\mu, \nu \leftarrow$ Solve QP (2.27) with $\mathbf{w}, C, G, H, \alpha$
        $\dot{\mathbf{w}} \leftarrow -D_{\mathbf{w}} c(\mathbf{w}) - \frac{\partial g(\mathbf{w})}{\partial \mathbf{w}}^\top \mu - \frac{\partial h(\mathbf{w})}{\partial \mathbf{w}}^\top \nu$
        $\mathbf{x}, \mathbf{u} \leftarrow \mathbf{w} + \xi \cdot \dot{\mathbf{w}}$
    end for
    Apply control input $\mathbf{u}_0$ from the optimized trajectory
    Roll forward control inputs:
    for $\tau = 0$ to $T-2$ **do** $u_\tau \leftarrow u_{\tau+1}$
    end for
    $u_{T-1} \leftarrow u_{T-2}$
    $t \leftarrow t + 1$
end while

Table 2.4

SIMULATION PARAMETERS

Parameter	Value	Cost Parameter	Value
R	0.0069 pu	m_p	2π
L	0.0196 pu	m_q	0.05
$I_{\max}$	1 pu	τ_p	0.1
Base	Value	Opt. Parameter	Value
ω_{base}	$2\pi 60$ rad s^{-1}	Δt	1 ms
V_{base}	120 V	ξ	1×10^{-3}
S_{base}	1500 V A	T	10
I_{base}	4.167 A	K	2

(2.28) using forward Euler with $\Delta t = 1$ ms. A control horizon of $T = 10$ is selected and each time step only a single iteration of the trajectory optimization problem is performed with $K = 2$ and $\alpha(z) = 20e^{10z}$. The STGF QP (2.27) is formulated with `CVXPY` [54] and solved using `OSQP` [80]. The dynamics of the inverter system are simulated using the `solve_ivp` function from `SciPy` [81]. We compare our approach to model predictive control (MPC) implemented with `do-mpc` [82, 83, 84], and a droop controller modified with a current limiter and an additional synchronization term [29, 85]. We assume the controllers have perfect knowledge of the grid voltage and frequency in each timestep.

Trajectory Performance & Steady-State Optimality

We test the system by simulating a step in the active and reactive power reference. The simulation lasts $N = 300$ time steps (i.e. 0.3 s) with initial values $\mathbf{x}_0 = \mathbf{0}, u_0 = \mathbf{0}$, and reference values $P^* = 0$ p.u., $Q^* = 0$ p.u., $V_{\text{nom}} = 1$ p.u., $\omega_{\text{nom}} = 2\pi \cdot 60$ rad s^{-1}. At $T = 100$, we step the active and reactive power references to infeasible values $P^* = 2.5$ p.u., $Q^* = -0.5$ p.u.. Fig. 2.14 shows the trajectory of I_{dq} with the STGF controller approaching the desired reference values while remaining within the current magnitude constraint. In Fig. 2.15, the output cost parameters are seen to converge to optimal values calculated by solving (2.31)

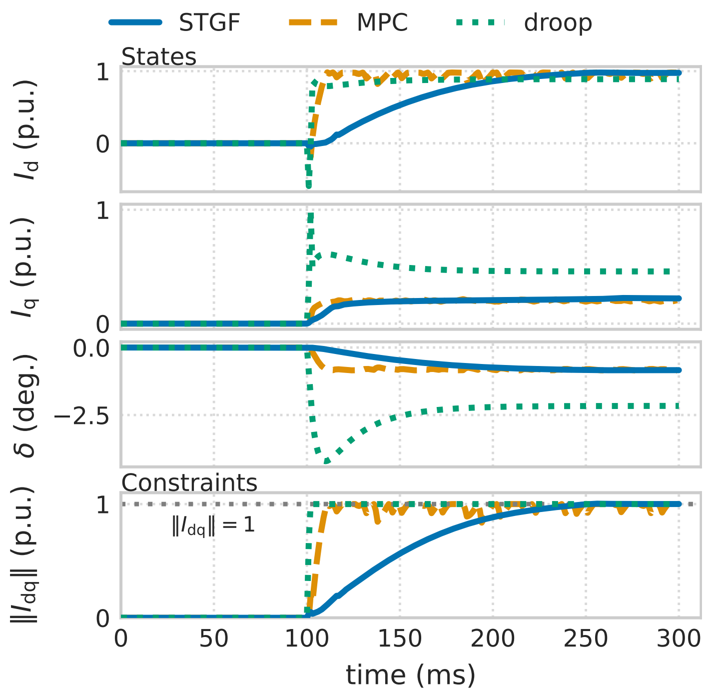

Figure 2.14: The states and current magnitude from the inverter system controlled with different control methods. While the MPC and droop controllers act are more aggressive, the STGF control drives the states smoothly to their new equilibria.

directly using a NLP solver. The trajectory of these output cost parameters along with the current magnitude constraint is visualized in Fig. 2.16. The MPC controller also approaches the optimal values and does so faster than the STGF controller, but struggles to maintain consistent control inputs when the nonlinear current limit constraint is active. The droop controller also acts faster, but does not settle to the optimal outputs.

Ideally, one would be able to select a horizon within which the system would be able to settle to a new equilibrium within. However, there is a trade-off that must be made in terms of computation time and numerical errors, and the length of the trajectory T, number of optimizing steps K, step size ξ, and the selected discretization time step Δt. These are important tuning-parameters to investigate when applying this approach. These simulations show that with a long enough horizon the future inputs can be continuously optimized as the system's state evolves even with only a single optimization step each time

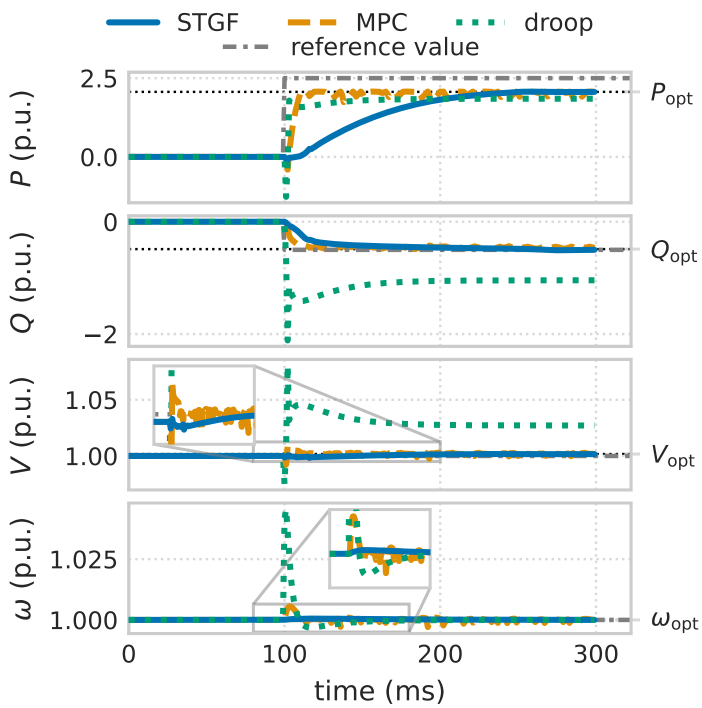

Figure 2.15: The outputs of the system included in the cost function for each control method. With STGF or MPC control the outputs converge to optimal values 'near' the given reference values. The droop control converges to values away from the optimal values.

step of the simulation. When selecting T it is important to ensure that $T > r$, were r is the relative degree of the system dynamics which is the number of times the output values used in $c(\mathbf{x}, \mathbf{u})$ must be differentiated before the input explicitly appears. For MPC one must select similar parameters to STGF relating to the trajectory optimization problem formulation. With the augmented droop controller there are multiple parameters in the synchronization term that must be tuned to ensure the droop controller remains stable, which can be challenging to guarantee for all cases.

At the moment, it is unclear what conditions must be met when selecting T for a given controllable system such that the quadratic program (2.27) is always feasible, however it is likely tied to the relative degree of the system dynamics with regard to the state and input constraints.

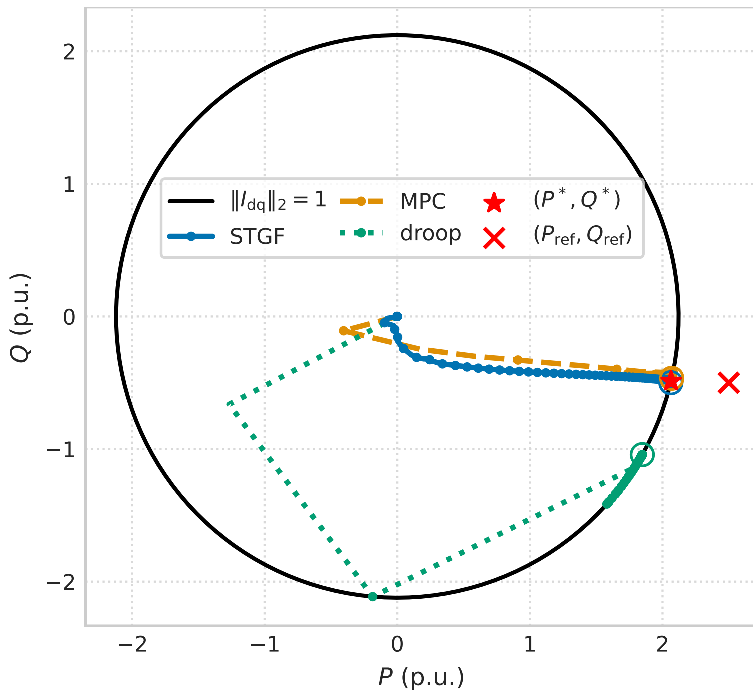

Figure 2.16: The trajectory of I_{dq} with the safe gradient flow trajectory optimization controller and $P^* = 2.5, Q^* = -0.5, V_{\mathrm{nom}} = 1, \omega_{\mathrm{nom}} = 2\pi \cdot 60$ rad s^{-1}. The states are initialized at $\mathbf{0}$. The trajectory is seen to be constrained by the current magnitude limit. Also plotted are the projected trajectory at each timestep.

Controller Computational Burden & Solve-time

Although QPs are considered to be simple optimization problems, using iterative solvers in real-time on inverters themselves is not trivial, because of microprocessor limitations and sampling/input delays [62, 63]. While, this form of QP has a closed-form solution [57], this closed-form solution requires checking which set of inequality constraints are active which becomes burdensome when there are many constraints. For these test we run the controllers and simulations using an Intel(R) Xeon(R) 2.20GHz CPU, which we note has significantly more processor power than a standard microprocessor. Fig. 2.17 shows the optimization solve times for each step of the simulation for the STGF and MPC controllers. The average STGF controller single step runtime is roughly 40 times faster than the MPC controller, demonstrating its potential to generate faster trajectories updates compared to fully solving the nonlinear trajectory optimization problem each time step. Further, the maximum solve time of the STGF QP is orders of magnitude smaller than that of MPC. This difference in solve time is most pronounced when the current magnitude limit is active. The average solve time of STGF was still greater than the simulation step sized used in these tests, indicating that without further computational speed up these approaches could not be used directly to control an inverter system with this setup. In comparison, the explicit droop controller implementation is significantly less computationally intensive.

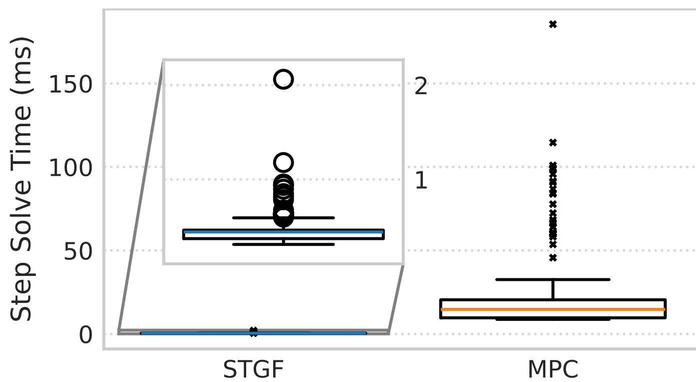

Figure 2.17: Single-step run-times for solving controller optimization problems. An inset enlarged view of STGF run-times shows that they are significantly shorter than most MPC run-times.

Chapter 3

SYSTEM STRENGTH & STABILITY OF POWER SYSTEMS WITH INVERTERS

This chapter introduces the problem of ensuring stability in power systems with high levels of IBRs, through the lens of system strength. Section 3.1 reviews the definition of system strength, modeling for system strength analysis, and existing system strength metrics; highlights their capabilities and limitations, and identifies future directions for developing new tools and metrics to support grid reliability and resilience. Section 3.2 then defines our own system strength metric based on calculations of the sensitivity of the system's eigenvalues to perturbations in power injections and shows how this metric can be used by system operators to find remedial actions in response to stability issues.

The increasing share of inverter-based resources (IBRs) in electrical power systems is causing a range of challenges related to maintaining power system stability. It is computationally challenging and time-intensive to evaluate the stability of power systems with high IBR penetration levels through time-domain simulation or small-signal analysis. Therefore, system strength metrics capable of screening systems for potential stability issues are valuable. The conventional grid strength metric, Short-Circuit Ratio (SCR), has proven useful. However, it is inaccurate when applied to power systems with high IBR penetration. To resolve this issue, a range of new metrics aimed at quantifying system strength have been proposed in the literature.

3.1 Introduction, Background & Modeling[1]

The ability of power systems to maintain stable voltage and frequency under disturbances can broadly be referred to as system strength. The large-scale integration of IBRs with conventional grid-following (GFL) control has led to a decrease in the strength of the grid [23]. Some stability challenges associated with IBRs and weak grids are voltage instability [24] and subsynchronous oscillations [25], which have been observed in power systems repeatedly over the past decade. The effect of weak system strength on wind turbine performance was seen in ERCOT when poorly damped voltage oscillations due to weak grid conditions were observed and eventually tripped the wind power plant [24]. Another event, which was in part due to a system strength shortfall, is a blackout that occurred in 2016 in South Australia [87]. Lastly, Ireland's power system operators face an ongoing grid strength challenge choosing to constrain the penetration of IBR generation to avoid stability concerns related to grid strength shortfalls [88].

Metrics that reflect the strength of the grid can be used to screen the weak system areas and identify potential instability issues. This is particularly important when new IBRs are connected to the grid. Electromagnetic transient (EMT) simulations can be used to assess their impact on grid stability, but the simulations can be computationally and time intensive. Therefore, there is a demand for screening tools or metrics that can be used to identify weak areas of the grid to focus these EMT studies [89]. Accessible and intuitive system strength metrics are essential for the reliable deployment and operation of systems with high levels of IBRs. However, the current literature lacks a clear consensus on which metric best characterizes system strength in such systems.

3.1.1 System Strength

System strength refers to the ability of an AC power system to maintain stable voltage and frequency under all operating conditions [22]. In the specific context of voltage stability, it relates to the sensitivity of bus voltages to perturbations in current injections [90] or

[1]Partially adapted from T. Joswig-Jones, S. Dong, J. Tan, D. Wu and B. Zhang, "A Review of Grid Strength Metrics for Inverter-based Power Systems," IEEE Power & Energy Society General Meeting, 2026, in press. [86]

the ability of the power system to maintain and control the voltage waveform at any given location in the power system, both during steady state operation and following a disturbance [91]. In a strong system the sensitivity of the voltage to system disturbances is low, while in a weak system this sensitivity is higher. Areas identified as having a low system strength are more susceptible to instability [92] and can not support the stable operation of conventional grid-following (GFL) IBRs [93].

Voltage stability in conventional power grids is typically related to slower timescale dynamics, while frequency stability relates to faster timescales, allowing for the decoupling of the stability analysis of these parameters. With increasing IBR penetration, traditional boundaries in timescales used to decouple stability analysis of voltage and frequency are becoming less distinct and both mechanisms are increasingly influenced by faster dynamics [21]. Reflecting this, the term system strength has progressively come to encompass broader ideas of system stability over the years [94]. The work in [95] proposes a new classification of system strength: steady-state system strength and dynamic-state system strength. Specifically, steady-state system strength deals with operation around nominal operating values, while dynamic-state system strength deals with stability of the system when subject to large disturbances. As system strength metrics exist in the literature, only steady-state system strength is considered; however, outside of the typical scope of system strength exist other metrics related to transient stability such as the critical clearing time of faults and system inertia, which can be used in conjunction with typical system strength metrics [96]. In fact, the Grid Strength Assessment Tool [96] – developed by EPRI – calculates the critical clearing time (CCT) of IBR devices for a three-phase fault at its POI to be used in conjunction with other SCR-related metrics to identify critical buses for further study [97]. In [98] the Southwest Power Pool uses CCT as a threshold to determine if further study was required; transient stability analysis was performed for events with a calculated critical clearing times less than 9 cycles. This tool has also been applied in case studies of the Ireland All Island system in [99]. In our work, we focus on existing system strength metrics relating to the stability of the system when operating in steady-state and subject to small disturbances.

When discussing system strength metrics, it is necessary to distinguish between *grid* *strength* and *system strength*, terms that are often used interchangeably in the literature [100, 95]. Grid strength refers to the sensitivity of the network based strictly on its passive characteristics. System strength, conversely, refers to the stability of the entire coupled system, encompassing both the power network and the active control dynamics of the connected IBR devices. For instance, traditional metrics like the Short-Circuit Ratio (SCR) measure grid strength at a specific bus, but a low SCR does not strictly guarantee that a well-tuned IBR connected at that location will be unstable.

In the literature, system strength metrics can generally be categorized into two groups: fault-current based metrics and admittance-model based metrics. Fault current based metrics are defined using short-circuit fault currents. To determine the system strength with these metrics, their value must be compared to its critical value. Industry experience informs approximate ranges for critical values, but their true critical value must be found through EMT simulation. This has led to recent efforts to develop metrics based on admittance models. Admittance-model based metrics use equivalent frequency-domain admittance/impedance models of devices and the power network in their formulation which capture the small-signal characteristics of the system around a steady-state operating point through frequency-domain analysis.

3.1.2 Short-Circuit Analysis Based System Strength

Short-circuit Ratio (SCR) is a metric that has conventionally been used to quantify the strength of the grid focusing on voltage stability [90]. [93]. SCR was originally defined to classify the stability characteristics of synchronous generators and was used as a measure of the electrical stiffness of the machine [101, 102]. The term was later used SCR was originally used with reference to the strength of the AC grid in the context of high-voltage direct-current lines, which can exhibit small-signal instability in weak grids [103]. More recently, SCR, described as (3.1), has been used as a screening tool when connecting IBRs to the system.

$$\mathrm{SCR} \coloneqq \frac{S_{\mathrm{SCMVA}}}{P_{\mathrm{RMW}}}, \tag{3.1}$$

where S_{SCMVA} denotes the short-circuit apparent power at a point of interconnection (POI); P_{RMW} represents the rated power of the connected IBR device.

Short-circuit Ratio (SCR) described in (3.1), is commonly used to quantify the grid strength in terms of voltage stability [90, 103]. Typically, an SCR greater than 3 implies a strong grid, an SCR between 2 and 3 implies a weak grid, and an SCR below 2 indicates a very weak grid [104]. However, a low SCR does not necessarily imply an IBR connected at that POI will be unstable. To evaluate the system strength in terms of its stability, both the power network and device dynamics need to be considered. Thus, the SCR needs to be compared with its critical value, i.e., the critical short-circuit ratio ($\mathrm{SCR}_{\mathrm{crit}}$), which depends on IBR device dynamics, control mode, control parameters, and dispatch. The system is considered stable when $\mathrm{SCR} > \mathrm{SCR}_{\mathrm{crit}}$ and unstable when $\mathrm{SCR} < \mathrm{SCR}_{\mathrm{crit}}$. In addition, the difference between these two values can be referred to as the stability margin.

SCR can not capture interactions between an IBR device and the rest of the system including other IBRs. When multiple IBRs are connected to a power system electrically close to each other, the short-circuit level of the system in that area is shared between the IBRs. This results in the grid strength being significantly less than the SCR calculated for a single one of these IBRs would imply [22]. Further, conventional GFL inverters are prone to instability in weak grids, while grid-forming (GFM) inverters are prone to instability when connecting to strong grids [105]. Thus, applying conventional SCR to systems with high IBR penetrations can lead to inaccurate measures of grid strength. Section 3.1.2 provides an overview of other grid strength metrics proposed in the literature to address this challenge.

System Model for Short-Circuit Analysis

To quantify grid strength using fault-current-based metrics, the power system is traditionally evaluated using fundamental-frequency, positive-sequence short-circuit analysis. Under this paradigm, the complex dynamics of the grid are abstracted into static, phasor-domain equivalent circuits [106].

For the assessment of a single IBR connected to the grid, the power network – as seen from the Point of Interconnection (POI) – is reduced to a Thevenin equivalent circuit [22].

This model consists of an ideal voltage source, representing the open-circuit grid voltage (E_{th}), in series with a Thevenin equivalent impedance (Z_{th}). Under the assumption of a bolted three-phase fault at the POI, the short-circuit apparent power capacity (S_{SCMVA}) provided by the grid is fundamentally defined by this equivalent impedance:

$$S_{\text{SCMVA}} = \frac{|V_{\text{nom}}|^2}{|Z_{th}|}, \tag{3.2}$$

where V_{nom} is the nominal voltage at the POI. In this framework, a "strong" grid is characterized by a minimal Z_{th}, yielding a high fault-current capacity that stabilizes the local voltage during disturbances.

When evaluating multi-infeed systems – where several IBRs are located in close electrical proximity – a single-port Thevenin equivalent is inadequate to capture the interactions between resources. Instead, the passive AC network is modeled using an $N \times N$ nodal impedance matrix ($\boldsymbol{Z}_{bus}$), where N is the number of IBR interconnection buses [49]. The diagonal elements, Z_{ii}, represent the self-impedance (or Thevenin impedance) at bus i, while the off-diagonal elements, Z_{ij}, represent the transfer impedance between bus i and bus j. These transfer impedances mathematically capture the electrical coupling between different nodes in the network. Specifically, they quantify how a voltage perturbation or power injection at one bus influences the voltage profile at another. Consequently, the nodal impedance matrix provides the generalized mathematical foundation necessary to evaluate the collective impact of multiple resources in multi-infeed fault-current metrics.

In practical power system studies, the elements of $\boldsymbol{Z}_{bus}$ and the resulting short-circuit capacities are calculated using standardized fault analysis algorithms [107]. These computational methods construct the network model by representing traditional synchronous generators using their subtransient reactances (X''_d) behind an internal voltage source, allowing for an accurate evaluation of the maximum initial fault current contributions [106]. This algorithmic approach forms the foundation of commercial short-circuit simulation platforms, providing the necessary short-circuit values to compute the fault-current-based grid strength metrics discussed below.

Review of Fault-current Based Metrics

Many SCR variants based on fault-current values have been proposed to more accurately define the grid strength to overcome the shortfalls of SCR, particularly evaluating systems with multiple IBRs. While these methods vary in how they are calculated and what they consider, many of these metrics are based upon fault-infeed values. Short-circuit calculation tools in RMS power system simulators, such as PSS®/E, are often used to calculate these values.

Weighted Short-circuit Ratio The Weighted Short-Circuit Ratio (WSCR) [108] – proposed by the Electric Reliability Council of Texas (ERCOT) – is found by assuming the electrical distance between IBRs can be ignored but interactions between these devices must be considered. It is defined as

$$\mathrm{WSCR} := \frac{\sum_i^{N_{\mathrm{IBR}}} S_{\mathrm{SCMVAi}} \cdot P_{\mathrm{RMW},i}}{\left(\sum_i^{N_{\mathrm{IBR}}} P_{\mathrm{RMW},i}\right)^2}, \tag{3.3}$$

where the subscript i indicates the value relating to the i-th of the N_{IBR} IBRs. This approach provides a conservative measure of SCR for the aggregated group of IBRs.

Composite Short-circuit Ratio The Composite Short-Circuit Ratio (CSCR) [109] – proposed by General Electric (GE) – assumes the IBRs of interest are electrically close and connects them all to a common bus while neglecting their electrical distance. The short-circuit MVA at this common bus excluding current contributions from IBRs, S_{CSCMVA}, is then calculated. CSCR is defined as

$$\mathrm{CSCR} := \frac{S_{\mathrm{CSCMVA}}}{\sum_i^{N_{IBR}} P_{\mathrm{RMW},i}}. \tag{3.4}$$

Like WSCR, CSCR calculates a single value for a group of IBRs.

Equivalent Short-circuit Ratio To account for the electrical distances between IBR devices, the Equivalent Short-Circuit Ratio (ESCR) – also referred to as SCR with interaction factors (SCRIF) [90] – is introduced in [110]. ESCR uses interaction factors defined as the observed change in the voltage at bus i for a given change in voltage at bus j, $IF_{ij} = \Delta V_i / \Delta V_j$, to determine the level of interaction between IBRs. ESCR is defined as

$$\mathrm{ESCR}_i \coloneqq \frac{S_{\mathrm{SCMVA},i}}{P_{\mathrm{RMW},i} + \sum_j (IF_{ji} \cdot P_{\mathrm{RMW},j})}. \tag{3.5}$$

ESCR gives a value at each bus, which differs from WSCR and CSCR that provide a single value for an area.

Site-Dependent Short-circuit Ratio The Site-Dependent Short-circuit Ratio (SDSCR) [111] considers how the IBRs in a multi-infeed system interact through their power injections. In deriving this metric the authors define the equivalent power injection at bus i, using Thevenin equivalent models, as $S_{\mathrm{eq},i} = P_{\mathrm{RMW},i} + \sum_{j \neq i} P_{\mathrm{RMW},j} \cdot w_{ij}$ where the power injections from other IBRs are considered through the interaction factor $w_{ij} = \frac{Z_{\mathrm{RR},ij}}{Z_{\mathrm{RR},ii}} \left(\frac{V_{\mathrm{R},i}}{V_{\mathrm{R},j}} \right)^*$, which depends on the ratio of the electrical distances and the ratio of the bus voltages of the i-th and j-th IBR resource. Using the static voltage stability limit and combining it with the change in voltage due to the equivalent power injection, SDSCR is then defined as

$$\mathrm{SDSCR}_i \coloneqq \frac{|V_{\mathrm{R},i}|^2}{|S_{\mathrm{eq},i}||Z_{\mathrm{RR},ii}|}. \tag{3.6}$$

SDSCR measures the system strength in terms of the distance to the static voltage stability limit of each bus; when SDSCR_i is smaller than 1, there is a heightened risk of voltage stability. [111] provides a case study of a 39-bus system comparing SCR, WSCR, CSCR, and SDSCR showing that SDSCR can accurately characterize static voltage stability of the network when the other methods cannot.

Inverter Penetration Short-circuit Ratio The Inverter Penetration Short-circuit Ratio (IPSCR) [112] – proposed by Electranix – is calculated solely using fault infeed data from short-circuit analysis and does not require the calculation of interaction factors. IPSCR is defined at each bus as the ratio of the short-circuit apparent power provided by conventional generators to that provided by IBRs:

$$\mathrm{IPSCR}_i \coloneqq \frac{S_{\mathrm{SCMVA,SG},i}}{S_{\mathrm{SCMVA,IBR},i}}, \tag{3.7}$$

where $S_{\text{SCMVA,SG},i}$ and $S_{\text{SCMVA,IBR},i}$ are the short-circuit powers evaluated on the bus i when all the IBRs turned off and when conventional generators turned off, respectively. In the IBR only case, each IBR's source impedance is adjusted until it contributes 1 p.u. of S_{SCMVA} at its POI based on its rated capacity. This approach naturally groups IBRs on the basis of their electrical distances and considers their power injection interactions.

Normalized Short-circuit Ratio Taking into consideration the impact that the power dispatch, control parameters, and control mode can have on the stability margin of an IBR, the Normalized Short-circuit ratio (NSCR) [113] is defined based on the stability margin found from comparing SCR to the critical SCR value of an IBR. NSCR is defined at bus i as

$$\text{NSCR}_i \coloneqq \frac{\text{SCR}_i - \text{SCR}_{\text{crit},i}^{\text{min}}}{\text{SCR}_{\text{crit},i}^{\text{max}} - \text{SCR}_{\text{crit},i}^{\text{min}}} \tag{3.8}$$

where $\text{SCR}_{\text{crit},i}^{\text{min}}$ and $\text{SCR}_{\text{crit},i}^{\text{max}}$ are the minimum and maximum critical SCR values for the IBR connected at bus i as its active power dispatch varies from 0.0 p.u. to 1.0 p.u. These critical SCR values are obtained through EMT simulations or small-signal analysis based on a single-IBR infinite-bus system.

3.1.3 Admittance Modeling Based System Strength

While fault-current-based metrics provide a useful heuristic for evaluating grid strength, they are inherently limited by their static formulation and inability to capture the high-frequency control interactions characteristic of IBRs. To accurately evaluate the stability of a power electronics-dominated system, the grid must be analyzed using dynamic frequency-domain transfer functions. This approach, first used for impedance-based stability criteria [114], models the small-signal dynamic characteristics of the coupled system around a specific steady-state operating trajectory.

Equivalent Admittance Modeling Formulation

To construct these frequency-domain models, the physical and control dynamics of the active devices (IBRs and synchronous machines) are first represented using detailed, nonlinear state-space equations in a local dq reference frame:

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{V}_{\mathrm{dq}}) \tag{3.9}$$

$$\mathbf{I}_{\mathrm{dq}} = g(\mathbf{x}, \mathbf{V}_{\mathrm{dq}}), \tag{3.10}$$

where $\mathbf{x}$ represents the internal state variables of the device (e.g., controller integrator states, phase-locked loop angles, and filter currents), while the inputs and outputs are the point of interconnection (POI) voltage ($\mathbf{V}_{\mathrm{dq}}$) and injected current ($\mathbf{I}_{\mathrm{dq}}$), respectively.

Because small-signal stability is strictly dependent on the system's current state, these nonlinear models are linearized around an equilibrium point $(\mathbf{x}_0, \mathbf{V}_{\mathrm{dq0}}, \mathbf{I}_{\mathrm{dq0}})$ established by the AC optimal power flow (ACOPF) equations. This yields the linearized state-space model:

$$\Delta\dot{\mathbf{x}} = \mathbf{A}\Delta\mathbf{x} + \mathbf{B}\Delta\mathbf{V}_{\mathrm{dq}} \tag{3.11}$$

$$\Delta\mathbf{I}_{\mathrm{dq}} = \mathbf{C}\Delta\mathbf{x} + \mathbf{D}\Delta\mathbf{V}_{\mathrm{dq}}, \tag{3.12}$$

where $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ are the respective Jacobian matrices evaluated at the operating point.

Applying the Laplace transform maps these time-domain state-space equations into the frequency domain, yielding the equivalent admittance model of the device:

$$\mathbf{Y}_{\mathrm{dq}}(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B} + \mathbf{D}. \tag{3.13}$$

For a given device, this formulation acts as a 2×2 multi-input multi-output (MIMO) transfer function matrix, explicitly relating voltage perturbations to current perturbations at the POI.

Whole-System Admittance Model To evaluate the system globally, the individual device models are aligned to a common, global dq-reference frame. Once aligned, the passive network and active devices can be combined to form a whole-system admittance model—a comprehensive frequency-domain transfer function matrix that characterizes the coupled dynamic behavior of the entire grid [115].

The whole-system admittance matrix, $\hat{\mathbf{Y}}(s)$, is defined analytically as:

$$\hat{\mathbf{Y}}(s) = \left(\mathbf{I} + \mathbf{Y}_{\text{net}}(s)\mathbf{Z}_{\text{dev}}(s)\right)^{-1} \mathbf{Y}_{\text{net}}(s), \tag{3.14}$$

where $\mathbf{Y}_{\text{net}}(s)$ is the nodal admittance matrix of the passive AC network, and $\mathbf{Z}_{\text{dev}}(s) = \text{diag}(Z_1(s), Z_2(s), \ldots, Z_K(s))$ is a block-diagonal matrix containing the equivalent impedances of all K active devices.

The poles of each element within $\hat{\mathbf{Y}}(s)$ are mathematically equivalent to the small-signal eigenvalues (λ) of the entire system [115]. If all system eigenvalues reside in the open left-half of the complex plane, the system is small-signal stable. Further, as discussed in [116], the sensitivity of these eigenvalues is related to $\hat{\mathbf{Y}}(s)$ as $\frac{\partial\lambda}{\partial\hat{\mathbf{Y}}^{-1}(\lambda)} = -\text{Res}_\lambda\hat{\mathbf{Y}}$ and the perturbation of these eigenvalues with respect to sufficiently small perturbations in $\hat{\mathbf{Y}}(s)$ is

$$\Delta\lambda = \langle -\text{Res}^*_\lambda\hat{\mathbf{Y}}, \Delta\hat{\mathbf{Y}}^{-1}\rangle,$$

where $\langle\cdot,\cdot\rangle$ is the Frobenius inner product of two matrices, $\text{Res}_\lambda\hat{\mathbf{Y}}$ is the residue of the whole-system admittance matrix evaluated at λ, and $*$ denotes the conjugate transpose [117].

A critical limitation of admittance-based modeling is its strict dependency on the linearization point. As the power injected by a device or its POI voltage magnitude fluctuates, the point at which the connected device is linearized shifts. Consequently, the device's equivalent impedance matrix, $\mathbf{Z}_{\text{dev}}(s)$, changes, subsequently altering the system eigenvalues. Because of this, traditional admittance metrics must be repeatedly recalculated to verify stability across varying dispatch scenarios.

Review of Admittance-Model-Based Metrics

The following metrics use equivalent admittance models in their formulation. These admittance models can capture the small-signal characteristics around a steady-state operating point. With this approach the stability of the whole system can be determined using methods applicable to MIMO transfer functions, such as the generalized Nyquist criterion [118] or disk margin [119].

Generalized SCR A natural extension of SCR to a multi-infeed system that reflects the interaction of IBRs led to the definition of generalized SCR (gSCR) in [120]. The approach is based on dividing the IBR system into decoupled subsystems using eigenvalue decomposition and then focusing on the weakest of these subsystems. Later studies have linked gSCR to small-signal stability [121], addressed heterogeneous IBR control [122], and considered GFM IBRs [123]. The gSCR value can be calculated from the steady-state system values as

$$\text{gSCR} \coloneqq \min_{i=1,..,n_{\text{GFL}}} \lambda_i(\boldsymbol{Y}_{eq}) \tag{3.15}$$

where $\boldsymbol{Y}_{eq} = \boldsymbol{S}_b^{-1}\boldsymbol{B}$ with $\boldsymbol{S}_b$ being a diagonal matrix containing the rated capacity of each GFL and $\boldsymbol{B}$ representing the reduced node susceptance matrix after eliminating passive and the infinite bus. gSCR provides a single value that represents the grid strength of a multi-infeed system; the grid strength of each POI can be further quantified based on the participation factor of the j-th device to the equivalent weakest equivalent subsystem, i.e., $p_{1j} = \boldsymbol{W}_{1j}\boldsymbol{W}_{j1}^{-1}$, where $\boldsymbol{W}$ contains the eigenvectors of $\boldsymbol{Y}_{\text{eq}}$. Calculating gSCR only requires phasor-domain parameters and GFM voltage support can be considered by embedding IBR device dynamics into network parameters [123]. The critical gSCR value can be obtained by either EMT simulations based on the weakest equivalent single-IBR subsystem or small-signal analysis by solving the expression

$$\det\left(\text{CgSCR}\cdot\boldsymbol{I} - \sum_{i=1}^{n} p_{1i}\boldsymbol{F}^{-1}\boldsymbol{Y}_{\text{GFL},i}(s_c) + \boldsymbol{\delta}_{\text{GFM}}(s_c)\right) = 0$$
$$s_c = j\omega_c,$$

the determinant of the return difference matrix of the weakest subsystem, for CgSCR. This expressions requires measurements of $\boldsymbol{Y}_{\text{GFL},i}(s)$ and $\boldsymbol{Y}_{\text{GFM},i}$, the admittance transfer function matrices of the grid-following inverters and grid-forming inverters. $\boldsymbol{\delta}_{\text{GFM}}(s)$ represents the shunt admittance models representing the angle dynamics of the grid-forming inverters, and $\boldsymbol{F}$ represents the admittance of a per-unit inductive line. This requires grey-box modeling of the GFMs as knowledge of the GFMs voltage and current control dynamics is required to calculate $\boldsymbol{\delta}_{\text{GFM}}(s)$. Another option is to ignore the perturbation term from the GFM devices and solve for CgSCR from

$$\det\left(\boldsymbol{Y}_{\text{GFL}}(s) + gSCR\cdot\boldsymbol{Y}_{\text{net}}(s)\right) = 0$$
$$s = j\omega.$$

gSCR has been proposed as being useful for solving a range of problems related to integrating IBRs [124, 125, 126, 127, 128]. In [124] and [125] they use gSCR to determine the optimal placement of GFM IBRs to increase grid strength and constrain grid strength to a desired stability margin, respectively. In [126] gSCR is applied to formulate an optimization constraint for assessing the maximum capacity of GFL IBRs in a system. In [127] they use gSCR to develop a grid strength constrained cooperative control approach for IBRs. A recent extension of gSCR is proposed in [128] – this index is called the extended generalized SCR (EgSCR) – to consider the influence of local loads on the strength of IBR systems.

For evaluating the stability of GFM IBRs, grid current strength (GCS), the dual concept of grid voltage strength, is introduced in [129], where the grid current strength is defined as

$$\mathrm{GCS} = \frac{S_{\mathrm{B}}}{S_{\mathrm{SCMVA}}}$$

with S_{B} representing the rated capacity of the IBR. Similar to gSCR, the generalized GCS (gGCS) metric is defined relating to the weakest decoupled GFM IBR subsystem. [129] proposes using both gSCR and gGCS in tandem to evaluate the stability of hybrid GFM-GFL system.

Grid Strength Impedance Metric The Grid Strength Impedance Metric (GSIM) defines a metric using a ratio including the eigenloci of the system admittance, $\boldsymbol{Y}_{\mathrm{sys}}(s)$, and the base impedance of the system, $\boldsymbol{Z}_{\mathrm{b}}(s)$ [100]. These eigenloci of the system admittance and base impedance are defined as

$$\lambda(\boldsymbol{Y}_{sys}(s)) = \begin{bmatrix} |\lambda_{\boldsymbol{Y}_{\mathrm{sys,q}}(s)}| \\ |\lambda_{\boldsymbol{Y}_{\mathrm{sys,d}}(s)}| \end{bmatrix}, \quad \lambda(\boldsymbol{Z}_{\mathrm{b}}(s)) = \begin{bmatrix} |\lambda_{\boldsymbol{Z}_{\mathrm{b,q}}(s)}| \\ |\lambda_{\boldsymbol{Y}_{\mathrm{b,d}}(s)}| \end{bmatrix},$$

respectively. The dq-axis GSIM values are calculated as

$$\begin{bmatrix} \mathrm{GSIM_q}(s) \\ \mathrm{GSIM_d}(s) \end{bmatrix} = \vec{\lambda}(\boldsymbol{Y}_{\mathrm{sys}}(s)) \odot \vec{\lambda}(\boldsymbol{Z}_{\mathrm{b}}(s)) \tag{3.16}$$

where $\vec{\lambda}(\boldsymbol{Y}_{\mathrm{sys}}(s))$ and $\vec{\lambda}(\boldsymbol{Z}_{\mathrm{b}}(s))$ are the eigenloci of the system admittance and base impedance, respectively, and $\odot$ represent the element-wise multiplication of two vectors. GSIM is then defined by combining values from (3.16) as

$$\text{GSIM}(s) \coloneqq \sqrt{\frac{\text{GSIM}_\text{q}(s)^2 + \text{GSIM}_\text{d}(s)^2}{2}}. \tag{3.17}$$

This metric is defined across the frequency spectrum and can vary significantly when evaluated at different frequencies. Interpreting GSIM can be challenging as it has no direct relationship to the stability margin of a system. However, a larger GSIM generally indicates a greater system strength. GSIM is applied in [130] for a case study of a 12-bus model to help identify stability challenges related to GFL IBRs.

Impedance Margin Ratio The Impedance Margin Ratio (IMR) [131] uses the fact that a change in the equivalent impedance of a device connected at a bus will result in a change in the modes of the system, based on the sensitivity of these modes to changes in device impedances. A low IMR identifies buses where perturbations in the impedance at that bus is associated with a large sensitivity in one of the system's modes. This modal sensitivity is shown in [116] to be $\partial\lambda/(\partial \boldsymbol{Z}_{Ak}(\lambda)) = -\text{Res}_\lambda \boldsymbol{Y}_{kk}^{\text{sys}}$, where $\text{Res}_\lambda \boldsymbol{Y}_{kk}^{\text{sys}}$ is the residue of the whole-system admittance matrix [115] measured at the POI of bus k, evaluated at the mode of interest, $\lambda = \sigma + j\omega$. This sensitivity is used to determine an impedance margin, which is defined as the maximum allowable change in impedance of a device such that the dominant mode will remain stable, as $\|\Delta \boldsymbol{Z}_{Ak}(\lambda)\|_\text{max} = |\sigma|/\|\text{Res}_\lambda^* \boldsymbol{Y}_{kk}^{\text{sys}}\|$, where $\|\cdot\|$ indicates the Frobenius norm of a matrix. IMR is defined as the magnitude of this impedance margin divided by the magnitude of the original impedance of the device,

$$\text{IMR} \coloneqq \frac{\|\Delta \boldsymbol{Z}_{Ak}(\lambda)\|_\text{max}}{\|\boldsymbol{Z}_{Ak}(\lambda)\|}. \tag{3.18}$$

To calculate IMR the modes of the system must first be identified and the strength of a bus is determined by the minimum IMR over the range of modes observed at that bus. It should be noted that the approximate impedance margin used in IMR is 'conservative' as it assumes the worst-case direction for the perturbation in the device impedance. Further, this margin is only accurate when small perturbations in the device impedance would result in the mode becoming unstable, as it is using a first-order sensitivity which does not capture the potentially nonlinear behavior of the paths of the modes for larger changes in device impedance. Lastly, the Admittance Margin (AM) is defined in [132] as a complement to IMR used to define grid strength at buses without any devices connected yet.

Dynamic Impedance Margin The authors of [133] proposed the Dynamic Impedance Margin (DIM) to address a gap between steady-state analysis tools and dynamic models. DIM provides insights into the dynamic voltage stability of systems by modeling generation resources using their *dynamic impedance* values. These values are extracted from the sensitivity of the current injected by a resource to changes in the voltage at that POI and is defined as

$$Z_{dyn} = \frac{1}{Y_{adj}}, \quad Y_{adj} = \frac{dI_{\mathrm{q}}}{dV_{\mathrm{mag}}} \cos\left(\phi_v - \phi_{I_{\mathrm{q}}} - \pi\right),$$

where I_{q} is the reactive current injected and $\phi_v, \phi_{I_{\mathrm{q}}}$ are the phase angles of the voltage and reactive current, respectively. The dynamic impedance varies, particularly for GFL IBRs, over the range of frequencies, so to evaluate the dynamic impedance requires selecting a perturbation frequency. Each device is then represented with a Thevenin equivalent model – an ideal voltage source behind its dynamic impedance value – in a steady-state power flow model and the system's power transfer limits are tested using P-V analysis. The authors show that the power transfer stability limit predicted using DIM is typically conservative compared to results found from EMT studies, where they evaluate stability through a list of criteria related to deviations in the active powers and voltages in the system. DIM extracts information from admittance models; however, it does not directly consider the small-signal stability of the system.

3.1.4 Comparison of Existing Metrics

In this section, we compare the reviewed grid strength metrics in terms of their characteristics, advantages, and limitations. Table 3.1 summarizes the characteristics of each metric.

The first distinction between the metrics presented aligns with our categorization. Fault-current based metrics cannot consider IBR control modes or GFM voltage support, except when compared to critical values obtained from EMT simulation. In contrast, admittance-model based metrics inherently consider device dynamics. This benefit of using admittance-model based metrics comes at the cost of acquiring admittance models, compared to fault-

Table 3.1

COMPARISON OF GRID STRENGTH METRICS. '-' IMPLIES A VALUE IS UNSPECIFIED OR UNDEFINED. (YES ✓, NO ✗)

Category	System Strength Evaluation					Applicable to IBR Dominated Systems		Considers Varying Operating Point	Metric Scope	Ease of Industry Adoption
	Metric		Critical Value							
	Name	Data Required for Calculation	Voltage Stability	Small-Signal Stability Approx.[2]	Small-Signal Stability Accurate[3]	Considers IBR Control	Considers IBR Interactions			
Fault-current based	SCR [90]	Phasor-domain Parameters	1	1-3	EMT	✗	✗	✗	Bus	★★★
	WSCR [108]	Phasor-domain Parameters	-	1-3	EMT	✗	✓	✗	System	★★★
	CSCR [109]	Phasor-domain Parameters	-	1-3	EMT	✗	✓	✗	System	★★★
	ESCR [110]	Phasor-domain Parameters	-	1-3	EMT	✗	✓	✗	Bus	★★★
	SDSCR [111]	Phasor-domain Parameters	1	-	-	✗	✓	✗	Bus	★★★
	IPSCR [112]	Phasor-domain Parameters	-	-	EMT	✗	✓	✗	Bus	★★★
	NSCR [113]	Phasor-domain Parameters	-	0-1	EMT	✓	✗	✓	Bus	★★☆
Admittance-model based	gSCR [123]	Phasor-domain Parameters	1	-	FDA or EMT[4]	✓	✓	✗	System	★★☆
	GSIM [100]	Device Admittance Models	-	-	-	✓	✓	✗	System	★☆☆
	IMR [131]	Device Admittance Models	-	-	0[5], FDA	✓	✓	✗	Bus	★☆☆
	DIM [133]	Device Admittance Models	-	-	EMT	✓	✗[6]	✗	Bus	★★☆

[2]Indicates approximate critical values of the metric based on industry experience or its formulation.

[3]Indicates whether the true critical value of a metric is determined using electromagnetic transient (EMT) simulation or frequency-domain analysis (FDA).

[4]The critical gSCR value of a system can be obtained by simulating the weakest equivalent single-inverter infinite-bus subsystem instead of the full system model.

[5]IMR does not define a critical value, but the IMR is zero when an eigenvalue's real part is zero.

[6]Interactions through the PV analysis are considered; however, IBR control interactions cannot be modeled with such an approach.

current based metrics which only require more readily available phasor-domain system parameters used in short-circuit analysis. Admittance models can be extracted from analytical models of device dynamics, which are not typically available from vendors, or can be obtained via frequency scanning techniques in black box scenarios [134]. Implementing frequency scanning in real time may need additional hardware, such as phasor measurement units, which complicates practical implementation. In this regard, gSCR is different from the other admittance-model based metrics as it is derived based on admittance models but its calculation only requires phasor-domain system parameters.

While all admittance-model based metrics consider the interactions between devices via admittance-based models, fault-current based metrics differ significantly in their consideration of IBR interactions. SCR ignores other IBRs, while WSCR and CSCR assume IBRs are connected at the same bus, thus ignoring the IBR interaction via the network. The fault-current based metrics can consider inverter interactions to some extent, but they cannot fully consider control interactions between IBRs that can lead to small-signal stability issues, such as subsynchronous oscillation. Admittance-model based approaches inherently reflect these control interactions.

To assess system stability, most metrics are interpreted relative to a critical value, which can be defined empirically (from industry experience), theoretically (based on voltage stability or small-signal stability theory), or via simulation (from EMT results). Most fault-current based metrics have approximate ranges indicating when a system is susceptible to instability based on industry experience. To accurately determine the critical value of these metrics, one must run EMT simulations of the full system being considered. Admittance-model based methods can instead use frequency-domain analysis to determine accurate critical values without requiring EMT simulation of the full system.

Within the admittance-model based metrics the robustness of the system's stability can be represented by not only a stability margin, but also through the sensitivity of the stability. Of the metrics reviewed, only IMR reflects both the stability margin and sensitivity. gSCR only indicates the stability margin when comparing it to the critical gSCR value and GSIM is not directly related to the stability of the system. DIM provides a stability margin in the form of power transfer limits, but it is unclear precisely how this relates to the overall stability of the system. While IMR reflects the sensitivity of the eigenvalues to perturbations in device impedances, it should be noted that it does so using first-order approximations of $\partial\lambda/\partial Z_{Ak}$, which does not capture the potentially nonlinear behavior of the paths of the modes for larger changes in device impedance, and the impedance margin calculation is 'conservative' as it assumes the worst-case direction for the perturbation in the device impedance.

Admittance-model and fault-current metrics depend on the system operating point and must be calculated over a range of operating points to understand how system strength varies with these values. Among the metrics considered, only NSCR explicitly incorporates the effect of varying device power set points in its definition. We discuss this challenge and propose our own system strength metric to address it in Section 3.2.

Lastly, in Table 3.1, we compare the scope and the ease of industry adoption of each metric. A metric's scope determines if it calculates a value for each bus in the system or a single value for the whole system being considered. Its ease of industry adoption gives an idea of how easily the metric can practically be implemented in an industry setting and is qualitatively rated out of three stars based on the modeling effort, required data availability, and ease of integration into existing short-circuit analysis programs.

3.1.5 Gaps and Future Directions for Metric Development

Current gaps and challenges related to grid strength metrics and possible future directions for research are as follows:

1. Fault-current metrics have not been developed to consider voltage support from GFM IBRs. The development of fault-current metrics that better consider GFM IBR capabilities would benefit IBR connection studies.

2. Fault-current and admittance-based metrics do not all provide clear critical values relating to system stability or practical approaches for calculating these critical values. Researchers should ensure metrics are developed such that they have clear definitions of critical values that can help inform system operator decisions.

3. Fault-current and Admittance-based metrics are dependent on the operating point of the system and must be recalculated as this value differs. Metrics that can reflect the grid strength and its sensitivity over a range of operating conditions could assist operators in identifying potential stability challenges in the system.

4. Admittance-model based metrics lack industrial documentation and guidance on their implementation. Collaboration between industry and academia to apply these metrics to real-world systems and the development of tools to automate or simplify the process of acquiring admittance models and implementing admittance-model metrics could accelerate adoption.

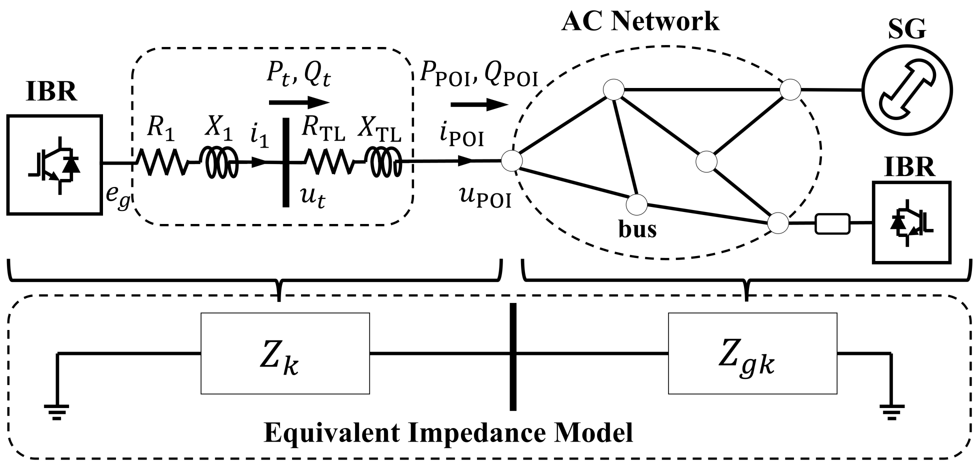

Figure 3.1: Illustration of a power system with its equivalent impedance model.

3.2 Sensitivity-Based System Strength Assessment[7]

3.2.1 System Model

We consider a system of devices connected in an AC network. A representation of such a network can be seen in Fig. 3.1. The network is defined with a set of N buses $\mathcal{N}$ and lines $\mathcal{E}$, where $Y_{ik} = G_{ik} + jB_{ik}$ denotes the admittance between buses i and k. Bus $i \in \mathcal{N}$ is characterized by power injections $S_i = P_i + jQ_i$ and voltage $V_i \angle \theta_i$; these are collected into vectors $\mathbf{V}, \mathbf{\Theta}, \mathbf{P}, \mathbf{Q}$. A slack bus maintains system power balance by varying its power injections.

Power Flow

In steady-state, the power flow through the network satisfies the AC power flow equations

$$P_i = V_i \sum_{i \sim k} V_k (G_{ik} \cos\theta_{ik} + B_{ik} \sin\theta_{ik}), \tag{3.19a}$$

$$Q_i = V_i \sum_{i \sim k} V_k (G_{ik} \sin\theta_{ik} - B_{ik} \cos\theta_{ik}), \tag{3.19b}$$

[7]Adapted from T. Joswig-Jones, B. Zhang, S. Dong and J. Tan, "System Strength Sensitivity to Power Flow Perturbations in AC Power Systems with Inverter-based Resources," IEEE Power & Energy Society General Meeting, 2026, in press. [135]

for each bus $i \in \mathcal{N}$. The first-order Taylor series expansion of (3.19) gives

$$\begin{bmatrix} \Delta\mathbf{P} \\ \Delta\mathbf{Q} \end{bmatrix} = J(\mathcal{M}) \begin{bmatrix} \Delta\mathbf{V} \\ \Delta\boldsymbol{\Theta} \end{bmatrix}, \quad J(\mathcal{M}) = \begin{bmatrix} \frac{\partial\mathbf{P}}{\partial\mathbf{V}} & \frac{\partial\mathbf{P}}{\partial\boldsymbol{\Theta}} \\ \frac{\partial\mathbf{Q}}{\partial\mathbf{V}} & \frac{\partial\mathbf{Q}}{\partial\boldsymbol{\Theta}} \end{bmatrix}, \tag{3.20}$$

where J is the Jacobian typically used to solve the power flow problem. Here we denote the power flow parameters at all buses in the system as $\mathcal{M}$, where $\mathcal{M}_i = [P_i, Q_i, V_i, \Theta_i]$ are the power flow parameters at bus i.

To solve power flow models, buses are specified as having two of their four power flow parameters being specified. Typically, bus types are PQ, PV, and $V\Theta$ (the slack bus) where the bus type indicates the parameters that are specified at that bus. We let $\overline{\mathcal{V}}$, $\overline{\Theta}$, $\overline{\mathcal{P}}$, and $\overline{\mathcal{Q}}$ denote the set of buses where the respective parameter is specified at that bus and we let $\widetilde{\mathcal{V}}$, $\widetilde{\Theta}$, $\widetilde{\mathcal{P}}$, and $\widetilde{\mathcal{Q}}$ denote the set of buses where the respective parameter value at that bus is dependent upon the specified values and (3.19).

Dynamics and Small-Signal Stability

As established in Section 3.1.3, the small-signal stability of the coupled power system is dictated by the poles of the whole-system admittance matrix, $\hat{\mathbf{Y}}(s)$. However, directly analyzing $\hat{\mathbf{Y}}(s)$ to determine grid strength can be mathematically cumbersome because the passive network admittance ($\mathbf{Y}_{\text{net}}$) and the active device impedance ($\mathbf{Z}_{\text{dev}}$) are deeply coupled through a matrix inverse, as seen in (3.14).

To evaluate the sensitivity of the system with respect to physical network parameters, it is highly advantageous to analyze the inverse whole-system admittance matrix. By inverting (3.14), the system components can be algebraically separated into a simple summation:

$$\hat{\mathbf{Y}}^{-1}(s) = \mathbf{Y}_{\text{net}}^{-1}(s) + \mathbf{Z}_{\text{dev}}(s). \tag{3.21}$$

This formulation fundamentally isolates the dynamics of the active IBR controllers ($\mathbf{Z}_{\text{dev}}$) from the passive topological characteristics of the AC grid ($\mathbf{Y}_{\text{net}}^{-1}$).

To quantify system strength, it is necessary to track how the system eigenvalues (λ) migrate in response to changes in this system matrix. As demonstrated by the authors in [116], the sensitivity of a system mode to perturbations in the inverse whole-system admittance matrix is analytically defined by its residue:

$$\frac{\partial \lambda}{\partial \hat{\mathbf{Y}}^{-1}(\lambda)} = -\mathrm{Res}_{\lambda} \hat{\mathbf{Y}}. \tag{3.22}$$

Consequently, the perturbation of a specific eigenvalue ($\Delta\lambda$) with respect to sufficiently small perturbations in the system matrix ($\Delta\hat{\mathbf{Y}}^{-1}$) can be evaluated using the Frobenius inner product:

$$\Delta\lambda = \langle -\mathrm{Res}^{*}_{\lambda} \hat{\mathbf{Y}}, \Delta\hat{\mathbf{Y}}^{-1} \rangle, \tag{3.23}$$

where $\langle \cdot, \cdot \rangle$ denotes the Frobenius inner product, $\mathrm{Res}_{\lambda}\hat{\mathbf{Y}}$ is the residue of the whole-system admittance matrix evaluated at the eigenvalue λ, and $*$ denotes the conjugate transpose [117].

While this analytical framework has been successfully applied to derive admittance-based metrics evaluating device impedance margins, such as the Impedance Margin Ratio (IMR) [131], it has not been fully leveraged to quantify grid strength regarding the physical network topology itself. Because admittance models vary with the system operating point, tracking stability as the physical network changes (e.g., during line outages or topological reconfigurations) traditionally requires recalculating the entire system state.

To address this limitation and provide an intuitive measure of structural grid strength, a new framework must be established that connects these eigenvalue sensitivities directly to the physical power flow and line admittance parameters of the network.

3.2.2 Method

The proposed methodology uses analytical sensitivities to evaluate how system eigenvalues respond to power flow parameter changes. As demonstrated in Fig. 3.2, while first-order approximations do not directly capture highly nonlinear trajectories under large parameter variations, they successfully indicate the initial direction and magnitude of modal perturbations. To formulate these metrics, this section first evaluates how parameter changes shift the operating point, and subsequently maps those shifts to the system's global eigenvalues. The resulting sensitivities are used to define system strength metrics.

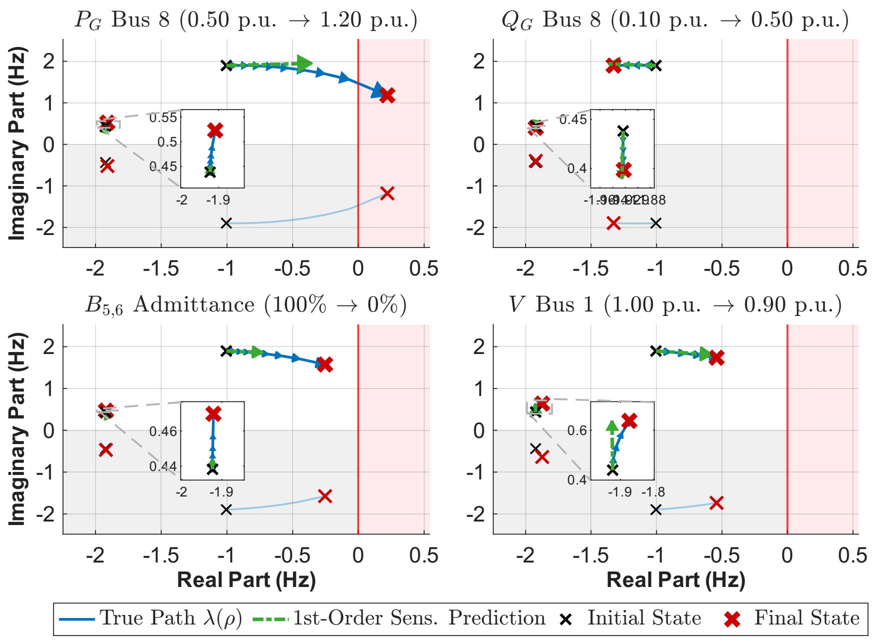

Figure 3.2: 14-bus case select eigenvalue traces with first-order sensitivity perturbation prediction for changes in power flow and network admittance parameters.

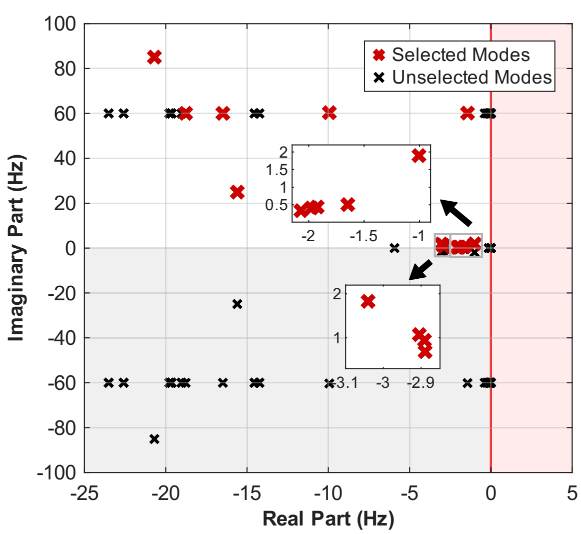

Figure 3.3: 14-bus case eigenvalues with considered eigenvalues highlighted. Considered eigenvalues are from the set $\{\lambda = \sigma + j\omega | \lambda \in \text{poles}(\hat{\mathbf{Y}}(s)), |\lambda| > 0.016 \text{ Hz}, -40 \text{ Hz} < \sigma, \omega \neq 0 \text{ Hz}, |\omega - \omega_{\text{nom}}| > 0.016 \text{ Hz}, -160 \text{ Hz} < \omega < 160 \text{ Hz}\}$.

Power Flow Sensitivity

Sensitivity to P, Q, V: When the power injection or voltage magnitude at a bus changes, the system adjusts to a new power flow solution satisfying (3.19). Accurately assessing the impact of such a perturbation on small-signal stability requires accounting for its effect on the overall power flow as this determines the operating point the devices in the system are linearized about [136]. Although this nonlinear problem is generally challenging, for sufficiently small perturbations it can be approximated using a first-order expansion. In this section we discuss how to calculate the power flow perturbation $\Delta\mathcal{M}$ for a perturbation to a chosen specified power flow parameter Δx (e.g. for a perturbation in the active power at bus 2 set $x = P_2$).

First, we partition the Jacobian relating to whether or not a power flow parameter is specified or dependent and define sub-matrices as shown in (3.24).

$$J = \left[\begin{array}{cc|cc} \frac{\partial\overline{\mathbf{P}}}{\partial\overline{\overline{\mathbf{V}}}} & \frac{\partial\overline{\mathbf{P}}}{\partial\overline{\overline{\boldsymbol{\Theta}}}} & \frac{\partial\overline{\mathbf{P}}}{\partial\widetilde{\widetilde{\mathbf{V}}}} & \frac{\partial\overline{\mathbf{P}}}{\partial\widetilde{\widetilde{\boldsymbol{\Theta}}}} \\ \frac{\partial\overline{\mathbf{Q}}}{\partial\overline{\overline{\mathbf{V}}}} & \frac{\partial\overline{\mathbf{Q}}}{\partial\overline{\overline{\boldsymbol{\Theta}}}} & \frac{\partial\overline{\mathbf{Q}}}{\partial\widetilde{\widetilde{\mathbf{V}}}} & \frac{\partial\overline{\mathbf{Q}}}{\partial\widetilde{\widetilde{\boldsymbol{\Theta}}}} \\ \hline \frac{\partial\widetilde{\mathbf{P}}}{\partial\overline{\overline{\mathbf{V}}}} & \frac{\partial\widetilde{\mathbf{P}}}{\partial\overline{\overline{\boldsymbol{\Theta}}}} & \frac{\partial\widetilde{\mathbf{P}}}{\partial\widetilde{\widetilde{\mathbf{V}}}} & \frac{\partial\widetilde{\mathbf{P}}}{\partial\widetilde{\widetilde{\boldsymbol{\Theta}}}} \\ \frac{\partial\widetilde{\mathbf{Q}}}{\partial\overline{\overline{\mathbf{V}}}} & \frac{\partial\widetilde{\mathbf{Q}}}{\partial\overline{\overline{\boldsymbol{\Theta}}}} & \frac{\partial\widetilde{\mathbf{Q}}}{\partial\widetilde{\widetilde{\mathbf{V}}}} & \frac{\partial\widetilde{\mathbf{Q}}}{\partial\widetilde{\widetilde{\boldsymbol{\Theta}}}} \end{array}\right] = \left[\begin{array}{c|c} J_{-,-} & J_{-,\sim} \\ \hline J_{\sim,-} & J_{\sim,\sim} \end{array}\right] \tag{3.24}$$

Here $J_{-,\sim}$ represents the mapping of dependent voltage parameters, $(\widetilde{\mathcal{V}}, \widetilde{\Theta})$, to specified power parameters, $(\overline{\mathcal{P}}, \overline{\mathcal{Q}})$. We further define $J_{-,\bullet} = [J_{-,-}, J_{-,\sim}]$, which selects the rows of J corresponding to specified power injections and the columns corresponding to all voltage values.

For a perturbation in the active power at buses in $\overline{\mathcal{P}}$ or reactive power at buses in $\overline{\mathcal{Q}}$ ($x = P_i, i \in \overline{\mathcal{P}}$ or $x = Q_i, i \in \overline{\mathcal{Q}}$) we can calculate the voltage magnitude and angle sensitivities at buses in $\widetilde{\mathcal{V}}$ and $\widetilde{\Theta}$, respectively, from

$$\left(\frac{\partial\mathbf{P}}{\partial x}, \frac{\partial\mathbf{Q}}{\partial x}\right) = J \cdot \left(\frac{\partial\mathbf{V}}{\partial x}\frac{\partial\boldsymbol{\Theta}}{\partial x}\right) \tag{3.25}$$

by taking $\frac{\partial \overline{\mathbf{V}}}{\partial x}, \frac{\partial \overline{\mathbf{\Theta}}}{\partial x} = \mathbf{0}$ to get

$$\begin{bmatrix} \frac{\partial \widetilde{\mathbf{V}}}{\partial x} \\ \frac{\partial \widetilde{\mathbf{\Theta}}}{\partial x} \end{bmatrix} = J_{-,\sim}^{-1} \begin{bmatrix} \frac{\partial \overline{\mathbf{P}}}{\partial x} \\ \frac{\partial \overline{\mathbf{Q}}}{\partial x} \end{bmatrix}. \tag{3.26}$$

We can then calculate the sensitivities of $\widetilde{\mathbf{P}}, \widetilde{\mathbf{Q}}$ as

$$\begin{bmatrix} \frac{\partial \widetilde{\mathbf{P}}}{\partial x} \\ \frac{\partial \widetilde{\mathbf{Q}}}{\partial x} \end{bmatrix} = J_{\sim,\bullet} \begin{bmatrix} \frac{\partial \mathbf{V}}{\partial x} \\ \frac{\partial \mathbf{\Theta}}{\partial x} \end{bmatrix}. \tag{3.27}$$

We combine these sensitivities to get the local sensitivities $\frac{\partial \mathcal{M}_i}{\partial x} = [\frac{\partial P_i}{\partial x}, \frac{\partial Q_i}{\partial x}, \frac{\partial V_i}{\partial x}, \frac{\partial \Theta_i}{\partial x}]$ at each bus in the system.

Similarly, for a perturbation in $\overline{\mathbf{V}}$ or $\overline{\mathbf{\Theta}}$ values, we can calculate the sensitivities such that $\frac{\partial \overline{\mathbf{P}}}{\partial x}, \frac{\partial \overline{\mathbf{Q}}}{\partial x} = \mathbf{0}$ from (3.25) to find $\frac{\partial \widetilde{\mathbf{V}}}{\partial x}, \frac{\partial \widetilde{\mathbf{\Theta}}}{\partial x}$ as

$$\begin{bmatrix} \frac{\partial \widetilde{\mathbf{V}}}{\partial x} \\ \frac{\partial \widetilde{\mathbf{\Theta}}}{\partial x} \end{bmatrix} = -J_{-,\sim}^{-1} \cdot J_{-,-} \begin{bmatrix} \frac{\partial \overline{\mathbf{V}}}{\partial x} \\ \frac{\partial \overline{\mathbf{\Theta}}}{\partial x} \end{bmatrix}. \tag{3.28}$$

Then we can find $\frac{\partial \widetilde{\mathbf{P}}}{\partial x}, \frac{\partial \widetilde{\mathbf{Q}}}{\partial x}$ from (3.27).

Note that these first-order power flow sensitivity calculations are linear, such that we can combine the sensitivities from multiple power flow parameter perturbations by summing the sensitivity values for each individual power flow parameter perturbation.

Sensitivity to Line Admittances: A change in a network line admittance value will also result in a change in the power flow values. Assuming that the line admittance is perturbed such that τ, the R/L line impedance ratio, is held constant then from $G_{ik}+B_{ik}j = \frac{R_{ik}-X_{ik}j}{R_{ik}^2+X_{ik}^2}$ we have $B_{ik} = \frac{1}{X_{ik}(1+(\tau/\omega_{\text{nom}})^2)}, G_{ik} = \tau/\omega_{\text{nom}} \cdot B_{ik}$. The first-order sensitivity of the dependent power injection values to a change in this admittance can be calculated from (3.19) for bus i by taking the partial derivatives

$$\begin{bmatrix} \frac{\partial P_i}{\partial B_{ik}} \\ \frac{\partial Q_i}{\partial B_{ik}} \end{bmatrix} = \begin{bmatrix} V_i^2 \left(\frac{\tau}{\omega_{\text{nom}}}\right) - V_i V_k \left(\frac{\tau}{\omega_{\text{nom}}} \cos\theta_{ik} - \sin\theta_{ik}\right) \\ V_i^2 - V_i V_k \left(\frac{\tau}{\omega_{\text{nom}}} \sin\theta_{ik} + \cos\theta_{ik}\right) \end{bmatrix} \tag{3.29}$$

Swapping the indices in 3.29 we can calculate the power flow sensitivities at the other bus k that the line connects to. Taking these sensitivities, we define $\mathbf{b} = (\partial\mathbf{P}/\partial B_{ik},\ \partial\mathbf{Q}/\partial B_{ik})$

The vector $\mathbf{b}$ is accordingly partitioned into $\overline{\mathbf{b}}$ (corresponding to the specified injections) and $\widetilde{\mathbf{b}}$ (corresponding to the dependent injections). Because the specified power parameters are held constant during the line perturbation, their variations are zero.

Applying these boundary conditions, the first-order Taylor expansion of the power flow equations reduces to:

$$\mathbf{0} = J_{-,\sim} \begin{bmatrix} \frac{\partial \widetilde{\boldsymbol{\Theta}}}{\partial B_{ik}} \\ \frac{\partial \widetilde{\mathbf{V}}}{\partial B_{ik}} \end{bmatrix} + \overline{\mathbf{b}} \tag{3.30}$$

where $J_{-,\sim}$ is the sub-matrix of the partitioned power flow Jacobian containing the partial derivatives of the specified injections with respect to the unknown state variables.

The sensitivities of the unknown voltage magnitudes and angles across the network can therefore be solved directly as

$$\begin{bmatrix} \frac{\partial \widetilde{\boldsymbol{\Theta}}}{\partial B_{ik}} \\ \frac{\partial \widetilde{\mathbf{V}}}{\partial B_{ik}} \end{bmatrix} = -J_{-,\sim}^{-1} \overline{\mathbf{b}}. \tag{3.31}$$

Finally, the sensitivities of the dependent power injections can be evaluated using the Jacobian sub-matrix $J_{\sim,\sim}$ corresponding to those dependent variables:

$$\begin{bmatrix} \frac{\partial \widetilde{\mathbf{P}}}{\partial B_{ik}} \\ \frac{\partial \widetilde{\mathbf{Q}}}{\partial B_{ik}} \end{bmatrix} = J_{\sim,\sim} \begin{bmatrix} \frac{\partial \widetilde{\boldsymbol{\Theta}}}{\partial B_{ik}} \\ \frac{\partial \widetilde{\mathbf{V}}}{\partial B_{ik}} \end{bmatrix} + \widetilde{\mathbf{b}}. \tag{3.32}$$

Eigenvalue Sensitivity

Sensitivities to Power Flow: Using these power flow sensitivities we can determine the system eigenvalue sensitivities using intermediate device impedance sensitivity to changes in the local power flow parameters. This requires the sensitivities $\frac{\partial Z_k}{\partial \mu}$ of the device impedances with respect to local power flow parameters. Equivalent device impedances and these sensitivities can be obtained from analytical models or from frequency-scanning measurements, as described in [116]. Because most black-box inverter models allow specification of power references, this approach is also compatible with black-box modeling when using frequency scanning to obtain device admittance models.

Using the chain rule, we can determine the sensitivity of a device's equivalent impedance to a power flow parameter x as

$$\frac{\partial Z_k}{\partial x} = \sum_{\mu \in \mathcal{M}_k} \frac{\partial Z_k}{\partial \mu} \frac{\partial \mu}{\partial x}. \tag{3.33}$$

We collect these as $\frac{\partial \mathbf{Z}_{\text{dev}}}{\partial x} = \text{diag}(\frac{\partial Z_1}{\partial x}, \frac{\partial Z_2}{\partial x}, \dots, \frac{\partial Z_K}{\partial x})$. Then using (3.14) and (3.23), noting that $\frac{\hat{\mathbf{Y}}^{-1}}{\partial x} = \frac{\hat{\mathbf{Z}}}{\partial x}$ (See Appendix B.1 for details) we calculate the sensitivity of the system eigenvalue λ as

$$\frac{\partial \lambda}{\partial x} = \left\langle -\text{Res}_\lambda^* \hat{\mathbf{Y}}, \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial x} \right\rangle. \tag{3.34}$$

Since $\mathbf{Z}_{\text{dev}}$ is a block-diagonal matrix, this formulation can be localized to

$$\frac{\partial \lambda}{\partial x} = \sum_{k \in \mathcal{N}} \langle -\text{Res}_\lambda^* \hat{\mathbf{Y}}_{kk}, \frac{\partial Z_k}{\partial x} \rangle,$$

where $\hat{\mathbf{Y}}_{kk}$ are 2×2 diagonal blocks of $\hat{\mathbf{Y}}$ corresponding to bus k, which can be calculated from the local device impedance and the grid impedance seen from that device's POI [116].

Sensitivities to Line Admittance Perturbing the admittance of a line, not only impacts the device equivalent admittances through perturbations to the power flow parameters, but it also directly impacts the eigenvalues through the perturbation to the network admittance transfer function matrix $\mathbf{Y}_{\text{net}}$. The sensitivity of the network admittance $\mathbf{Y}_{\text{net}}$ to a specific line susceptance B_{ij} is

$$\frac{\partial \mathbf{Y}_{\text{net}}(s)}{\partial B_{ij}} = \mathcal{L}_{ij} \otimes \begin{bmatrix} \beta(s) & \alpha(s) \\ -\alpha(s) & \beta(s) \end{bmatrix} \tag{3.35}$$

where $\mathcal{L}_{ij}$ is the Laplacian matrix with only the entries for the line being perturbed, B_{ij} is the admittance of the line, $\otimes$ represent the Kronecker product, and $\alpha(s) = \frac{\omega_0}{(s+\tau)^2/\omega_0+\omega_0}$, $\beta(s) = \frac{s+\tau}{(s+\tau)^2/\omega_0+\omega_0}$. Then using (3.14) and (3.23), we calculate the eigenvalue sensitivity as

$$\frac{\partial \lambda}{\partial B_{ij}} = \left\langle -\text{Res}_\lambda^* \hat{\mathbf{Y}}, \mathbf{Y}_{\text{net}}^{-1} \frac{\partial \mathbf{Y}_{\text{net}}(s)}{\partial B_{ij}} \mathbf{Y}_{\text{net}}^{-1} + \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial B_{ij}} \right\rangle. \tag{3.36}$$

where $\frac{\partial \mathbf{Y}_{\text{net}}(s)}{\partial B_{ij}}$ comes from (3.35) and $\frac{\partial \mathbf{Z}_{\text{dev}}}{\partial B_{ij}}$ is obtained by calculating (3.33) for each device using the power flow sensitivities described in Section 3.2.2. The derivation of this can be found in Appendix B.2.

Sensitivities of Load Models

In small-signal stability analysis, static constant power loads are modeled as equivalent parallel shunts. Linearized around a steady-state operating point, the equivalent shunt admittance $Y_{Lk} = G_{Lk} + jB_{Lk}$ at bus k is calculated from the specified power demand (P_{Lk}, Q_{Lk}) and steady-state voltage magnitude V_k as: $G_{Lk} = P_{Lk}/V_k^2$, $B_{Lk} = -Q_{Lk}/V_k^2$. Loads can also be modeled as ZIP loads, where portions of the demand are classified as constant impedance, constant current, and constant power. The selection of these ZIP proportions significantly impacts the damping of the system; constant impedance components increase system damping, whereas constant power components introduce negative incremental resistance that reduces damping.

These equivalent load admittances are embedded within the network matrix $\mathbf{Y}_{\text{net}}(s)$. Therefore, a change in a local power flow parameter x induces a sensitivity in the network matrix relating to the load:

$$\frac{\partial \mathbf{Y}_{\text{net}}(s)}{\partial x} = E_{kk} \otimes \frac{\partial L_{kk}(s)}{\partial x} \tag{3.37}$$

where E_{kk} is an $N \times N$ standard basis matrix with a 1 at the k-th diagonal, and $L_{kk}(s)$ is the 2×2 admittance transfer function matrix of the load. $\frac{\partial L_{kk}(s)}{\partial x}$ is found by differentiating G_{Lk} and B_{Lk} with respect to x.

The eigenvalue sensitivity to this load perturbation is then evaluated using the network sensitivity formulation in (3.36). For a structural network perturbation (like a line outage) that simultaneously alters load voltages, the total $\frac{\partial \mathbf{Y}_{\text{net}}}{\partial B_{ij}}$ is the sum of the direct line contribution from (3.35) and this indirect load contribution.

Note that loads can also be represented with dynamic models (e.g., induction motors or inverters) rather than static steady-state equivalents. In such cases, the dynamic load is modeled as a standard power electronic or electromechanical apparatus, and its power-flow-driven impedance sensitivities are incorporated into the device matrix $\mathbf{Z}_{\text{dev}}(s)$ and evaluated using (3.2.2).

3.2.3 Small-Signal Stability Metrics

In this section we define system strength metrics based on the first-order approximation of eigenvalue shifts to changes in system parameters. Using the eigenvalue sensitivity $\frac{\partial \lambda}{\partial \mu}$ we can calculate a scaled eigenvalue perturbation $\Delta_\mu \lambda$ as

$$\Delta_\mu \lambda = \frac{\partial \lambda}{\partial \mu} \Delta \mu,$$

where $\Delta \mu$ can reflect the sensitivity to an absolute change in the parameter ($\Delta \mu = 1$ pu) or a proportion of the base parameter value ($\Delta \mu = \mu_0$). We define the sensitivity index for an eigenvalue λ and parameter μ at bus i as

$$\mu \mathrm{SI}_{i,\lambda} = \frac{\mathrm{re}(\Delta_\mu \lambda)}{|\mathrm{re}(\lambda)|}. \tag{3.38}$$

This normalizes the sensitivity of the real part of an eigenvalue by the magnitude of the real part of that eigenvalue. The smaller the real part of an eigenvalue and the higher its sensitivity, the larger the sensitivity index will be. From this we can consider the maximum normalized eigenvalue shift towards the right hand plane for an increase (to) or decrease (from) in the power value and define the to sensitivity index (μTSI) and from sensitivity index (μFSI) as

$$\mu \mathrm{TSI}_i = \max_\lambda (\mu \mathrm{SI}_{i,\lambda}), \quad \mu \mathrm{FSI}_i = \max_\lambda (-\mu \mathrm{SI}_{i,\lambda}). \tag{3.39}$$

The inverse of the μTSI and μFSI sensitivity indices give a first-order approximate of how large a perturbation to the parameter can be before the eigenvalue will cross into the RHP and the system will become unstable. With this intuition, it is clear that larger sensitivities indicates a weaker system which we show in the case study results in the following section.

3.2.4 Calculation Scalability

This sensitivity-based approach is significantly faster than recalculating the full system eigenvalues at a perturbed operating point, as it leverages system residues to compute sensitivities through simple linear algebraic operations and can be calculated for a selected subset of the system eigenvalues. Utilizing selective eigenvalue algorithms [137] to acquire the base eigenvalues and sparse state-space representations to avoid the numerical precision challenges of high-order transfer function matrices allows the method to be applied to large-scale systems. Acquiring equivalent admittance models and their sensitivities from black-box models through vector fitting [138] is the most intensive step; however, these device-level admittances are independent and can be calculated in parallel or pre-calculated offline for a range of operating points.

When evaluating large-scale systems with high state dimensionality, specific computational strategies are required. Large time-scale separation in power systems yields highly stiff state matrices, while the dense clustering of network modes causes their corresponding eigenvectors to become nearly parallel. Together, these factors generate an ill-conditioned global eigenbasis that induces severe numerical instability and precision loss in standard dense eigensolvers [139, 140]. To overcome this, selective eigenvalue algorithms [137, 141] can be utilized to project the system into targeted frequency bands, allowing for the accurate extraction of critical eigen-pairs. Furthermore, sparse state-space representations are employed to avoid the numerical precision challenges inherent to high-order transfer function matrices.

The computational scalability of the proposed framework is demonstrated with the systems described in Table 3.2. Fig. 3.4 compares the runtimes for calculating the eigenvalues sensitivities with our analytical approach against sampling the perturbed eigenvalues for a perturbation in the parameter of interest and see a significant increasing reduction in computation time as the system size grows. The total runtimes consider sampling all specified power flow parameters and all non-bridging line's network admittances for a selection of critical eigenvalues that lie within a specified set.

3.2.5 Case Study

We perform case studies to validate the eigenvalue sensitivity calculations and test the performance of our proposed metrics in identifying weak buses and informing remedial actions. These case studies use a modified version of the Simplus Grid Tool [142] available at `https://github.com/TragerJoswig-Jones/Simplus-Grid-Tool-System-Sensitivity`.

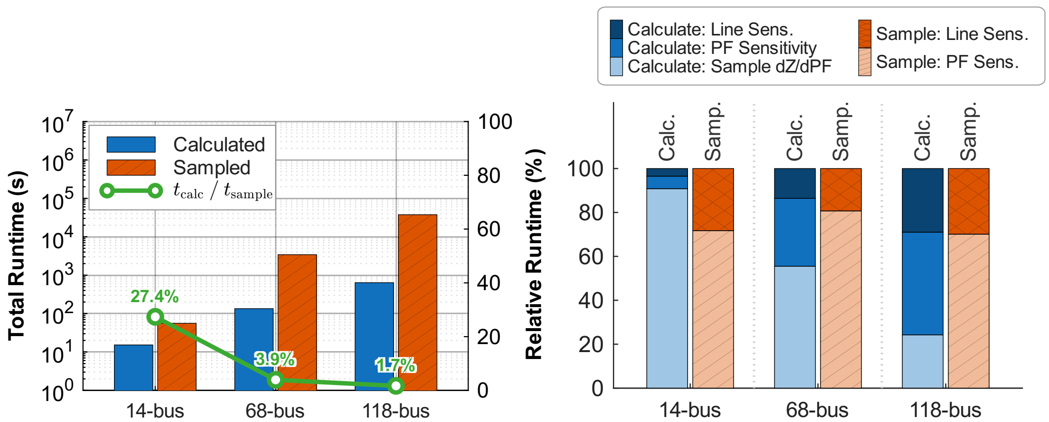

Figure 3.4: Runtimes comparing calculating the eigenvalue sensitivities to resampling the eigenvalues for a perturbation in a parameter for a range of systems. P_G, Q_G, V sensitivities are tested for all specified power flow values at buses, and the B_{ij} sensitivities are tested for all non-bridging lines.

Table 3.2

Test System Parameters

System	Buses	Lines	Devices	States	Considered Eigenvalues	Considered Lines
14-bus	14	6	20	249	16	19
68-bus	68	83	29	1041	35	65
118-bus	118	179	26	2130	66	170

Table 3.3

IEEE 14-BUS CASE STUDY CASE DESCRIPTIONS

Case	**Changes from Base Case**
Case A	$P_G : 0.5 \rightarrow 1.2$ pu at bus **8** and bus **10**
Case B	$P_G : 0.5 \rightarrow 1.2$ pu at bus **6** and bus **10**
Case A-2	$P_G : 0.5 \rightarrow 1.2$ pu at bus **8** and bus **10**, Q_G: $0.05 \rightarrow 0.5$ pu at bus **8**.
Case A-3	$P_G : 0.5 \rightarrow 1.2$ pu at bus **8** and bus **10**, Bus **8** GFL to GFM control.

3.2.6 14-Bus Test System Results

First, we use a modified IEEE 14-bus system where the synchronous condensers are replaced by identical GFL IBRs, with one additional IBR at bus 10.For this system, we choose to focus on perturbations to specified active and reactive power injections and model all buses other than the slack bus as PQ buses. The inverters are all selected to have identical GFL control schemes and vary their current injections to align with active and reactive power set points. We compare different generation cases detailed in Table 3.3.

First, we validate that the proposed methodology for calculating eigenvalue perturbations produces accurate results. We calculate the actual eigenvalue perturbations by perturbing the power injection at a bus, calculating the new steady-state power flow for the system, and then computing the resulting eigenvalue perturbation. As shown in Table 3.4, the actual and predicted eigenvalue perturbations align closely.

Next, we calculate the proposed power flow sensitivity metrics proposed in Section 3.2.3 with a selection of eigenvalues for the Base Case scenario. We also plot the system's Impedance Margin Ratio (IMR) values in Fig. 3.6, and compare this to our metrics shown in Fig. 3.5. We see that IBR buses identified as weak by IMR are also identified as being weak by the PTSI and QFSI metrics; however, these metrics link this weakness to shifts in power injections at these buses. Although similar to IMR in their use of admittance residues, the

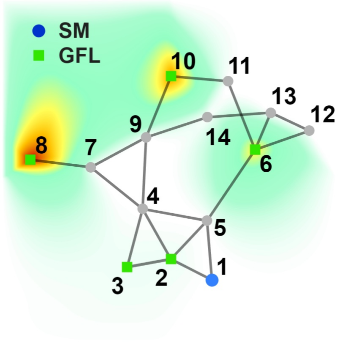

(a) PTSI

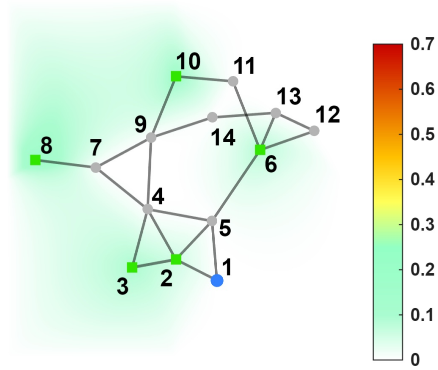

(b) PFSI

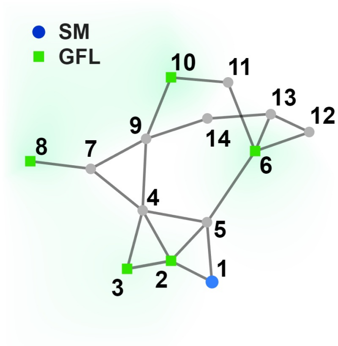

(c) QTSI

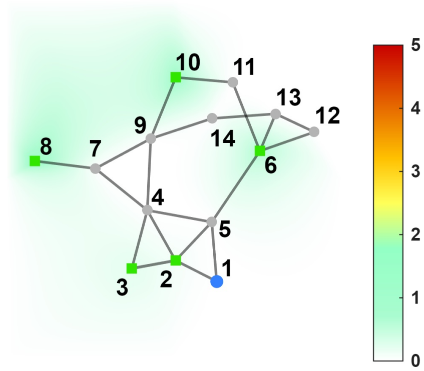

(d) QFSI

Figure 3.5: 14-bus system base case generator active and reactive power injection sensitivities.

proposed metrics distinguish themselves by linking eigenvalue sensitivities directly to power flow perturbations rather than relying on conservative, worst-case impedance margins.

Next, we apply increases in active power injection at buses with differing PTSI values and simulate the nonlinear time-domain system to see how this impacts its response to a step in the active power load at bus 5. Both cases increase the power injection at bus 10, but Case A increases generation at the more sensitive bus 8 IBR and Case B selects the less sensitive IBR at bus 6. Fig. 3.7 shows that the increase in power injected at bus 8 with the higher PTSI value results in a maintained oscillation post-disturbance, while the same increase in power injected at the less sensitive bus 6 results in the system maintaining

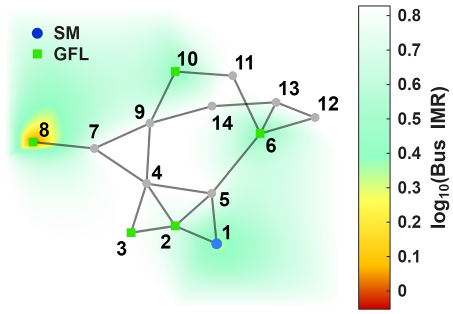

Figure 3.6: 14-bus system base case IMR values.

adequate damping of this mode.

Lastly, we explore two remedial actions to mitigate the subsynchronous oscillation introduced in Case A. Our first action is informed by looking at the reactive power sensitivities for Case A in Fig. 3.8. According to the QFSI values, an increase in the reactive power injection at bus 8 will result in some eigenvalue shifting to the left and increase the damping of its associated mode. The QTSI value at bus 8 being low indicates that this increase will not significantly shift any eigenvalues towards the right. Following this, in Case A-2 we increase the reactive power injection of the IBR at bus 8 and observe a significant increase in the damping of the oscillatory mode in Fig. 3.7. Next in Case A-3, we change the control mode of the IBR at bus 8 from GFL to a grid-forming (GFM) droop controller and plot the active power sensitivity injections for the Base Case in Fig. 3.9. Fig. 3.9a shows that the system is no longer sensitive to increases in active power injection at bus 8, which we validate by simulating Case A with GFM control at bus 8 and seeing an improved response in Fig. 3.7.

3.2.7 118-Bus Test System Results

For further studies on the network admittance and voltage sensitivities, we use the IEEE 118-bus test system [143, 144]. We modify the system to contain a mix of SM and IBR devices including GFLs and GFMs. The SMs are represented with 8-th order generator models with exciters, power stabilizers, and governors. Load buses and most GFLs are

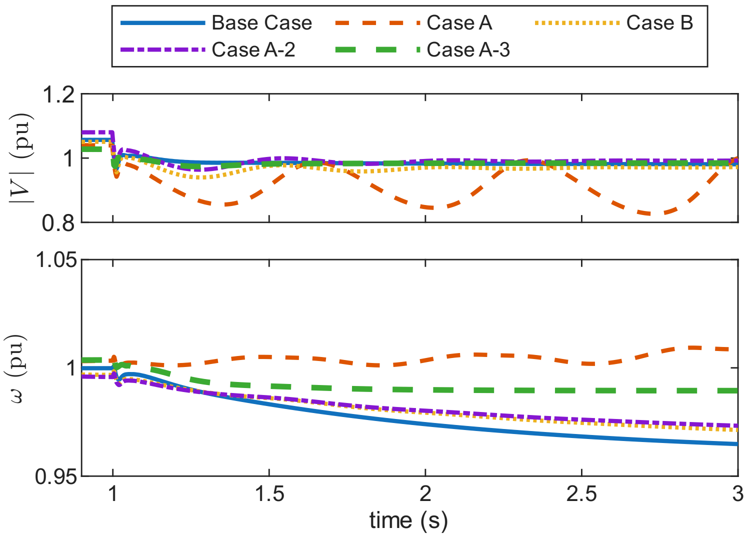

Figure 3.7: Time-domain response of bus 10 voltage magnitude and frequency for the cases in Table 3.3. A 0.67 pu load step at bus 5 occurs at $t = 1$ s. Case A exhibits an undamped 1.4 Hz oscillation, which is not prevalent in Case B. Increasing reactive power injected at bus 8 (Case A-2) or changing its control from GFL to GFM (Case A-3) mitigate the oscillation introduced in Case A.

Table 3.4

Comparison of predicted and actual eigenvalue perturbations for three modes of the Base Case 14-bus test system.

Perturbation		**Eigenvalue Perturbation** ($\Delta\lambda \times 10^5$)		
		$\lambda = -10.610 + 0.000\,j$	$\lambda = -10.791 + 3.382\,j$	$\lambda = -159.545 + 427.018\,j$
$\Delta P_8 = 1$ pu	Actual	0.016	0.302 - 0.183j	0.456 + 0.073j
	Pred.	0.016	0.302 - 0.183j	0.455 + 0.073j
	Error	0.28%	0.02%	0.20%
$\Delta Q_8 = 1$ pu	Actual	0.522	-0.645 + 0.008j	2.985 + 0.935j
	Pred.	0.522	-0.645 + 0.008j	2.983 + 0.933j
	Error	0.00%	0.01%	0.09%
$\Delta V_1 = 1$ pu	Actual	2.330	-3.455 - 5.890j	18.035 + 4.269j
	Pred.	2.330	-3.455 - 5.890j	18.036 + 4.266j
	Error	0.00%	0.00%	0.02%
$\Delta B_{5,6} = 1$ pu	Actual	0.112	0.188 - 0.075j	20.251 + 1.385j
	Pred.	0.112	0.188 - 0.075j	20.248 + 1.386j
	Error	0.04%	0.04%	0.02%

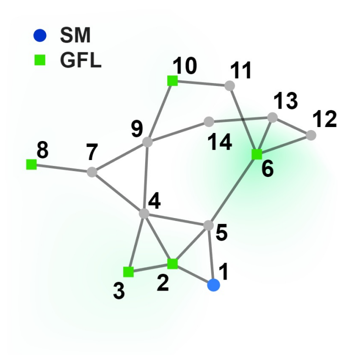

(a) QTSI

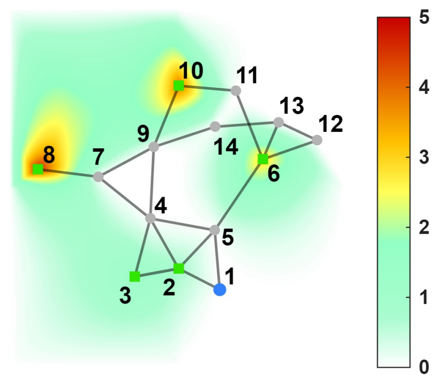

(b) QFSI

Figure 3.8: 14-bus generator reactive power injection sensitivities for case A.

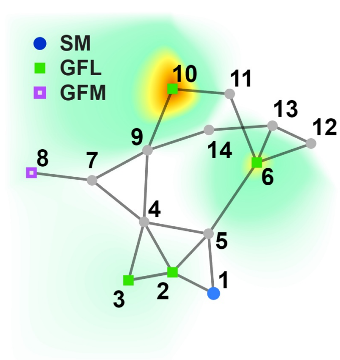

(a) PTSI

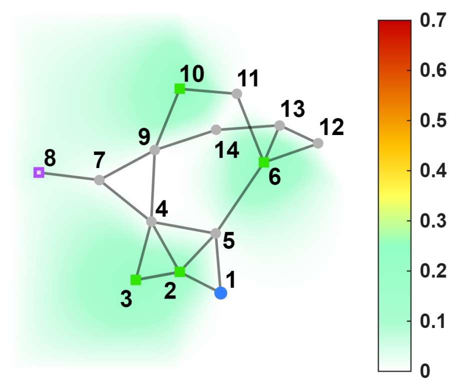

(b) PFSI

Figure 3.9: 14-bus generator active power injection sensitivities for base case with a GFM at bus 8.

modeled as PQ buses, while SMs and GFMs are modeled as PV buses (besides bus 1, which is taken to be the slack bus). The system is initialized with a set of lines that are switched off. We linearize the loads around the initial operating point and maintain this constant impedance between perturbations.

We calculate the line admittance sensitivities for eigenvalues in the set $\{\lambda = \sigma + j\omega \,|\, \lambda \in \text{poles}(\hat{\mathbf{Y}}(s)), |\lambda| > 0.016 \text{ Hz}, -40 \text{ Hz} < \sigma, \omega \neq 0 \text{ Hz}, |\omega - \omega_{\text{nom}}| > 0.016 \text{ Hz}, |\omega| < 160 \text{ Hz}\}$. We plot the sensitivity indices for a decrease and increase in the line admittances in Fig. 3.10a and Fig. 3.10b, respectively. Fig. 3.10a shows that the system has a high sensitivity to decreases in the admittance of the line from bus 103 to 110, while Fig. 3.10b shows that it has low sensitivity to increases in this line's admittance. Looking at line 100-103, Fig. 3.10 shows that the system has a moderate sensitivity to increases in the admittance and a low sensitivity to decreases in its admittance. From these values we can anticipate that switching off (i.e. decreasing the admittance to zero) line 103-110 will have a more detrimental impact on the system's small-signal stability, than switching off line 100-103.

In Fig. 3.11, we test the system's time-domain response to switching of these lines. We switch line 100-103 off and on, at $t = 0.5$ s and $t = 1.0$ s, and see the system return to a new stable steady-state damping out any induced oscillations. At $t = 2$ s we switch off line 103-110 and see that the system becomes unstable with growing undamped oscillation of 95 Hz. With 'No Action' taken, the system states diverge.

Looking for remedial actions, we calculate the reactive power injection and voltage magnitude sensitivities of only the unstable eigenvalue ($\lambda = 1.59 + 600.85j$ Hz) with line 103-110 off. Fig. 3.12, shows that with this line off the system has high sensitivities to increases in the reactive power at bus 112 and the voltage at bus 100. In the 'Remedial Action' simulation, we decrease these values at $t = 2.2$ s and the oscillation is dampened out stabilizing the system. At $t = 2.5$ s we switch line 103-110 back on.

Returning to Fig. 3.2, we can observe that the first-order eigenvalue sensitivities may not provide perfect predictions of the nonlinear trace of the eigenvalues for large steps in a parameter, such as when switching a line admittance off, however indices defined using these sensitivities can provides valuable insights into which lines have little affect on the system stability and which lines the system is sensitive to and should be further studied.

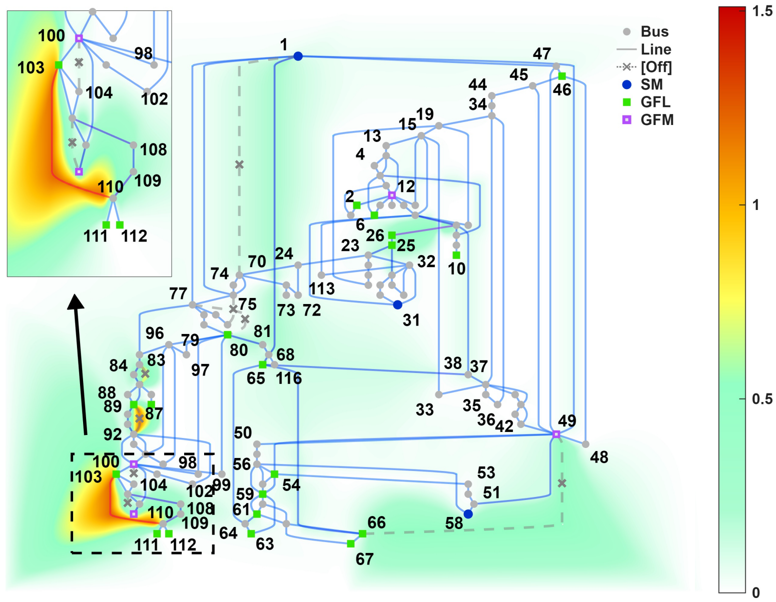

(a) $\mathrm{B_{ij}FSI}$

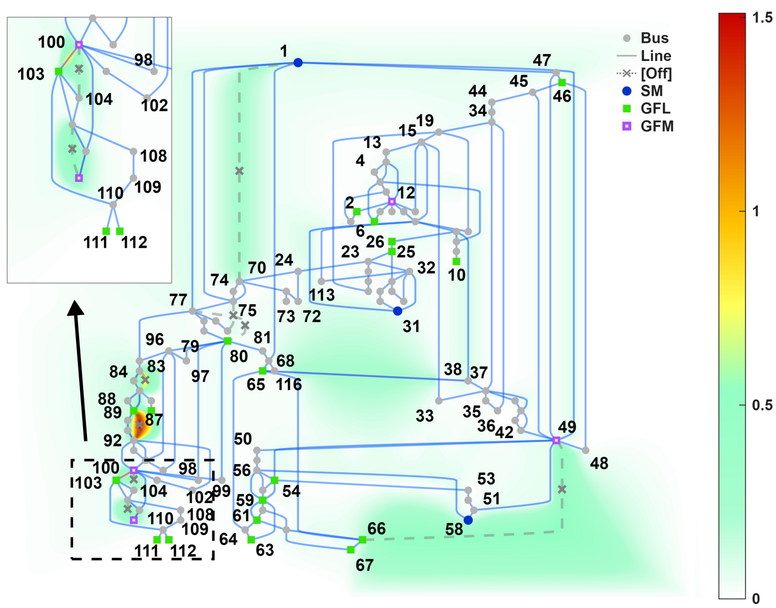

(b) $\mathrm{B_{ij}TSI}$

Figure 3.10: 118-bus line admittance sensitivities.

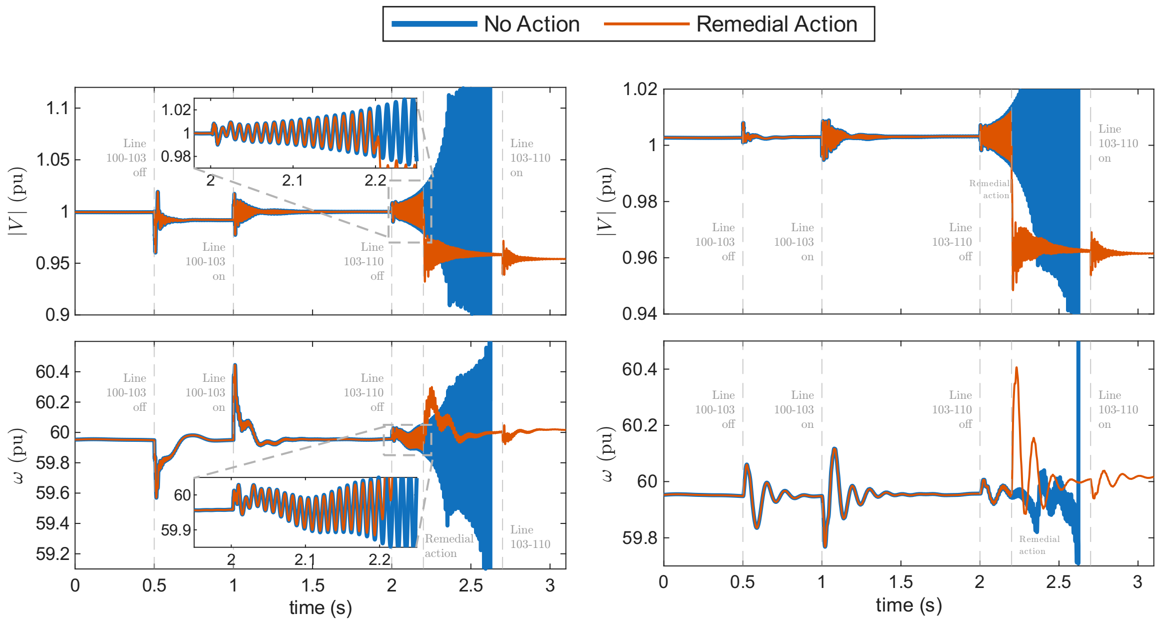

(a) Bus 103 GFL

(b) Bus 100 GFM

Figure 3.11: Time-domain response of bus 103 and bus 100 voltage magnitude and measured frequency when switching the line between bus 100 and bus 103 off at time $t = 0.5$ s and on at $t = 1.0$ s, then switching line 103-110 off at $t = 0.2$ s and on at $t = 2.7$ s. Following the line 103-110 switch, the system becomes unstable if no action is taken; the simulation diverges at $t \approx 2.6$ s. Applying a remedial action decreasing the voltage at bus 100 by 0.1 p.u and decreasing Q_G at bus 112 by 0.3 p.u at $t = 2.2$ s resolves the instability.

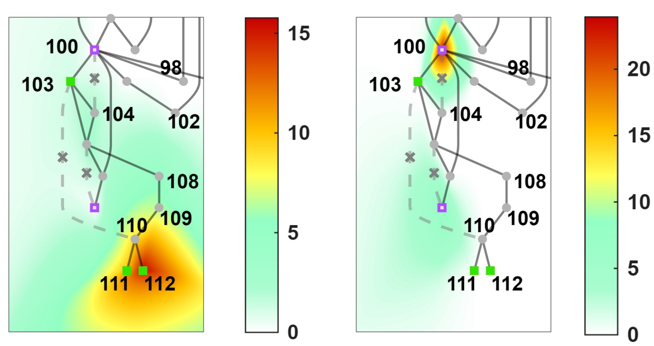

(a) QTSI

(b) VTSI

Figure 3.12: Generator reactive power injection and voltage increase unstable eigenvalue ($\lambda = 1.59 + 600.85j$ Hz) sensitivities for a portion of the 118-bus system with line 103-110 switched off.

Chapter 4

UTILIZING INVERTERS FOR VOLTAGE REGULATION IN DISTRIBUTION SYSTEMS

In this section we focus on the control of inverter-based resources as it relates to voltage regulation. In Section 4.1, we discuss the challenge of regulating voltages with large data center loads and then introduce the system model used in the following studies. In Section 4.2, we propose a switching-reference voltage control approach for regulating the voltage at the grid interconnection of a data center. Then in Section 4.3, we explore the viability of using data center uninterruptible power supply energy storage headroom to provide active power regulation of the data center load.

4.1 Introduction, Background & Modeling

The rapid growth of artificial intelligence (AI) has led to the deployment of increasingly large training workloads in hyperscale data centers. Empirical scaling laws show that model performance improves predictably with increasing model size, data, and computational budget [145, 146]. As a consequence, electricity demand from AI training has also grown rapidly [147]. At the same time, the scale and geographic concentration of modern data centers make their power consumption an increasingly important factor in power system operation.

Data center loads themselves are interfaced to the grid through inverters that can be actively leveraged to help regulate the impact that data center loads have on the grid voltage by adjusting their reactive power output. Further, data centers are often equipped with significant portions of battery-energy storage that sits idle waiting to power the data center in the event that the center decides to disconnect from the grid and switch to its backup power supply. Data center loads can be volatile, significantly impacting the voltage at their point of interconnection, but are inherently built with non-insignificant resources that can be used to mitigate their impact on the grid.

4.1.1 Data Center Loads

Data center loads vary significantly over relatively short timescales, particularly when running Large Language Model (LLM) inference and training tasks [148]. A representative AI training power trace is shown in Fig. 4.1, characterized by rapid ramps and periodic fluctuations. This distinctive pattern arises from the distributed training of large AI models across multiple GPU racks: power demand increases substantially during computation in each training iteration and drops sharply during communication. These abrupt load transitions constitute the dominant source of voltage deviations along the distribution feeder.

Since the most severe voltage violations arise during these abrupt transitions, existing regulation methods designed for smoothly varying loads often fail to respond sufficiently fast to mitigate the worst violations. Furthermore, responding to repeated large power fluctuations using traditional methods can incur substantial control costs.

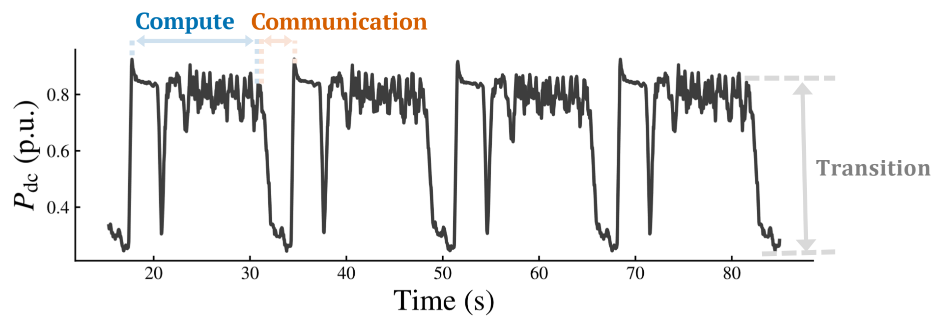

Figure 4.1: Per unit power readings from an at-scale training job on DGX H100 racks [160, 161].

One common approach to voltage regulation in distribution systems is network-wide reactive power control. Traditionally achieved using mechanical devices such as tap-changing transformers, this method is unsuitable for frequent adjustments due to equipment lifespan limitations. Inverter-based resources have therefore emerged as valuable assets for voltage regulation due to their ability to update reactive power setpoints rapidly. Accordingly, voltage droop control and its variants have been widely proposed [149, 150, 151, 152]. However, these methods are primarily designed for slowly varying loads.

Beyond grid-side regulation, recent studies have treated data centers as active participants by leveraging internal infrastructure for voltage support. This includes workload-aware power capping [153], server-level power modulation [154, 148], coordinated control of UPS and cooling systems [155], and GPU utilization control [156]. Some works have further considered coordination with external grid assets, such as grid-forming battery systems [157, 158] and EV charging stations [159]. However, these approaches often overlook the underlying power dynamics of modern AI training workloads, leading to unnecessarily large control efforts with limited improvement in worst-case voltage deviations.

4.1.2 Distribution System Modeling

To evaluate how active power transitions from AI workloads impact bus voltages across the network, the physical power flow must be mathematically defined. We model a radial distribution feeder system as a directed tree graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ represents the set of buses and $\mathcal{E}$ represents the set of distribution lines. The substation at the head of the feeder is designated as the slack bus, establishing the reference voltage.

DistFlow Power Flow Formulation

For the accurate physical simulation of the distribution network, we employ the DistFlow model [162], which represents the true nonlinear branch flow equations. For each line $(i, j) \in \mathcal{E}$ connecting bus i to bus j at time t, the real and reactive power flows are governed by:

$$P_{ij,t} = p_{j,t} + \sum_{k:(j,k)\in\mathcal{E}} P_{jk,t} + r_{ij} l_{ij,t} \tag{4.1a}$$

$$Q_{ij,t} = q_{j,t} + \sum_{k:(j,k)\in\mathcal{E}} Q_{jk,t} + x_{ij} l_{ij,t} \tag{4.1b}$$

$$v_{j,t}^2 = v_{i,t}^2 - 2(r_{ij} P_{ij,t} + x_{ij} Q_{ij,t}) + (r_{ij}^2 + x_{ij}^2) l_{ij,t}, \tag{4.1c}$$

where $p_{j,t}$ and $q_{j,t}$ are the net active and reactive power injections at bus j, r_{ij} and x_{ij} are the resistance and reactance of line (i, j), and $v_{i,t}$ is the voltage magnitude at bus i. The term $l_{ij,t} = (P_{ij,t}^2 + Q_{ij,t}^2)/v_{i,t}^2$ represents the squared magnitude of the complex line current, which explicitly captures the nonlinear power losses. These equations are used to simulate the physical system voltages. For radial networks, the non-convex equality defining the line current can be relaxed to an inequality ($l_{ij,t} \geq (P_{ij,t}^2 + Q_{ij,t}^2)/v_{i,t}^2$), forming a convex Second-Order Cone (SOC) constraint [163]. Although the SOC relaxation provides a useful approach for power flow routing, it can still introduce significant computational overhead for receding-horizon control, and it is mathematically cumbersome for deriving closed-form stability guarantees.

LinDistFlow Approximation

While the exact DistFlow model is necessary for accurate physical simulation, its non-linearities make it computationally burdensome for receding-horizon optimal control and mathematically intractable for deriving closed-form stability guarantees.

To facilitate controller synthesis—specifically when establishing the error-tracking contraction of voltage references or incorporating network-wide voltage constraints into a convex framework—the system is linearized. By assuming the line losses ($l_{ij,t}$) are relatively small and the bus voltages remain close to the nominal 1.0 p.u. reference, the higher-order loss terms are neglected. This yields the LinDistFlow approximation [162, 150]:

$$v_{j,t} \approx v_{i,t} - (r_{ij}P_{ij,t} + x_{ij}Q_{ij,t}). \tag{4.2}$$

By aggregating these linearized equations across the entire radial network, the relationship between nodal power injections and bus voltages can be expressed compactly in matrix form:

$$\boldsymbol{v}_t = \mathbf{R}\boldsymbol{p}_t + \mathbf{X}\boldsymbol{q}_t + \mathbf{1}, \tag{4.3}$$

where $\boldsymbol{v}_t$, $\boldsymbol{p}_t$, and $\boldsymbol{q}_t$ are the vectors of bus voltage magnitudes, active power injections, and reactive power injections, respectively, and the all-ones vector $\mathbf{1}$ represents the nominal slack bus voltage. The positive definite matrices $\mathbf{R}, \mathbf{X} \in \mathbb{R}^{n\times n}$ capture the network-wide sensitivity mapping of power injections to bus voltages.

Traditional voltage regulation relies on decentralized Volt-VAR droop control, where inverter-based resources adjust $\boldsymbol{q}_t$ in response to (4.3) to drive voltages toward a fixed reference. However, because fixed-reference control responds only after a voltage change has occurred, it is fundamentally ill-equipped to preempt the severe step-transitions induced by AI workloads.

Data Center Workload Modeling

Modern synchronized LLM training workloads exhibit temporal structure in active power consumption that differs fundamentally from conventional slowly varying demand. As shown in Fig. 4.1, the aggregate data center power exhibits repeated switching between two dominant operating levels with abrupt transitions. We refer to the high-power phase (forward and backward propagation) as the *compute phase*, and to the lower-power phase (synchronization and checkpointing) as the *communication phase*. Because accelerators operate synchronously, the aggregate demand switches collectively between these two phases.

Motivated by this temporal structure, we model active power injection at the data center bus i as:

$$p_{i,t} = \bar{p}_i(m_t) + w_t, \qquad |w_t| \leq \bar{w}(m_t), \tag{4.4}$$

where $m_t \in \{0, 1\}$ denotes the mode, with 0 for communication and 1 for computation.

The injections $\bar{p}_i(0)$ and $\bar{p}_i(1)$ are the corresponding average demands, while w_t captures intra-phase fluctuations bounded by $\bar{w}(m_t)$.

During large-scale AI training, short-term active power variation is dominated by the training workload within the data center, and we regard the power injections at other buses as remaining nearly constant. Let $\bar{\boldsymbol{p}}(0), \bar{\boldsymbol{p}}(1) \in \mathbb{R}^n$ denote the network average active power injections associated with the communication and computation modes, respectively. For a data center connected at bus i, the two network injection profiles differ only at that bus:

$$\bar{\boldsymbol{p}}(0) - \bar{\boldsymbol{p}}(1) = (\bar{p}_i(0) - \bar{p}_i(1))\boldsymbol{e}_i, \tag{4.5}$$

where $\boldsymbol{e}_i$ is the i-th standard basis vector.

4.2 Regulation of Data Center Loads Using Switching-Reference Voltage Control[1]

In this work, we aim to answer the following question: *How should we design a voltage control law to mitigate voltage deviations induced by AI training workloads, while avoiding excessive control effort?*

We propose a decentralized switching-reference voltage control for distribution systems with AI-training data centers. The key idea is to exploit the structured switching dynamics of such loads to reduce worst-case voltage deviations with minimal control effort. We first characterize the two-mode structure of large-scale AI training workloads, and then design a switching-reference droop controller that aligns the voltage reference with the dominant operating level of the data center load. On this basis, we establish conditions for voltage convergence and characterize when voltage reference placement can effectively reduce voltage deviations. We further develop a rule for switching voltage references based solely on local voltage measurements, enabling simple local implementation while significantly reducing control effort. Case studies demonstrate that the proposed approach substantially reduces voltage deviations and reactive power control effort compared with conventional fixed-reference droop control. It remains effective under multiple data centers and data-center-level load smoothing, without the need for extensive coordination.

The decentralized switching-reference is defined at each bus i as:

$$v_{i,t}^{ref} = v_{i,t}^{bias} + s_{i,t} v_{i,t}^{amp}, \tag{4.6}$$

where $v_{i,t}^{bias}$ tracks the midpoint of the two operating levels, $v_{i,t}^{amp}$ captures half the magnitude of the mode-induced voltage shift, and the sign selector $s_{i,t} \in \{-1, +1\}$ acts as a local estimate of the unknown data center mode m_t. This control law is updated completely autonomously using only local voltage measurements. The rigorous mathematical proofs establishing the closed-loop contraction of the reference-tracking error and the optimization of the worst-case steady-state deviation are thoroughly detailed in [161].

[1] Adapted from M. Yan, T. Joswig-Jones, B. Zhang, Y. Chen and W. Cui, "Switching-Reference Voltage Control for Distribution Systems with AI-Training Data Centers," arXiv, 2026. [161]

By aligning the voltage reference with the dominant operating level, this approach significantly mitigates the worst-case voltage violations while drastically reducing the reactive power control effort required by grid-side inverters.

4.2.1 Grid-Side Voltage Control

Voltage Control and Challenges

Large and rapid variations in data center power usage affect voltages throughout the system and may drive them toward or beyond acceptable operating limits (e.g., within $\pm 5\%$ of nominal) [164].

Reactive power control is commonly used for voltage regulation in distribution systems. A widely adopted approach is decentralized Volt–VAR droop control [149, 151, 165], in which inverter-based resources adjust reactive power injections based on local voltage measurements. In discrete time, a typical droop update can be written as

$$\boldsymbol{q}_{t+1} = \boldsymbol{q}_t - \mathbf{K}\big(\boldsymbol{v}_t - \mathbf{1}\big), \tag{4.7}$$

where $\mathbf{K} = \mathrm{diag}(k_1, \ldots, k_n)$ with $k_i > 0$.

This droop control only uses current information to drive voltages toward a fixed reference (e.g., 1 p.u. in (4.7)). However, when applied to step power transitions induced by AI training workloads, the controller repeatedly reacts to each transition without accounting for the underlying two-mode structure of the training workload. As a result, large voltage deviations can persist during frequent mode transitions.

To address the above challenges, we propose a decentralized grid-side voltage control framework that exploits the structured power dynamics of AI training workloads. The key idea is to incorporate the characteristic step power changes of data centers into the design of voltage control, as illustrated in Fig. 4.1. As we will show, by aligning voltage regulation with this structured load behavior, voltages can be maintained within the desired range with much smaller and less frequent reactive power injections. Importantly, the proposed controller operates independently of internal power management in data centers. If internal resources such as battery storage partially smooth the active power profile [166], the grid-side controller remains effective and compatible with internal data center control strategies.

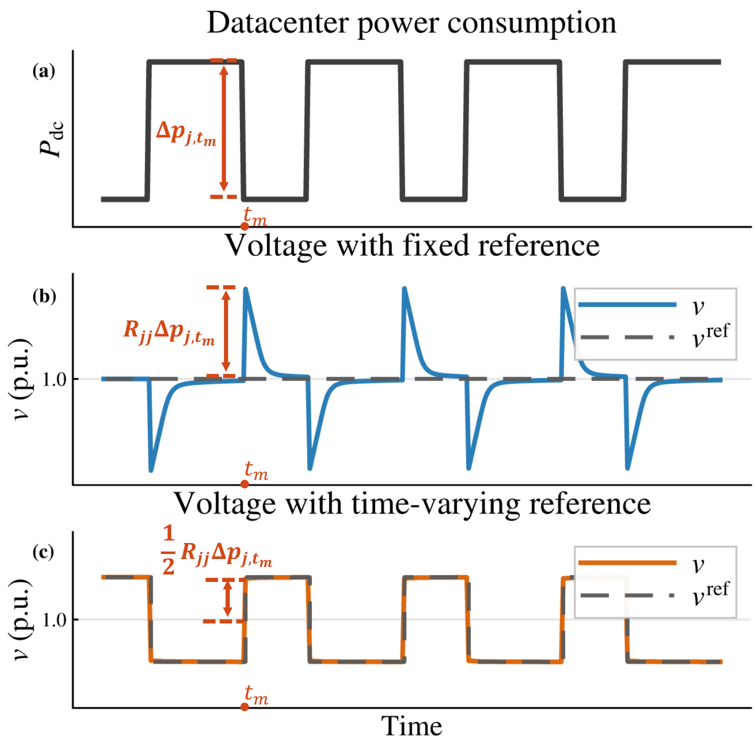

Figure 4.2: Voltage response to step-like data center power transitions. (a) Active power profile at the data center bus, where the power switches between two operating levels and exhibits a step change of $\Delta p_{j,t_m}$ at each transition. (b) Voltage trajectory under decentralized droop control with a fixed reference $v^{\text{ref}} \equiv 1$. Each transition induces a voltage shift approximately $R_{jj}\Delta p_{j,t_m}$. (c) Voltage trajectory with a time-varying reference aligned with the data center mode switching. By shifting the reference, the worst post-transition voltage deviation is reduced to approximately $\frac{1}{2}R_{jj}\Delta p_{j,t_m}$.

Control Law with Time-varying Voltage References

Our proposed controller retains the form

$$\boldsymbol{q}_{t+1} = \boldsymbol{q}_t - \mathbf{K}\big(\boldsymbol{v}_t - \boldsymbol{v}_t^{\text{ref}}\big), \tag{4.8}$$

with the key difference from (4.7) being that the reference used to compute the feedback error term, $\boldsymbol{v}_t^{\text{ref}}$, is now time-varying. To see why this reference should be set dynamically rather than be held constant, we note that the voltage change between two neighboring time steps satisfies

$$\boldsymbol{v}_{t+1} = \boldsymbol{v}_t + \mathbf{R}\,\Delta\boldsymbol{p}_t + \mathbf{X}\,\Delta\boldsymbol{q}_t, \tag{4.9}$$

where $\Delta\boldsymbol{p}_t := \boldsymbol{p}_{t+1} - \boldsymbol{p}_t$ and $\Delta\boldsymbol{q}_t := \boldsymbol{q}_{t+1} - \boldsymbol{q}_t$.

Motivated by this observation, we assign different voltage reference levels for the two data center operating modes so that the reference effectively shifts between two levels to preempt the large load changes. As shown in Fig. 4.2(c), this time-varying reference effectively shifts between two levels, allowing part of the voltage change induced by the power transition to be absorbed by the reference shift itself. For example, right after a transition, the peak voltage deviation is approximately $|R_{jj}\Delta p_{j,t_m}|$. However, choosing the reference as $v^{\text{ref}}_{j,t_m} = 1 - \frac{1}{2}R_{jj}\Delta p_{j,t_m}$ yields a post-transition voltage $v_{j,t_m+1} \approx 1 + \frac{1}{2}R_{jj}\Delta p_{j,t_m}$, reducing the peak deviation to roughly half that under fixed-reference droop control.

In particular, we implement this time-varying reference as a switching reference that alternates between two levels. More formally, we parameterize the reference as

$$v^{\text{ref}}_{i,t} = 1 + s_{i,t} v^{\text{amp}}_i, \tag{4.10}$$

where $s_{i,t} \in \{-1, +1\}$ selects one of the two reference levels, and $v^{\text{amp}}_i \in \mathbb{R}_{\geq 0}$ determines their magnitude. We will later show how to set $s_{i,t}$ such that the reference of each bus i switches between $1 \pm v^{\text{amp}}_i$ depending on the data center mode m_t.

In practice, voltage fluctuations exist within each operating mode and may have different magnitudes (e.g., voltage variations during the compute phase are larger than those during the communication phase). To accommodate this asymmetry, we introduce a bias term that allows the midpoint of the two reference levels to shift. Accordingly, we generalize (4.10) to

$$v^{\text{ref}}_{i,t} = v^{\text{bias}}_i + s_{i,t} v^{\text{amp}}_i, \tag{4.11}$$

where $v^{\text{bias}}_i \in \mathbb{R}$ is the bias of voltage reference at bus i.

Under this parameterization, the two reference levels for each bus i are centered at v^{bias}_i and separated by $2v^{\text{amp}}_i$. The remaining question is how these reference levels should be chosen so that the steady-state voltage shift induced by active power variation is incorporated into the reference itself.

Switching Law and Closed-Loop Voltage Deviation Analysis

We now analyze the closed-loop reference-tracking error, and show how the references should be chosen to reduce the worst-case voltage deviation from its nominal value.

Substituting (4.8) into the LinDistFlow model in (4.3), the voltage deviation relative to the reference evolves as

$$\boldsymbol{v}_{t+1} - \boldsymbol{v}_{t+1}^{\mathrm{ref}} = (\mathbf{I} - \mathbf{X}\mathbf{K})\big(\boldsymbol{v}_t - \boldsymbol{v}_t^{\mathrm{ref}}\big) + \boldsymbol{d}_t, \tag{4.12}$$

where $\boldsymbol{d}_t := \mathbf{R}\Delta\boldsymbol{p}_t - \Delta\boldsymbol{v}_t^{\mathrm{ref}}$, $\Delta\boldsymbol{p}_t := \boldsymbol{p}_{t+1} - \boldsymbol{p}_t$, and $\Delta\boldsymbol{v}_t^{\mathrm{ref}} := \boldsymbol{v}_{t+1}^{\mathrm{ref}} - \boldsymbol{v}_t^{\mathrm{ref}}$. The detailed derivation is provided in [161].

The reference-tracking error $\boldsymbol{v}_{t+1} - \boldsymbol{v}_{t+1}^{\mathrm{ref}}$ depends on both terms on the right-hand side in (4.12). Therefore, we design $\mathbf{K}$ to ensure closed-loop contraction and bound the size of the disturbance $\boldsymbol{d}_t$. We have the following theorem[2]:

Theorem 5 (Convergence of voltage deviation)**.** *If the diagonal gain matrix* $\mathbf{K}$ *satisfies* $0 \prec \mathbf{K} \prec 2\mathbf{X}^{-1}$, *then there exist a nonnegative matrix* $\mathbf{C}$ *and a constant* $0 < \epsilon < 1$ *such that for any initial time* t_0 *and all* $t \geq t_0$,

$$|\boldsymbol{v}_t - \boldsymbol{v}_t^{\mathrm{ref}}| \leq \mathbf{C}\,\epsilon^{t-t_0}|\boldsymbol{v}_{t_0} - \boldsymbol{v}_{t_0}^{\mathrm{ref}}| + \sum_{\tau=t_0}^{t-1}\mathbf{C}\,\epsilon^{t-1-\tau}|\boldsymbol{d}_\tau|. \tag{4.13}$$

In particular, for any $\bar{\boldsymbol{d}}$ *where* $|\boldsymbol{d}_t| \leq \bar{\boldsymbol{d}}$ *for all* t,

$$\limsup_{t\to\infty} |\boldsymbol{v}_t - \boldsymbol{v}_t^{\mathrm{ref}}| \leq \frac{1}{1-\epsilon}\mathbf{C}\,\bar{\boldsymbol{d}}. \tag{4.14}$$

Theorem 5 shows that the voltage deviation $|\boldsymbol{v}_t - \boldsymbol{v}_t^{\mathrm{ref}}|$ is driven by the disturbance $\boldsymbol{d}_t = \mathbf{R}\Delta\boldsymbol{p}_t - \Delta\boldsymbol{v}_t^{\mathrm{ref}}$. For conventional fixed-reference droop control, $\boldsymbol{d}_t = \mathbf{R}\Delta\boldsymbol{p}_t$ and therefore the large fluctuation of load (i.e., large $\Delta\boldsymbol{p}_t$) results in large voltage deviations. The proof of Theorem 5 is presented in [161].

Next, we show how designing switching references can reduce the magnitude of $\boldsymbol{d}_t$, thereby limiting the steady-state voltage deviation. The active-power increment at the data center can be decomposed as $\Delta p_{j,t} = (\bar{p}_j(m_{t+1}) - \bar{p}_j(m_t)) + (w_{t+1} - w_t)$, where the first term represents the structured shift of power caused by mode transitions of data center loads, and the second term captures the intra-mode fluctuations. To cancel the voltage

[2] Here $|\cdot|$ denotes element-wise absolute value. Vector and matrix inequalities are interpreted element-wise, and $\limsup$ applied to vectors is also understood element-wise.

deviation induced by mode transitions, we design the reference to depend only on the data center mode m_t rather than directly on time t,

$$\boldsymbol{v}_t^{\mathrm{ref}} = \bar{\boldsymbol{v}}^{\mathrm{ref}}(m_t), \tag{4.15}$$

where $\bar{\boldsymbol{v}}^{\mathrm{ref}}(0)$ and $\bar{\boldsymbol{v}}^{\mathrm{ref}}(1)$ denote the reference voltages associated with the two data center modes. From (4.3), the voltage change induced by a data center power step is $\Delta\boldsymbol{v}_{\mathrm{mode}} := \mathbf{R}\big(\bar{\boldsymbol{p}}(0) - \bar{\boldsymbol{p}}(1)\big)$. The next proposition shows that a bias–shift parameterization of the two mode-dependent reference levels cancels the disturbance induced by mode switching, corresponding to (4.11) with $v^{\mathrm{bias}} = \boldsymbol{b}$ and $v^{\mathrm{amp}} = \frac{1}{2}\Delta\boldsymbol{v}_{\mathrm{mode}}$.

Proposition 6 (Voltage reference switching)**.** *Let the reference design follow* (4.15) *and* $\Delta\boldsymbol{v}_{\mathrm{mode}} := \mathbf{R}\big(\bar{\boldsymbol{p}}(0)-\bar{\boldsymbol{p}}(1)\big)$. *If, for some bias vector* $\boldsymbol{b} \in \mathbb{R}^n$*,the two mode-dependent reference levels are chosen as*

$$\bar{\boldsymbol{v}}^{\mathrm{ref}}(0) = \boldsymbol{b} + \frac{1}{2}\Delta\boldsymbol{v}_{\mathrm{mode}}, \quad \bar{\boldsymbol{v}}^{\mathrm{ref}}(1) = \boldsymbol{b} - \frac{1}{2}\Delta\boldsymbol{v}_{\mathrm{mode}}, \tag{4.16}$$

then the disturbance term $\boldsymbol{d}_t$ *in* (4.12) *reduces to* $\boldsymbol{d}_t = \boldsymbol{r}_j(w_{t+1} - w_t)$, *where* $\boldsymbol{r}_j$ *denotes the* j*th column of* $\mathbf{R}$.

Proposition 6 is proven in [161] and shows that a proper choice of the references can reduce the disturbance magnitude $|\boldsymbol{d}_t|$, and consequently the voltage deviation $|\boldsymbol{v}_t - \boldsymbol{v}_t^{\mathrm{ref}}|$. Our ultimate objective, however, is to reduce the worst-case steady-state voltage deviation from the nominal value of 1 p.u. We next characterize how shifting the reference bias can achieve this objective.

In particular, under the reference design (4.15)–(4.16), the steady-state extrema $\boldsymbol{v}^+(\boldsymbol{b})$ and $\boldsymbol{v}^-(\boldsymbol{b})$ induced by the current bias $\boldsymbol{b}$ provide the information needed to correct the bias and reduce the worst-case steady-state deviation. This observation is formalized in the following proposition.

Proposition 7 (Bias shifting)**.** *Under the reference design* (4.15)–(4.16)*, define for each mode* $m \in \{0,1\}$ *the steady-state voltage extrema* $\boldsymbol{v}^+(m,\boldsymbol{b}) := \limsup_{t\to\infty} \boldsymbol{v}_t$ *and* $\boldsymbol{v}^-(m,\boldsymbol{b}) := \liminf_{t\to\infty} \boldsymbol{v}_t$ *along trajectories with fixed mode* $m_t \equiv m$*. The corresponding component-wise* *extrema across two modes are*

$$\boldsymbol{v}^{+}(\boldsymbol{b}) := \max\{\boldsymbol{v}^{+}(0,\boldsymbol{b}),\, \boldsymbol{v}^{+}(1,\boldsymbol{b})\}, \quad \boldsymbol{v}^{-}(\boldsymbol{b}) := \min\{\boldsymbol{v}^{-}(0,\boldsymbol{b}),\, \boldsymbol{v}^{-}(1,\boldsymbol{b})\}.$$

Then the bias that minimizes the worst-case steady-state voltage deviation from the nominal value is

$$\boldsymbol{b}^{\star} = \boldsymbol{b} - \Big(\frac{\boldsymbol{v}^{+}(\boldsymbol{b}) + \boldsymbol{v}^{-}(\boldsymbol{b})}{2} - \mathbf{1}\Big). \tag{4.17}$$

The proof of Proposition 7 can be found in [161].

Measurement-Based Reference Adaptation

In practice, neither the data center mode m_t, the mode-induced voltage shift $\Delta\boldsymbol{v}_{\mathrm{mode}}$ from the mode switching, nor the optimal bias $\boldsymbol{b}^{\star}$ is directly available to the grid controller. We therefore construct a measurement-based reference using only local voltage observations.

Fortunately, both the mode-induced voltage shift $\Delta\boldsymbol{v}_{\mathrm{mode}}$ and the voltage extrema needed for bias correction can be approximated from recent voltage measurements. We therefore maintain the most recent L voltage samples for each bus i, $\mathcal{W}_{i,t} := \{v_{i,t-\ell} \mid \ell = 0, \ldots, L-1\}$.

Specifically, for each bus i, we use the reference structure characterized in Proposition 6 and 7 to construct the time-varying reference

$$v_{i,t}^{\mathrm{ref}} = v_{i,t}^{\mathrm{bias}} + s_{i,t} v_{i,t}^{\mathrm{amp}}, \tag{4.18}$$

where $v_{i,t}^{\mathrm{bias}}$, $s_{i,t}$, and $v_{i,t}^{\mathrm{amp}}$ serve as estimates of the bias b_i, mode m_t and $\frac{1}{2}\Delta v_{\mathrm{mode},i}$, respectively.

Amplitude Update Guided by Proposition 6, we set $v_{i,t}^{\mathrm{amp}}$ to $\frac{1}{2}\Delta v_{\mathrm{mode},i}$, that is, half of the step voltage change induced by a mode transition. We therefore adapt this value online using the largest per-step voltage change in the past window. Let $L_a \leq L$ denote the window length used for amplitude estimation. If L_a is sufficiently long to include a mode transition, the window $\mathcal{W}_{i,t}$ contains at least one recent transition, and the largest one-step voltage change within this window $\max_{0\leq\ell\leq L_a-2} |v_{i,t-\ell} - v_{i,t-\ell-1}|$ provides an estimate of the mode-induced voltage shift $\Delta v_{\mathrm{mode},i}$. The amplitude is therefore updated as

$$v_{i,t}^{\mathrm{amp}} = \tfrac{1}{2} \max_{0 \le \ell \le L_a - 2} \left| v_{i,t-\ell} - v_{i,t-\ell-1} \right|. \tag{4.19}$$

Bias Update Proposition 7 shows that the optimal bias can be determined from the steady-state voltage extrema using (4.17). Since these extrema are not directly observable, we approximate them from recent voltage measurements over a horizon $L_b \le L$. Let $\mathcal{W}_{i,t}$ denote the corresponding measurement window. If L_b is sufficiently long relative to the typical interval between mode transitions, the window $\mathcal{W}_{i,t}$ captures the voltage range produced by the two operating modes. The midpoint of this range therefore provides an estimate of the bias correction term $\Delta \boldsymbol{b}$ in Proposition 7

$$\Delta b_{i,t} = \tfrac{1}{2} \Big(\max_{0 \le \ell \le L_b - 1} v_{i,t-\ell} + \min_{0 \le \ell \le L_b - 1} v_{i,t-\ell} \Big) - 1.$$

The bias is then updated incrementally as

$$v_{i,t+1}^{\mathrm{bias}} = v_{i,t}^{\mathrm{bias}} - \eta_b \, \Delta b_{i,t}, \tag{4.20}$$

where $\eta_b \in (0, 1]$ is a smoothing gain. Equation (4.17) corresponds to a full correction, i.e., $\eta_b = 1$. Here we instead use a smaller gain to reduce the influence of outdated extrema that may remain in the finite window until replaced by newer samples. A smaller η_b helps prevent the bias update from overreacting to such stale extrema.

Sign Selection The sign $s_{i,t} \in \{-1, +1\}$ acts as a local estimate of the unknown data center mode m_t, enabling a decentralized approximation of (4.15) using only local voltage measurements. We choose:

$$s_{i,t} = \begin{cases} +1, & v_{i,t} > v_{i,t}^{\mathrm{bias}}, \\ -1, & \text{otherwise}. \end{cases} \tag{4.21}$$

Because the bias approximates the midpoint between the two operating levels and intra-phase fluctuations are typically much smaller than the level separation, the voltage remains consistently above or below the bias within each phase, providing a simple decentralized rule for reference selection.

The above construction does not depend on a particular workload, model scale, or switching period. Different AI training jobs or internal data center controls may lead to different phase durations and transition magnitudes. Since these changing power patterns are reflected in the most recent voltage measurements, the measurement-based implementation enables the controller to adapt effectively to variations in the power profile. As will be demonstrated in the case study, the proposed controller can accommodate disturbances from multiple data centers, as well as changes in the data center load profile induced by internal power-smoothing control.

4.2.2 Experiments

Simulation Setup

We evaluate the proposed method on the IEEE 33-bus distribution feeder [162]. The base unit for power and voltage is 10 MVA and 12.66 kV, respectively. The system is simulated with time step $\Delta t_{\mathrm{sim}} = 0.1$ s, at which voltage states and active power injections are updated. Reactive power control operates on a slower time scale, with updates applied every $\Delta t_{\mathrm{ctrl}} = 1.0$ s, while reactive injections are held constant between control instants, consistent with practical inverter-based regulation [167].

For the proposed switching-reference method, both the bias-update and the amplitude-estimation window are set to 120 s. We compare it with the fixed-reference droop controller in (4.7), which employs a fixed voltage reference of 1 p.u.. For both methods, we adopt a symmetric deadband of 0.02 p.u. and per-update reactive power saturation of 0.01 p.u., consistent with IEEE 1547 [149]. Moreover, the linear droop gain $\mathbf{K}$ is identical for both methods.

Simulations use an approximately 45-minute power trace collected during synchronous distributed training of the LLaMA-2-70B model [168] with a QLoRA configuration [169] on a 4×H200 GPU server. We use this trace as a representative AI-training power profile to capture the temporal structure induced by synchronized training. Such a workflow produces repeatable compute/communication phase alignment and periodic patterns in the aggregate power trace. Scaling the number of accelerators primarily changes the magnitude of the

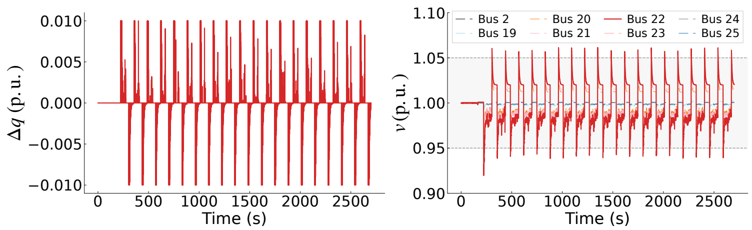

(a) Fixed-reference control

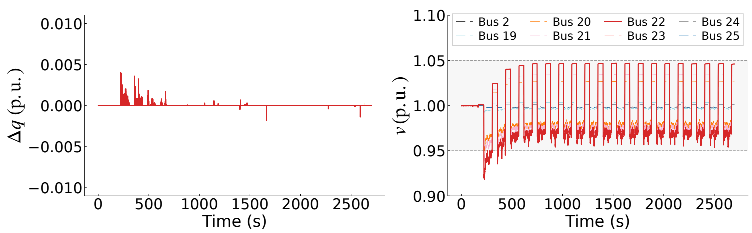

(b) Switching-reference control

Figure 4.3: Single data center scenario: comparison of incremental reactive adjustment $\Delta \boldsymbol{q}$ (p.u.) and voltage trajectory $\boldsymbol{v}$ (p.u.) under fixed-reference and switching-reference control. The proposed switching-reference control substantially reduces reactive power control effort while maintaining voltage within the admissible $\pm 5\%$ voltage-deviation band.

aggregate load, while largely preserving its underlying temporal structure.

Decentralized Grid-Side Control

Fig. 4.3 compares the conventional fixed-reference droop control with the proposed switching-reference control when a data center is connected to bus 22. Under fixed-reference control, control actions repeatedly reach the per-update limit, indicating sustained saturation of the control action. The resulting voltage trajectory exhibits persistent oscillations synchronized with the load fluctuations. In contrast, the switching-reference control substantially reduces both the magnitude and persistence of control actions after a short transient. The voltage remains consistently within the admissible $\pm 5\%$ voltage-deviation band and closer to the nominal value.

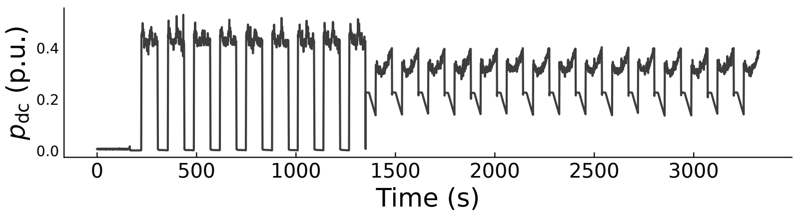

(a) Data center power profile

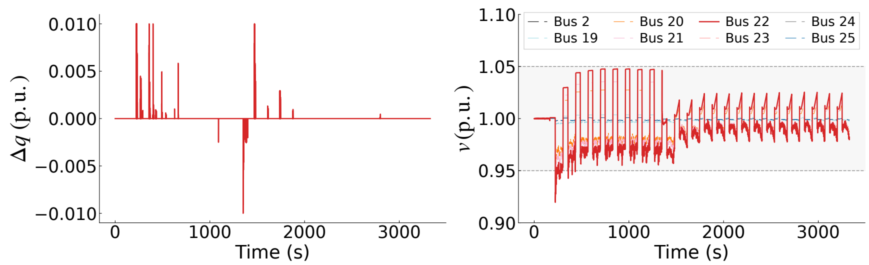

(b) Switching-reference control

Figure 4.4: Decentralized grid-side voltage control with internal data center load smoothing. (a) Data center power profile with internal load smoothing starting at $t \approx 1400$s. (b) Reactive control actions $\Delta \boldsymbol{q}$ (p.u.) and voltage trajectories $\boldsymbol{v}$ (p.u.) under the proposed switching-reference controller.

Compatibility with Data Center Load Smoothing

To examine compatibility with potential load-smoothing control in the data center, we simulate a case in which the data center initially operates without internal power regulation and activates internal load smoothing at $t \approx 1400$ s. The power trace therefore allows us to assess whether the proposed controller can promptly adapt to changes in the power profile induced by internal load-smoothing control.

Specifically, the data center is modeled as a single aggregated load equipped with an energy storage system, where the storage is sized to supply the maximum data center power for 60s, corresponding to the time required for backup generators to come online. The storage dispatch changes the net active power injection at the data center bus as $\hat{p}_{j,t} = p_{j,t} - z_{j,t}$, where $z_{j,t}$ denotes the storage power dispatched to offset deviations of the instantaneous data center load from its recent average. Its dispatch follows a simple droop-type rule: $z_{j,t} = k_p\big(p_{j,t} - \frac{1}{T}\sum_{\tau=t-T}^{t} p_{j,\tau}\big)$, with droop gain $k_p = 0.6$ and averaging horizon $T = 100$s. To preserve backup capability, the storage is restricted to operate within a narrow state-of-charge band around 95% (0.93–0.97), and the dispatch is forced to zero whenever these energy limits are reached.

The data center power profile is shown in Fig. 4.4(a), where the power trace initially follows the original switching pattern of the training workload and exhibits a reduced magnitude of power transitions after internal regulation begins. The corresponding voltage trajectory and control actions are shown in Fig. 4.4(b). After internal load smoothing is activated, the reduced power fluctuations lead to smaller voltage deviations. These results illustrate that the proposed controller remains effective without requiring extensive coordination when the load switching pattern is altered by load smoothing in the data center.

4.3 Regulation of Data Center Loads Using Uninterruptible Power Supply Energy Storage

While Section 4.2 addresses grid-side reactive power compensation, we can also leverage active power compensation using the internal energy storage capacities of data centers in the form of Uninterruptible Power Supplies (UPS). If dispatched strategically, this UPS storage can directly smooth the active power profile before the fluctuations propagate to the distribution feeder. In this section, we investigate whether the energy storage units in data center UPSs can be effectively leveraged for active voltage regulation without compromising their core backup redundancy functions.

4.3.1 Data Center System Modeling and Reliability Constraints

To leverage the data center for active power regulation, the internal topology of the facility must be considered. The facility is modeled as a set of uninterruptible power supplies (UPSs), $\mathcal{S}$, and server racks, $\mathcal{D}$. For a data center located at bus $i \in \mathcal{N}$, its UPSs are collected in the subset $\mathcal{S}_i$, and its server rack loads in the subset $\mathcal{D}_i$.

The connections between UPSs and server racks are modeled via the biadjacency matrix $\mathbf{A} \in \{0,1\}^{N_\mathcal{D} \times N_\mathcal{S}}$, where $\mathbf{A}_{j,k} = 1$ indicates that the j-th server rack can receive power from the k-th UPS. The energy stored in the k-th UPS at time t is denoted $e_{k,t}$, bounded by its maximum capacity $\bar{e}_k$. The active power output of each UPS is represented by the decision variables u_k^P.

UPS N-1 Redundancy Constraints

A fundamental operational limitation is that a data center must maintain sufficient battery-energy storage in its UPS units to power all server racks for T_{backup}, the time required for backup generators to come online, even during an N-1 UPS contingency event [170]. In each timestep t, rack j demands $D_{j,\tau}$ power. The backup energy demand for rack j over this horizon is defined as $d_{j,t} = \sum_{\tau=t}^{t+T_{\text{backup}}} D_{j,\tau} \cdot \Delta t$.

To ensure adequate energy reserves in any N-1 contingency, generalized routing constraints introduce an energy allocation matrix $\boldsymbol{E}_t \in \mathbb{R}^{N_\mathcal{D} \times N_\mathcal{S}}$, where element $e_{k \to j,t}$ represents the energy allocated from UPS k to rack j. For each failed UPS scenario $f \in \mathcal{S}$, the weighted allocations $\hat{\boldsymbol{E}}_t^f = \boldsymbol{A}^f \odot \boldsymbol{E}_t$ (where $\boldsymbol{A}^f$ is the biadjacency matrix with column f zeroed out [171]) must simultaneously satisfy the rack demands and respect the surviving UPS capacities:

$$\sum_{k \in \mathcal{S}} \hat{e}^f_{k \to j,t} = d_{j,t} \qquad \forall j \in \mathcal{D}, \forall f \in \mathcal{S} \tag{4.22a}$$

$$\sum_{j \in \mathcal{D}} \hat{e}^f_{k \to j,t} \leq e_{k,t} \qquad \forall k \in \mathcal{S}, \forall f \in \mathcal{S}. \tag{4.22b}$$

While this generalized formulation accommodates arbitrary routing topologies, data centers typically employ standardized physical architectures that eliminate the need for the dense $N_{\mathcal{D}} \times N_{\mathcal{S}}$ allocation variables entirely.

Typical Data Center Layouts

Centralized UPS: If all server racks are supported by a common power bus fed by a centralized pool of UPS units, the biadjacency matrix is an all-ones matrix ($\boldsymbol{A} = \mathbf{1}$). The point-to-point allocation variables vanish, and the N-1 redundancy requirement simplifies to $N_{\mathcal{S}}$ aggregate bounds ensuring the surviving fleet can meet the total facility demand:

$$\sum_{i \in \mathcal{S} \setminus \{k\}} e_{i,t} \geq \sum_{j \in \mathcal{D}} d_{j,t} \quad \forall\, k \in \mathcal{S}. \tag{4.23}$$

Dedicated N+1 (Swing UPS): A typical N+1 backup layout is shown in Fig. 4.5. If each rack is supported by a dedicated primary UPS with a single standby "swing" UPS s acting as a universal backup, the biadjacency matrix becomes $\boldsymbol{A} = [\mathbf{I} \mid \mathbf{1}]$. This structure decouples the contingencies into two strict bounds: every primary UPS must independently support its own rack, and the swing UPS must possess enough energy to cover the maximum demand of any single rack in the facility:

$$e_{k,t} \geq d_{k,t} \quad \forall k \in \mathcal{S} \tag{4.24a}$$

$$e_{s,t} \geq d_{j,t} \quad \forall j \in \mathcal{D}. \tag{4.24b}$$

Leveraging these structurally simplified N-1 constraints, the computational burden of solving a storage dispatch problem with energy storage limits can be significantly reduced.

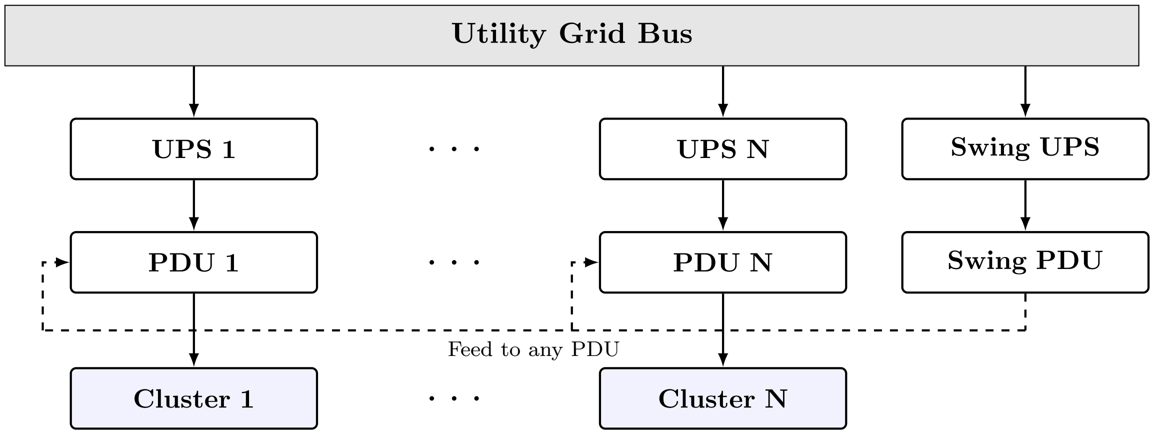

Figure 4.5: Diagram illustrates the N+1 UPS Layout. There are N primary power trains (Utility → UPS → PDU → Cluster), and one single swing UPS and swing PDU pair acting as an isolated redundant backup. In the event of a primary UPS failure, its dedicated PDU can switch to receive power from the backup UPS via the swing PDU, as shown by the dashed lines.

4.3.2 Ideal MPC Storage Dispatch Problem

Subject to these reliability constraints assuming an N+1 UPS layout, we formulate the baseline, deterministic voltage regulation optimization problem as:

$$
\begin{aligned}
\min_{\boldsymbol{u}^P, \boldsymbol{u}^Q} \quad & \sum_{i \in \mathcal{N}} \sum_{t \in T} \frac{1}{2} (v_{i,t} - v_{\text{nom}})^2 && && (4.25a) \\
\text{s.t.} \quad & p_{i,t} = P^L_{i,t} - \sum_{k \in \mathcal{S}_i} u^P_{k,t} + \sum_{j \in \mathcal{D}_i} D_{j,t} && \forall\, i \in \mathcal{N}, \forall\, t \in T && (4.25b) \\
& q_{i,t} = Q^L_{i,t} + \sum_{k \in \mathcal{S}_i} u^Q_{k,t} && \forall\, i \in \mathcal{N}, \forall\, t \in T && (4.25c) \\
& \boldsymbol{v}_t = \mathbf{R} \boldsymbol{p}_t + \mathbf{X} \boldsymbol{q}_t + \mathbf{1} && \forall\, t \in T && (4.25d) \\
& u^P_k \in [\underline{P}_k, \bar{P}_k] && \forall\, k \in \mathcal{S}_i && (4.25e) \\
& d_{k,t} \leq e_{k,t} \leq \bar{e}_k && \forall\, k \in \mathcal{S}_i \setminus \{s\}, \forall\, t \in T && (4.25f) \\
& d_{j,t} \leq e_{s,t} \leq \bar{e}_s && \forall\, j \in \mathcal{D}_i, \forall\, t \in T && (4.25g) \\
& \boldsymbol{e}_{t+1} = \boldsymbol{e}_t - \boldsymbol{u}^P_t \cdot \Delta t && \forall\, t \in T && (4.25h)
\end{aligned}
$$

where here $P^L_{i,t}$ is the background load at bus i from loads outside the data center.

Solving (4.25) globally over a long horizon is computationally intractable for fast control, and requires knowledge of the distribution system's parameters, loads at other buses, and the future demand of the data center racks.

4.3.3 Predictive Control Modeling Considerations

To transition from the idealized deterministic problem (4.25) to an implementable real-world controller, the Model Predictive Control (MPC) algorithm requires a finite-horizon forecast of the data center's future load.

Because the true future active power demand of the server racks at future time step $\tau > t$, denoted as $D_{j,\tau}$, is unknown at the current control time step t, the controller must rely on a predictive forecasting model, $\mathcal{F}$. This predictor utilizes a historical observation window of the realized load to generate a sequence of future load predictions over the horizon H.

The prediction sequence for a given rack over the horizon is thus defined as:

$$[\tilde{D}_{j,t}, \tilde{D}_{j,t+1}, \dots, \tilde{D}_{j,t+H-1}] = \mathcal{F}(D_{j,t-W}, \dots, D_{j,t}) \tag{4.26}$$

where W is the historical look-back window utilized by the predictor.

Consequently, any state or constraint that depends on the future load must also be treated as a predicted variable. Most critically, the dynamic N-1 backup energy requirement integrated over the backup horizon T_{backup} transforms from a known deterministic bound into a predicted requirement. For a given prediction step τ, this is:

$$\tilde{d}_{j,\tau} = \sum_{m=\tau}^{\tau+T_{\text{backup}}} \tilde{D}_{j,m} \cdot \Delta t. \tag{4.27}$$

The accuracy of the controller is intrinsically linked to the prediction error, defined as the difference between the realized value and the forecasted value (e.g., $D_{j,t+\tau} - \tilde{D}_{j,\tau}$).

However, predicting LLM training workloads presents unique forecasting and operational challenges. Unlike traditional aggregate residential or commercial loads, the raw active power demand of a synchronized LLM training cluster exhibits a highly volatile, distinctly bimodal distribution. It jumps abruptly between low-draw communication phases and maximum-draw compute phases. This unmitigated bimodal base load behavior is illustrated in Fig. 4.6(a).

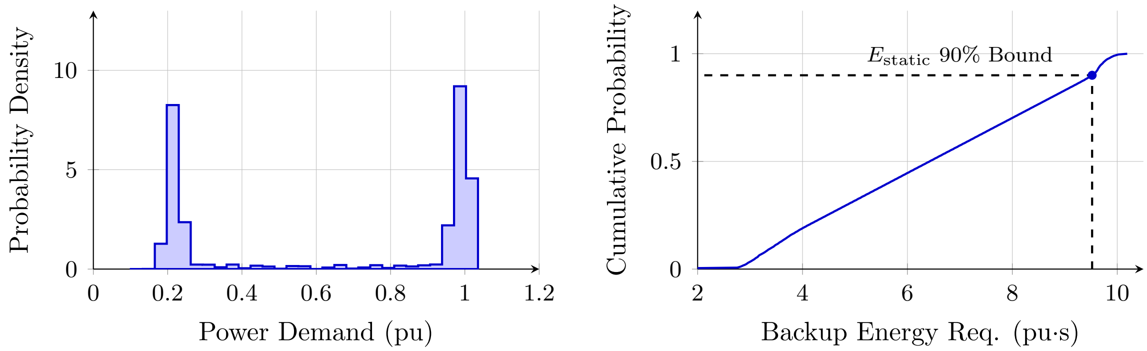

(a) Instantaneous Power Demand PDF

(b) True Backup Energy CDF

Figure 4.6: Empirical characterization of the raw LLM training workload. Because the load is highly bimodal (a), standard Gaussian assumptions fail. Consequently, extracting a robust 99^{th} percentile limit from the historical energy requirement (b) yields a large static lower energy bound for the UPS storage.

To successfully regulate this extreme profile, the predictive control architecture must be carefully augmented to address four distinct challenges: predicting required energy bounds, step-timing misalignments, finite-horizon trajectory inconsistencies, and integrated energy prediction uncertainties.

Static Energy Capacity Bounds

Because the bimodal LLM distribution is not Gaussian, standard statistical assumptions regarding load volatility are invalid. If a grid operator employs a standalone, reactive controller without predictive foresight, they must size the N-1 emergency reserve against a static, absolute worst-case scenario.

By evaluating the cumulative distribution function (CDF) of the historical backup energy requirement (Fig. 4.6(b)), a robust percentile-based lower bound can be extracted. However, utilizing this static bound locks away a significant percentage of the battery capacity that would otherwise be available when the future energy requirements are lower, leaving it unavailable for active power regulation.

Load Step Timing Errors

To utilize this additional capacity, the controller can use a time-series predictor to dynamically forecast the baseline requirement. However, directly applying MPC to these predictions introduces a vulnerability due to the sharp, bimodal step-changes of the workload. As illustrated in the stylized power demand diagram in Fig. 4.7, even if a predictor perfectly estimates the magnitude of an impending compute spike, it will inevitably lead or lag the physical event by some amount of time. This misalignment, an error in predicting the

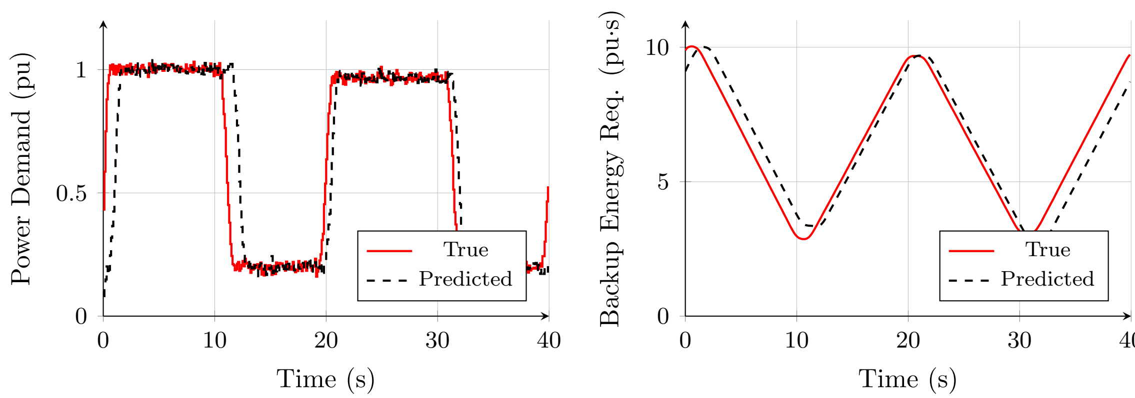

(a) Instantaneous power horizon.

(b) Integrated backup energy.

Figure 4.7: Stylized illustration of a prediction phase-shift error. A slight timing lag in anticipating an abrupt physical load step (a) immediately propagates into a persistent mismatch in the integrated backup energy requirement (b).

timing of the load step, severely degrades the performance of a predictive controller. If the MPC expects the load to drop, but the physical servers remain in a high-compute phase for an extra second, the controller may start to charge the UPS early, exacerbating power profile steps and grid voltage deviations.

A single-layer MPC operating at a coarse predictive timestep (e.g., 1 to 10 seconds) cannot physically react quickly enough to correct these instantaneous step timing errors. However, attempting to execute a predictive optimization at a sub-second resolution to address these step changes is computationally intractable. Therefore, to safely control these volatile workloads, the control architecture must be decoupled into multiple levels. An outer predictive layer is used to plan power and energy trajectories over a long horizon, while a faster inner control loop is necessary to track the physical load and quickly reject the timing errors of the predictor.

Finite Horizon Trajectory Anchoring

In addition to step timing misalignments, the periodic nature of LLM workloads makes the average value of the load vary across finite computational horizons if the prediction window does not align with the load period and is insufficiently long to capture the true average of the periodic workload. As the receding horizon slides forward in time, the local average load captured within the finite window fluctuates, causing the unanchored controller to repeatedly step the power trajectory to align with this shifting local mean. Since our controllers performance objectives are to regulate the power demand to reduce its ramp rate and variance, these steps in the predicted average value degrade the performance of the control when evaluated over the full simulation period.

To resolve this without increasing the computational burden of an extended prediction window, we include an anchor cost penalty to the MPC formulation objective. By explicitly penalizing large step deviations from the immediate prior control action, and simultaneously evaluating both the recent historical past and the predicted future against a shared, optimized mean variable, the controller can tie its discrete optimization windows to the previously established power trajectory.

Predicted Required Energy Buffering

To utilize the predictive foresight necessary to leverage additional battery capacity while remaining robust against errors in the predicted backup storage energy, we propose adding a buffer on the energy constraint.

While the instantaneous timing errors cause sharp power spikes, the N-1 contingency constraint is fundamentally an energy requirement integrated over the backup horizon, T_{backup}. Fig. 4.8 evaluates the statistical distributions of the prediction errors. Because timing errors are auto-correlated, they can compound during integration over the control and planning windows.

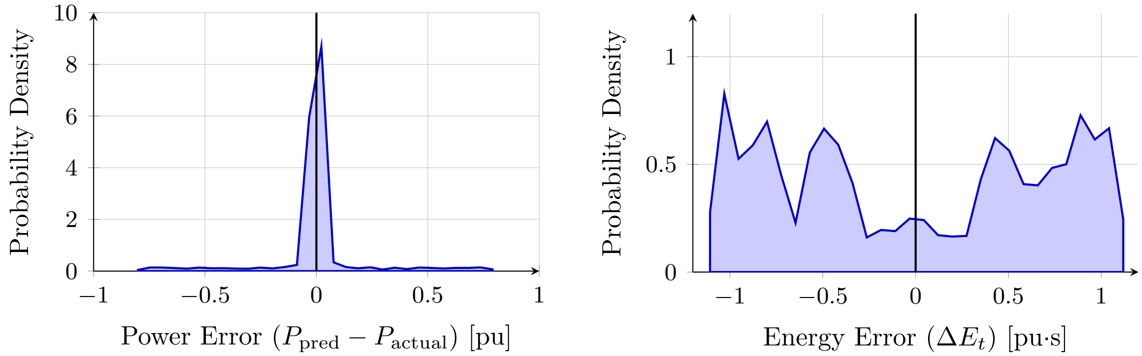

(a) Power prediction error PDF.

(b) Energy prediction error PDF (ΔE_t).

Figure 4.8: Stylized illustration of potential errors in the predicted load. (a) The PDF of the instantaneous power error highlighting the small long tails caused by step timing errors. (b) The PDF of the absolute integrated energy prediction error (ΔE_t) over the backup horizon.

Furthermore, the directionality of the prediction error dictates different operational risks, requiring asymmetric buffers. Underestimating the load creates an accumulated energy shortage, risking violation of the N-1 contingency minimums. Conversely, overestimating the load creates an accumulated energy surplus, risking battery overcharge or clipping if the controller leaves insufficient headroom.

Fig. 4.9 demonstrates the resulting cumulative distribution functions (CDFs) of these integrated energy errors. By extracting the 99^{th} percentiles, we define two distinct buffers: a lower safety buffer (B^{lower}) derived from the accumulated shortages over the control and planning horizons, and an upper headroom buffer (B^{upper}) derived from the accumulated surplus resulting from synchronous updates.

Note that when a single UPS supports multiple server racks, the aggregate prediction error inherently depends on the correlation between the individual rack workloads. If the racks process independent computational tasks, their prediction errors partially cancel each other out when aggregated, effectively reducing the total required safety buffer. In our

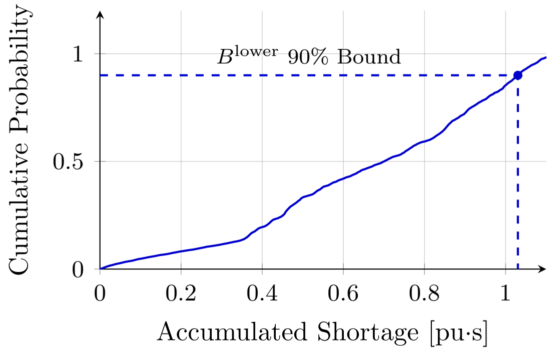

(a) Lower buffer requirement (Shortage).

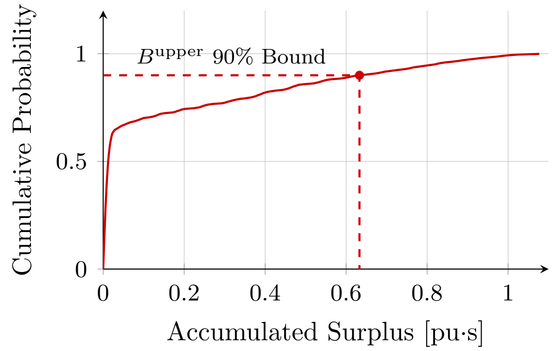

(b) Upper buffer requirement (Surplus).

Figure 4.9: Statistical evaluation of the asymmetric energy buffer requirements. (a) The CDF of the accumulated shortage (underestimates) used to size the lower safety buffer (B^{lower}). (b) The CDF of the accumulated surplus (overestimates) used to size the upper headroom buffer (B^{upper}).

formulation, however, we assume the server racks operate as a unified cluster training a large language model in temporal synchrony. Because these workloads are highly correlated, the aggregated prediction error simplifies to a direct summation. Therefore, we can safely size the stochastic buffers directly from the aggregate synchronized load profile.

In instances of severe predictor overshoot, the sum of the forecasted requirement and the lower stochastic buffer might artificially exceed the absolute maximum energy the physical system could theoretically consume. To prevent the controller from locking away more capacity than the load could possibly demand, this dynamic lower limit is strictly capped by the static 99$^{\text{th}}$ percentile historical energy requirement ($E^{99\%}_{\text{static}}$) derived previously from the raw load profile. The finalized, dynamically bounded contingency constraint for any UPS k becomes:

$$e_{k,t} \geq \min\left(\tilde{E}^{\text{req}}_{k,t} + B^{\text{lower}},\ E^{99\%}_{\text{static}}\right), \tag{4.28}$$

where $\tilde{E}^{\text{req}}_{k,t}$ is the predicted energy requirement lower limit at time t for storage unit k as defined in (4.22b) or (4.24b) depending on the data center architecture. This formulation ensures a survival probability of $P(e_{k,t} \geq \tilde{E}^{\text{req}}_{k,t}) \geq 0.99$ while firmly bounding the maximum reserved capacity. It grants the controller the freedom to safely utilize the battery for load smoothing, temporarily using the B^{lower} margin to absorb any underestimation inaccuracies from the predictor, while guaranteeing that the system defaults to the safety of the static historical limit for worst-case prediction failures.

4.3.4 Hierarchical Control Architecture

To reconcile the need for long-term predictive peak-shaving with the requirement for instantaneous disturbance rejection, we propose a two-layered hierarchical control architecture that implements the solutions motivated in the previous sections. This structure decouples the optimization horizons, allowing a supervisory MPC to handle optimizing the storage dispatch over time and maintaining the energy storage constraints while a decentralized inner loop manages sub-second load volatility. This control structure is illustrated in Fig. 4.10.

The architecture is divided into two primary loops:

1. **The Outer Planning Loop (Predictive):** Operating at a coarse timestep, this layer utilizes the predictor forecasts to compute the optimal, grid-smoothing power baselines. It incorporates the asymmetric stochastic energy buffers (B^{lower} and B^{upper}) derived in Section 4.3.3 to probabilistically guarantee N-1 contingency survival and prevent overcharge saturation despite predictor uncertainties.

2. **The Inner Tracking Loop (Reactive):** Operating at a high-frequency, sub-second timestep, this layer responds to the actual current power demand. It utilizes a low-pass filter to track the physical instantaneous load, rejecting the timing misalignments of the predictor and ensuring the UPS hardware executes a smooth, safe AC power injection.

Through this hierarchical structure, the system achieves optimal long-term planning without being compromised by the load step timing errors inherent in high-volatility load predictions.

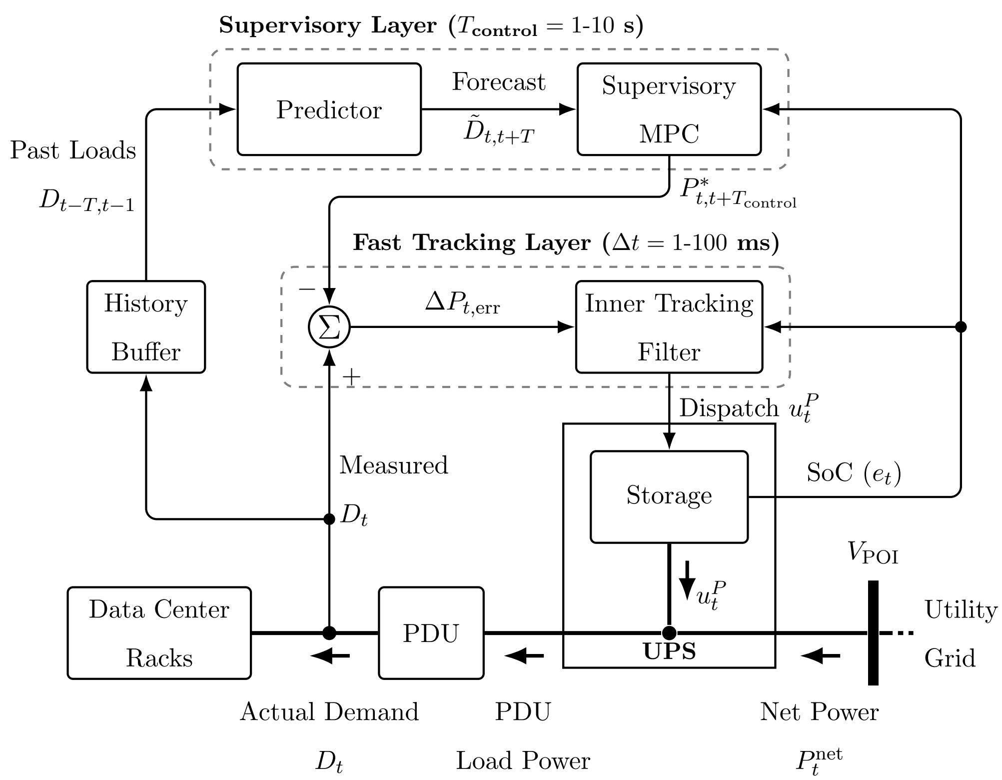

Figure 4.10: Proposed hierarchical control architecture. The supervisory MPC operates on a slow timescale ($T_{\text{control}} = 1$ s - 10 s) to establish smoothed grid targets (P^*_{nom}), while the inner tracking filter operates at the hardware timescale ($\Delta t = 1$ ms-100 ms) to reject instantaneous prediction errors.

Outer Loop: Optimal Trajectory Planning

The outer loop consists of the supervisory MPC, executing at a slower timescale T_{control} solving a finite-horizon (T) trajectory optimization problem. At each control step t, it utilizes the finite-horizon (H) predictions of the LLM training workload. The objective of this layer is to calculate an active power trajectory, P^*_{nom}, for the grid point of interconnection (POI).

To explicitly balance the competing interests of grid impact minimization and battery preservation, the supervisory MPC minimizes a multi-objective cost function J over the prediction horizon H. The optimization considers the future prediction horizon $\tau \in \{0, \ldots, T-1\}$. The total cost function is formulated as:

$$
\begin{aligned}
J(u^P, \mu, \epsilon, P_{\text{net},\tau}) = \sum_{\tau=t}^{t+T-1} \Bigg[& \underbrace{\alpha \left(P_{\text{net},\tau} - \mu\right)^2}_{\text{Power Deviations}} + \underbrace{\beta \left(P_{\text{net},\tau} - P_{\text{net},\tau+1}\right)^2}_{\text{Power Ramp Rate}} \\
& + \underbrace{\gamma \sum_{k\in\mathcal{S}} \Phi_{\text{deg}}(u^P_{k,\tau}, e_\tau)}_{\text{Battery Degradation}} + \underbrace{\eta \sum_{k\in\mathcal{S}} \|\epsilon_{k,\tau}\|^2}_{\text{Soft Constraints}} \Bigg] + J_{\text{anchor}}
\end{aligned}
\tag{4.29}
$$

The primary summation evaluates the predictive objective over the future horizon. The first term penalizes power deviations from the mean power demand (μ). Rather than fixing this mean as a rigid parameter, μ is treated as a free decision variable within the optimization. The second term penalizes step-to-step ramp rates of the power demand.

The third term introduces a battery degradation penalty, $\Phi_{\text{deg}}(u^P_{k,\tau}, e_{k,\tau}) = (u^P_{k,\tau})^2 + (e_{k,\tau} - e^*)^2$, applying a quadratic cost to the storage dispatch and deviations from the specified SoC (e^*) to discourage aggressive power draws and deep discharge cycles, and naturally promote proportional load sharing among the available UPS units. The fourth term penalizes an energy constraint slack variables ($\epsilon_k = (\epsilon_k^{\text{lower}}, \epsilon_k^{\text{upper}})$) introduced later, ensuring that the buffered energy storage limits are only violated if necessary to maintain solver feasibility during unpredicted transients.

To reduce large steps between MPC control iterations, an anchoring term is appended to the objective function:

$$J_{\text{anchor}} = \underbrace{\alpha \left(\sum_{\tau=t-T_{\text{control}}}^{t-1} (P_{\text{net},\tau} - \mu)^2 \right)}_{\text{Historical Mean Anchor}} + \underbrace{\beta \left(P_{\text{net},t} - P_{\text{net},t-1}\right)^2}_{\text{Ramp Anchor}}. \tag{4.30}$$

This term applies the power deviation penalty (α) to the fixed historical trajectory ($P_{\text{net},\tau}$ for $\tau \in \{t - T_{\text{control}}, \ldots, t-1\}$.) against the same shared free variable, μ. By minimizing the combined variance of both the historical past and the predicted future against this single decision variable, the optimization ties the future load trajectory to the recent past average value. Further, we penalize the deviation of the power in the first step from the last achieved power value to prevent large steps to occur in the desired power trajectory between the MPC finite-horizon updates.

At each timestep t, the supervisory layer solves the following constrained optimization problem over the prediction horizon $\mathcal{T} \coloneqq \{t, \ldots, t+T-1\}$:

$$\begin{aligned}
\min_{u^P,\mu,\epsilon} \quad & J(u^P, \mu, \epsilon, \tilde{P}_{\text{net}}) && && (4.31\text{a})\\
\text{s.t.} \quad & \tilde{P}_{\text{net},\tau} = \sum_{j\in\mathcal{D}_i} \tilde{D}_{j,\tau} - \sum_{k\in\mathcal{S}_i} u^P_{k,\tau} && \forall\, \tau \in \mathcal{T} && (4.31\text{b})\\
& e_{k,\tau+1} = e_{k,\tau} - u^P_{k,\tau} \cdot \Delta T_{\text{control}} && \forall\, k \in \mathcal{S}, \forall\, \tau \in \mathcal{T} && (4.31\text{c})\\
& \underline{P}_k \le u^P_{k,\tau} \le \overline{P}_k && \forall\, k \in \mathcal{S}, \forall\, \tau \in \mathcal{T} && (4.31\text{d})\\
& e_{k,\tau} \ge \left(\tilde{E}^{\text{req}}_{k,\tau} + B^{\text{lower}}\right) - \epsilon^{\text{lower}}_{k,\tau} && \forall\, k \in \mathcal{S}, \forall\, \tau \in \mathcal{T} && (4.31\text{e})\\
& e_{k,\tau} \le (\overline{e}_k - B^{\text{upper}}) + \epsilon^{\text{upper}}_{k,\tau} && \forall\, k \in \mathcal{S}, \forall\, \tau \in \mathcal{T} && (4.31\text{f})\\
& \epsilon^{\text{lower}}_k, \epsilon^{\text{upper}}_k \ge 0 && \forall\, k \in \mathcal{S}, && (4.31\text{g})
\end{aligned}$$

Equation (4.31c) enforces the discrete-time battery energy dynamics, with the initial state $e_{k,0}$ set to the currently measured battery level at time t. Equation (4.31e) enforces the dynamically calculated N-1 contingency lower bounds, protected by the lower energy buffer.

Equation (4.31f) establishes the upper State of Charge (SoC) limitation. To mitigate long-term battery degradation (such as calendar aging) caused by floating at maximum physical capacity, the operational limit $\overline{e}_k$ is purposefully set below 100%. The upper stochastic buffer (B^{upper}) is then subtracted from this ceiling to guarantee sufficient head-room to absorb instantaneous energy surpluses resulting from load overestimates. Finally, integrating the slack variables ($\epsilon_{k,\tau}^{\text{upper}}, \epsilon_{k,\tau}^{\text{lower}}$) converts these limits into soft bounds, preventing problem infeasibility during large transients. If set to be sufficiently large, the buffered energy constraints will only be violated if necessary for problem feasibility when there are significant steps in the predicted energy requirement.

The optimal trajectory, P_{nom}^*, is passed to the inner tracking layer as the reference target over the following T_{control} period.

Inner Loop: Fast Disturbance Rejection

Because the power trace predictions are inherently imperfect, the instantaneous aggregate data center demand, $\sum D_{j,t}$, will inevitably deviate from the MPC's expected profile between the slower T_{control} optimization updates. To prevent these high-frequency prediction errors from propagating to the grid, the inner tracking layer operates at a fast control rate to reject these instantaneous disturbances.

The tracking controller first calculates the total unpredicted active power residual, defined as the difference between the actual required tracking power and the sum of the MPC's planned dispatch commands, $u_{k,t}^{P*}$:

$$\Delta P_{\text{err}}(t) = \left(\sum_{j \in \mathcal{D}_i} D_{j,t} - P_{\text{nom}}^* \right) - \sum_{k \in \mathcal{S}_i} u_{k,t}^{P*}. \tag{4.32}$$

To safely absorb this disturbance, the residual error is dynamically distributed among the available UPS units proportionally to their available physical headroom. This available headroom calculation evaluates both the instantaneous active power ratings ($\underline{P}_k, \overline{P}_k$) and the remaining energy capacity relative to the lower energy bounds and upper energy limit, preventing the inner loop from attempting to draw power from saturated units. Let $\tilde{u}_{k,t}^P$ denote the intermediary assigned power target for UPS k after this energy-aware distribution.

To achieve smooth hardware actuation and consider storage ramp rates, the controller applies a discrete-time first-order low-pass filter to these intermediary targets to yield the final active power dispatch commands:

$$u_{k,t}^P = \alpha \tilde{u}_{k,t}^P + (1 - \alpha) u_{k,t-1}^P \tag{4.33}$$

where the filter coefficient is defined as $\alpha = \frac{\Delta t}{\tau_f + \Delta t}$, with τ_f being the filter time constant and Δt representing the execution timestep of the inner loop.

4.3.5 Case Study and Experimental Setup

To evaluate the effectiveness of the proposed hierarchical architecture under realistic forecasting uncertainties, we simulate a data center connected to an IEEE 33-bus radial distribution system. The data center is modeled with 3 active server racks and 4 UPS storage units to satisfy N+1 redundancy requirements. The aggregate data center load is scaled to 2.0 MW and driven by an LLM training workload profile exhibiting abrupt compute-to-communication phase transitions. The required backup period of the data center is 60 seconds.

To generate the requisite active power forecasts ($\tilde{D}_{j,t}$) utilized by the supervisory MPC, we employ a Long Short-Term Memory (LSTM) recurrent neural network as the specific predictive model for this study. The LSTM was trained offline using power traces of the LLM training workloads, enabling it to predict the compute-communication cycles of the workload.

During the simulation, the supervisory MPC horizon is set to 60 seconds using 120 second long predictions with a control update interval of 10 seconds. The costs and simulation parameters can be found in Table 4.1. We evaluate three scenarios:

1. **Base Load:** The unmitigated data center active power profile.

2. **Standard Proportional Droop:** A standalone, purely responsive linear control law that dispatches active power proportional to the instantaneous load error. Its target, P_{nom}, is a simple historical moving average, and its safety bound is derived from the absolute system load that is passed to the inner loop controller low-pass filter. To guarantee survival without foresight, this requires a static, immutable offline lower limit (e.g., 0.7816 pu) derived from the 99$^{\text{th}}$ percentile global maximum of the historical energy requirement.

Table 4.1

SIMULATION PARAMETERS AND MPC TUNING

Parameter	Value	Parameter	Value
Temporal & Tracking Horizons		*Energy Limits* (pu·s)	
Sim. Duration	52 min	Lower Limit ($E^{99\%}_{\text{static}}$)	0.7816
Backup Horizon (T_{backup})	60 s	Upper Limit ($\overline{e}_j$)	0.8750
MPC Horizon (T_{pred})	60 s	Lower Buffer (B^{lower})	0.1590
Update Interval (T_{control})	6 s	Upper Buffer (B^{upper})	0.0180
Inner Timestep (Δt)	0.1 s	*MPC Objective Weights*	
Filter Constant (τ_f)	0.5 s	Smooth (α): 10^0	Ramp (β): 10^2
SoC Reference Pnlty.	10^{-2}	Degrad. (γ): 10^{-1}	Slack (η): 10^1

3. **Hierarchical MPC:** The proposed supervisory controller utilizing the LSTM predictor and asymmetric stochastic buffering.

Stochastic Bounds and Prediction Error

Prior to evaluating the closed-loop control performance, we first analyze the empirical forecasting accuracy of the trained LSTM model on a dataset of LLM load profiles. Characterizing these distributions is a prerequisite for the simulation, as they directly dictate the sizing of the asymmetric stochastic buffers utilized by the Hierarchical MPC.

To intuitively understand the nature of forecasting uncertainty in this domain, it is instructive to first examine the LSTM's behavior in the time domain. Fig. 4.11 illustrates a single prediction window captured at $t = 550$ s during the evaluation of a LLaMA-2 70B training workload. As shown in Fig 4.11(a), the model successfully anticipates the impending transition from a low-power communication phase to a maximum-draw compute phase. However, slight misalignments due to mistiming the steps, are inevitable during these exceptionally abrupt transitions. Fig. 4.11(b) demonstrates how these instantaneous power mismatches integrate over the 60-second backup horizon, causing the predicted energy requirement to deviate from the true requirement.

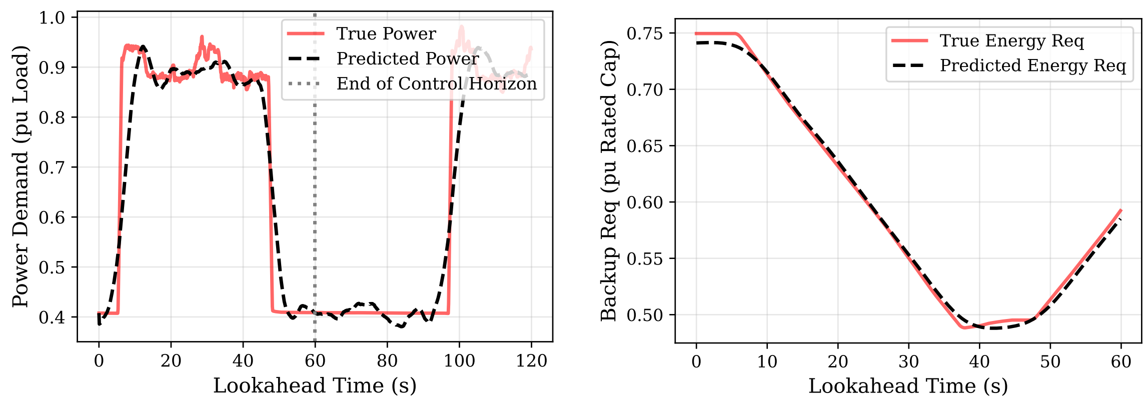

(a) Instantaneous power horizon (120s).

(b) Integrated backup energy horizon (60s).

Figure 4.11: Time-domain alignment of the LSTM predictor evaluating a LLaMA-2 70B training workload. While the predictor successfully captures the magnitude and general timing of compute spikes, minor step timing errors occur at the transition boundaries, accumulating into the integrated energy deviations shown in (b).

To guarantee operational safety, the control architecture must robustly account for these step timing deviations. Therefore, we aggregate the statistical distributions of these errors across the entire dataset. First, we characterize the physical demands of the unmitigated data center load, illustrated in Fig. 4.12.

As shown in Fig. 4.12(a), the raw active power demand exhibits the expected bimodal distribution, jumping abruptly between low-draw idle states and near-maximum capacity during synchronized compute phases. Fig. 4.12(b) presents the cumulative distribution function (CDF) of the absolute energy required to survive a worst-case T_{backup} contingency. A standalone, purely reactive controller must size its permanent static reserve against the 99^{th} percentile of this global requirement (evaluated at 0.7816 pu), permanently locking away storage capacity.

To safely use this capacity, the Hierarchical MPC utilizes the LSTM to forecast the baseline requirement. Fig. 4.13 evaluates the statistical accuracy of these forecasts across the different integration horizons.

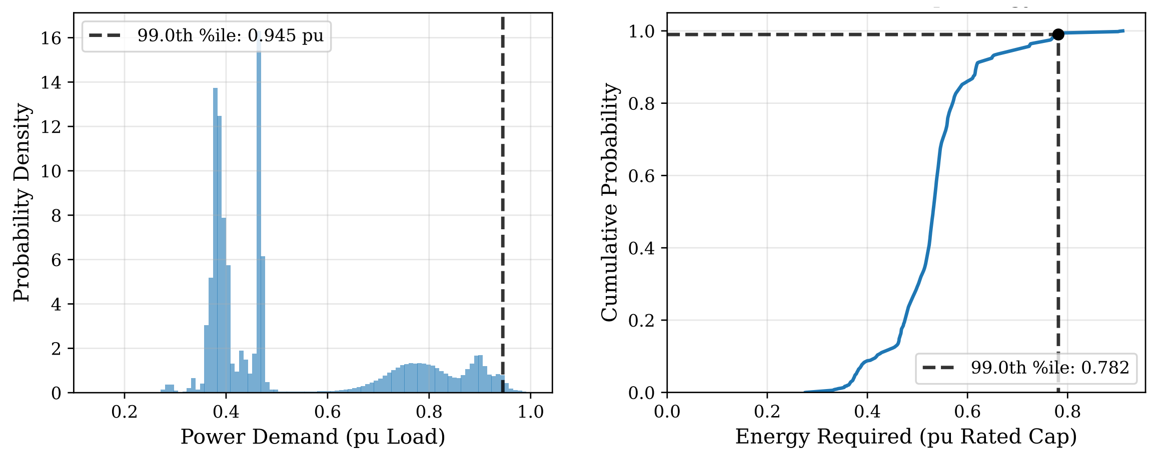

(a) Instantaneous power demand PDF.

(b) Required backup energy CDF.

Figure 4.12: Empirical characterization of 63 different LLM training workloads. (a) The bimodal nature of the instantaneous active power demand. (b) The CDF of the integrated energy required over T_{backup}, which establishes the conservative static safety limits required by standalone reactive controllers.

When the LSTM predictor is applied to the testing data, the instantaneous active power prediction error (Fig. 4.13(a)) remains tightly centered around zero, indicating that the model successfully captures steady-state load magnitudes. However, the presence of long tails shows the occurrence of the step timing errors observed in the time domain.

For the enforcement of the N-1 redundancy constraints and upper capacity hardware limits, Fig. 4.14 isolates the directional risks of these errors. By extracting the 99$^{\text{th}}$ percentiles from these empirical CDFs, we precisely size the dual stochastic buffers. For this specific simulation workload, the 99$^{\text{th}}$ percentile of accumulated shortages yields a required lower safety buffer of $B^{\text{lower}} = 0.1591$ pu-s. Conversely, the 99$^{\text{th}}$ percentile of accumulated surplus from synchronous updates yields an upper headroom buffer of $B^{\text{upper}} = 0.0179$ pu-s. These empirical margins are applied to the Hierarchical MPC, guaranteeing the system gracefully absorbs prediction uncertainties.

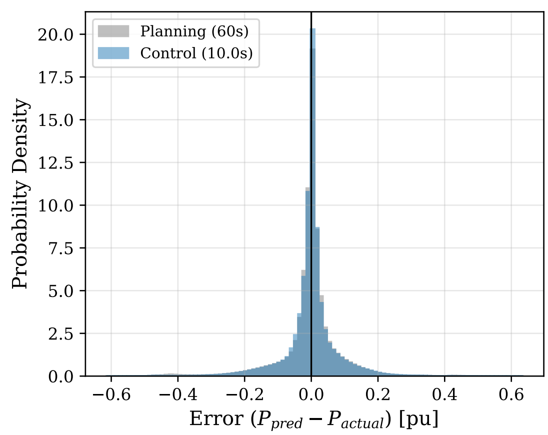

(a) Power prediction error PDF.

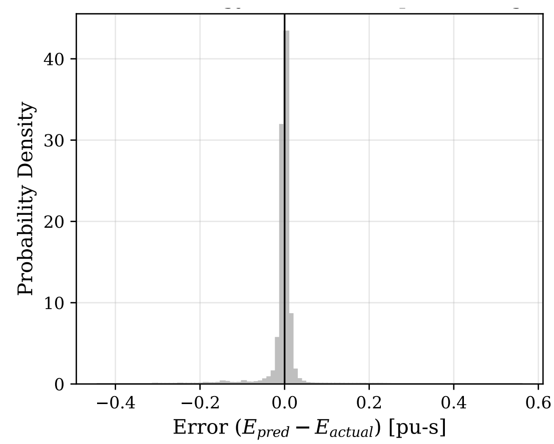

(b) T_{backup} energy prediction error PDF (T window).

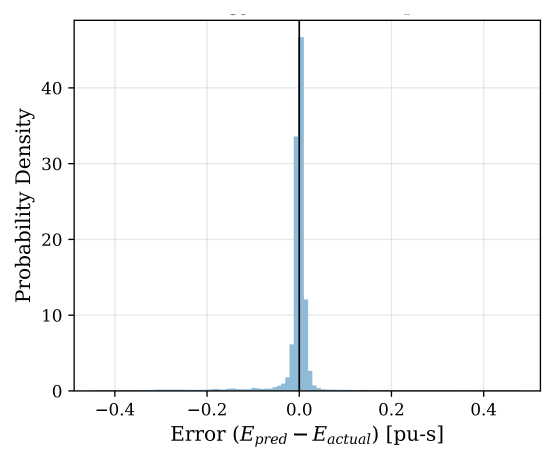

(c) T_{backup} energy prediction error PDF (T_{control} window).

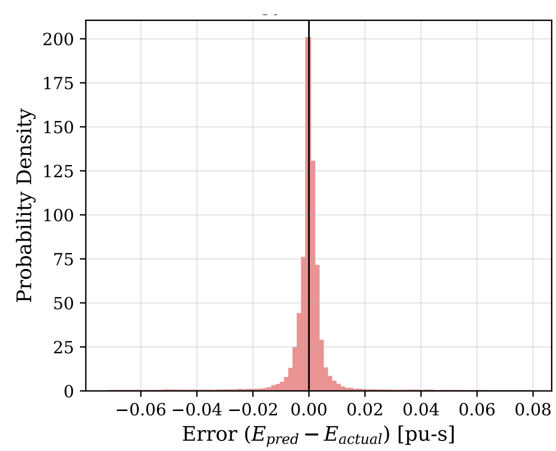

(d) T_{control} energy prediction error PDF ($t = 0$ update).

Figure 4.13: Probability distributions of the LSTM predictor uncertainties when tested on 63 LLM workload traces. While instantaneous power errors (a) are tightly zero-centered, integrating these auto-correlated timing errors over the various backup and control horizons (b-d) results in wider energy error dispersions that necessitate protective buffering.

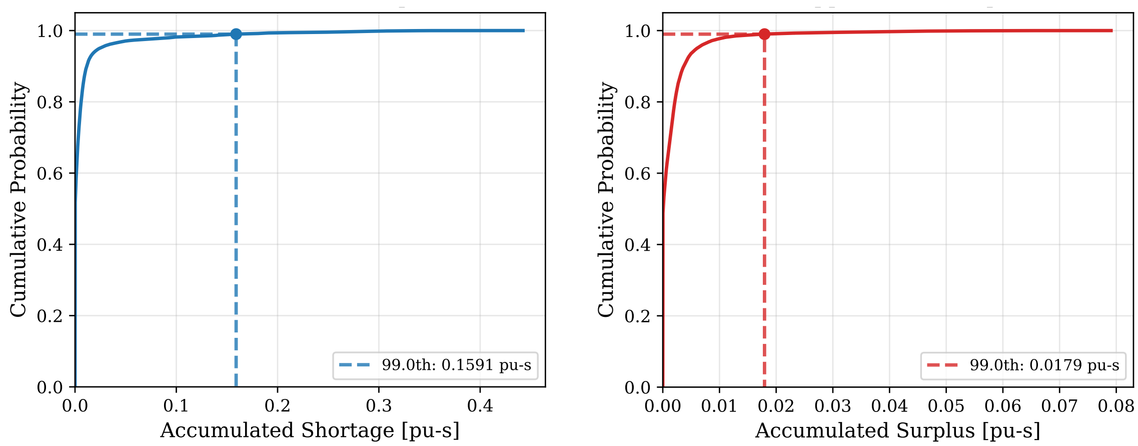

(a) Lower buffer requirement (Accumulated shortage).

(b) Upper buffer requirement (Accumulated surplus).

Figure 4.14: Cumulative Distribution Functions (CDFs) used to size the asymmetric safety margins. The 99^{th} percentile of the accumulated shortages (a) dictates B^{lower}, while the overestimates (b) dictate B^{upper}.

Control Performance and Constraint Management

The performance of the evaluated control strategies is quantified using the defined objective function over the whole simulated trajectory. The proposed Hierarchical MPC achieves a total realized cost of 0.0153, outperforming both the standard proportional droop controller (0.0436) and the unmitigated base load (0.1303), as detailed in Table 4.2. This cost reduction is primarily driven by minimizing power deviations from the average load and significantly reducing the transient ramp rates. As illustrated in the top panel of Fig. 4.16, the MPC establishes a smoothed power profile using its mean-ramp objective, whereas purely reactive penalization results in continuous drifting within the load states. While the MPC utilizes higher control effort—resulting in a slightly increased battery degradation cost compared to the droop controller—the substantial improvements in load smoothing outweigh this penalty under the selected tuning parameters.

A primary advantage of the predictive architecture is its ability to safely maximize battery utilization. Lacking foresight, standalone reactive controllers must rely on a conservative, static N-1 lower bound (teal dashed line, Fig. 4.15), which severely restricts the permissible depth of discharge. Conversely, the Hierarchical MPC dynamically calculates instantaneous backup requirements using load forecasts. By applying predefined stochastic buffers to these forecasts, the MPC probabilistically guarantees N-1 compliance while safely operating below the conservative static bound to reclaim stranded capacity.

Table 4.2

Final Realized Cost Evaluation by Control Strategy ($\times 10^{-2}$)

Scenario	**Total Cost**	**Smooth**	**Ramp**	**Degrad.**	**SoC**
Base Load	13.0300	6.4800	6.5500	0.0000	0.0000
Prop. Droop	4.3600	1.6400	2.6700	0.0400	0.0002
Hierarchical MPC	1.5300	0.5500	0.8900	0.0900	0.0006

Achieving comparable smoothing with a proportional droop controller requires aggressively increasing its gain and manually lowering its static bounds, risking hardware limit violations. During sustained compute phases, this aggressive tuning rapidly saturates the battery at maximum capacity (green line, Fig. 4.15), forcing instantaneous curtailment that produces unmitigated secondary load steps. Tuning the MPC objective weights presents similar challenges in balancing performance trade-offs; however, the optimization framework inherently incorporates system boundaries as constraints, automatically decelerating the control action to prevent limit-induced saturation.

To improve the droop control method in the face of these challenges would require significant structural augmentation, including state-of-charge feedback and predictive dynamic bounds. While the evaluated droop controller naturally charges and discharges in line with periodic workloads, it remains vulnerable to volatile load variations without explicit predictive boundary awareness. Ultimately, the MPC framework robustly smooths the active power drawn from the grid, effectively mitigating severe voltage sags at the point of interconnection (POI) while maintaining operational limits.

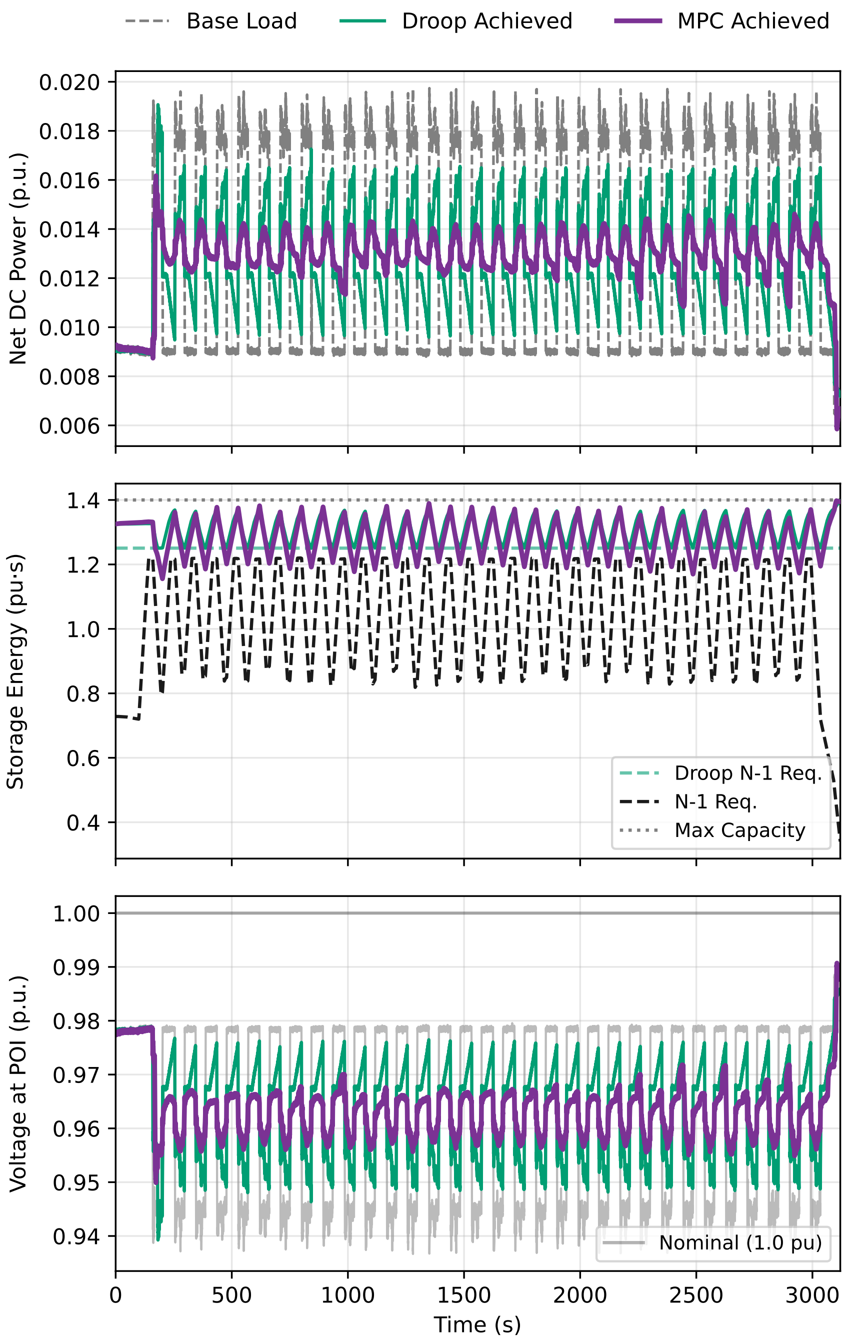

Figure 4.15: Simulation results comparing control strategies. Top: Net DC active power drawn from the grid. Middle: Aggregate State-of-Charge for the UPS batteries. Bottom: Resulting Voltage at the Point of Interconnection (POI). The MPC uses additional capacity while avoiding the dynamic buffered bounds (purple shaded region), whereas the Droop controller occasionally saturates against the static limits.

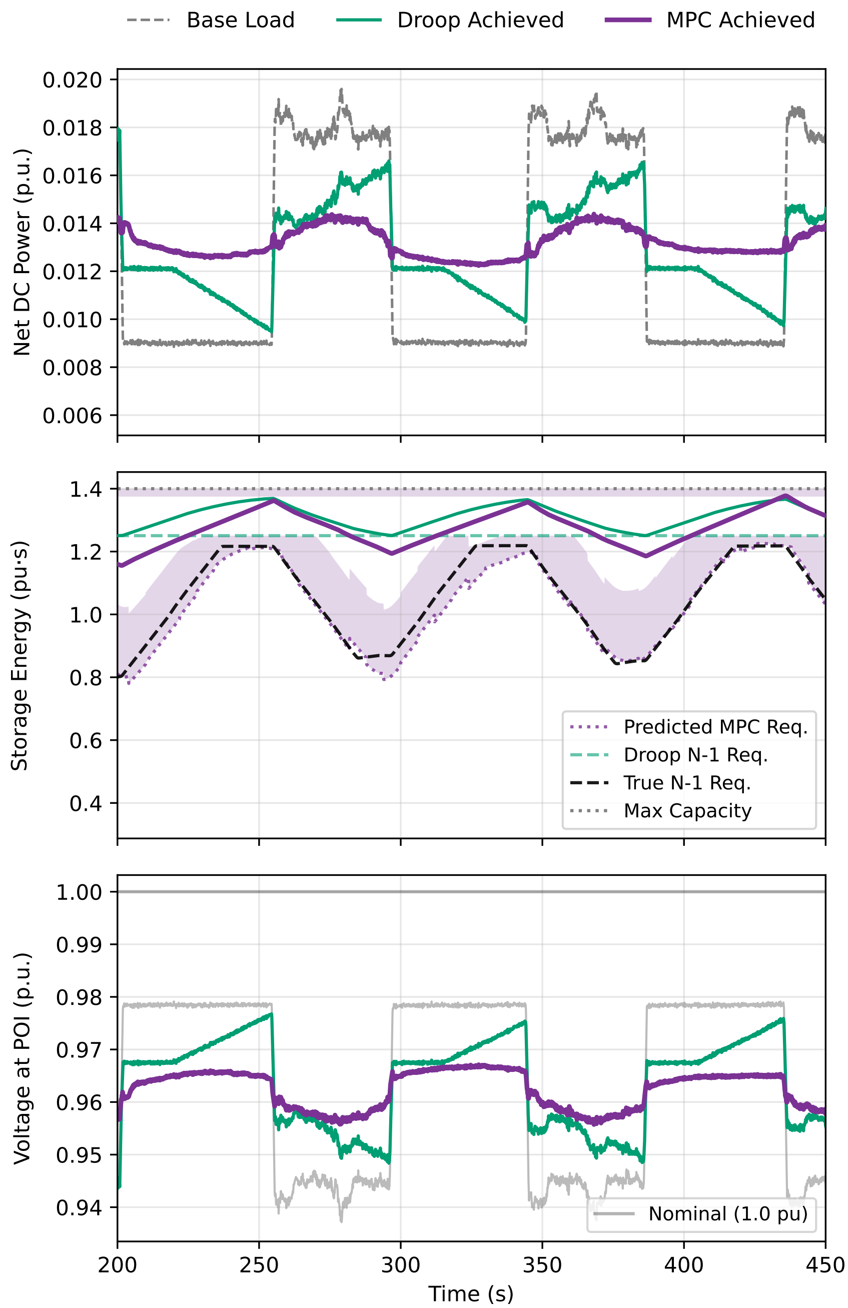

Figure 4.16: Zoomed in simulation results from $t = 200$ s to $t = 450$ s. The proposed Hierarchical MPC effectively reduces power deviation (top) while smoothly managing energy throughput (middle) to prevent the voltage steps (bottom) characteristic of the unmitigated base load.

Chapter 5

CONCLUSION

This dissertation explored three main topics relating to IBRs: the control of inverters with current magnitude constraints, the stability and sensitivity of systems with inverters, and the regulation of large inverter-based loads.

In Chapter 2, we covered a range of approaches for modeling and controlling inverters with current limits. We identified stability conditions for the control of an IBR with a current limiter giving insights into how uneven control actions can cause convergences challenges for systems with circular magnitude limits. We proposed a safety filter approach for IBR control that strictly enforces current limits with minimal intervention, advancing our understanding of how to manage inverters when transients preclude stiff voltage source operation. We showed that the geometry of the feasible output of an IBR constrained by current limits is convex opening the door for the application of control approaches and optimization methods that require convexity to systems of current constrained inverters. Lastly, we proposed a safe trajectory optimizing controller applicable to a grid-interfacing inverters.

In Chapter 3, we focused on the small-signal stability of power systems containing IBRs. First, we reviewed the topic of system strength and evaluated the applicability of existing system strength metrics to systems with high levels of IBRs. We found that traditional fault current based metrics are unable to identify critical control interactions in systems with IBRs that admittance model based metrics can. However, clear and intuitive admittance model based metrics still have yet to developed and adopted in industry. Next, we proposed a new grid strength metric based on the sensitivity of a system's eigenvalues to perturbations in power flow parameters. This approach to quantifying system strength through the sensitivity of the system's small-signal stability characteristics provides intuitive insights into what changes in the system can cause instability challenges and what remedial actions a system operator could take to resolve such issues.

In Chapter 4, we considered the impact that large data center loads have on system voltages and proposed two approaches for regulating voltages in systems with these loads. We developed a switching-reference reactive power controller for IBRs that can preemptively mitigate the rapid, step-change load fluctuations characteristic of LLM training. Next, we established that by utilizing the idle headroom in Uninterruptible Power Supply (UPS) energy storage—while strictly enforcing N-1 redundancy limits—data centers can actively smooth their net active power draw, significantly buffering the distribution feeder from severe compute volatility without compromising internal reliability.

5.1 Directions for Future Work

Based upon the results presented in this dissertation, there are several avenues for subsequent research.

Control of Inverter-Based Resources under Current Constraints

Relating to the stability of IBRs with current limiters in Section 2.2, future work includes considering the impact of disturbances in the grid voltage, investigating other objective functions for the reference MPC trajectories, and considering the cost of actuation required to maintain the current magnitude saturation bound.

Relating to the safety filter approach presented in Section 2.3, future work includes considering the impact of disturbances and parameter uncertainty, extending these results to include voltage magnitude control, modifying this method for integral control, and studying the stability of networks of inverters using this approach.

Lastly, relating to the Safe Trajectory Gradient Flow control presented in Section 2.5, future work includes, our future work includes finding conditions to guarantee the convergence of the rolling horizon trajectory problem, and how to handle the impact of disturbances on the ability of the approach to maintain safety. This could involve incorporating autoregressive filters for the real-time estimation of exogenous inputs [172]. In addition, reducing the computational complexity to make it easier to implement on microcontrollers or a hierarchical control structure with a slower outer loop controller that is driven with this safe trajectory gradient descent approach would be required to make this approach to the control of grid-interfacing inverters practical. Lastly, it is unclear what conditions are required such that the quadratic program always remains feasible as the trajectory horizon rolls forward.

Finally, analyzing the stability of interconnected networks of these current constrained inverters remains an open and highly relevant challenge.

System Strength and Small-Signal Stability

Building upon the system strength and eigenvalue sensitivity metrics introduced in Chapter 3, future studies could explore extending the method to consider second-order eigenvalue sensitivities to more accurately predict the impact of large perturbations. Additionally, future work could more accurately represent droop dynamics or automatic generation control (AGC) when determining how the steady-state power flow solution is perturbed. Because sensitivity-based stability approaches have proven highly effective for evaluating margins in converter-interfaced systems [173, 174], other directions include utilizing these proposed sensitivities to optimize generator dispatches to increase modal damping and strengthen the system's small-signal stability. In particular, these analytical sensitivities could be highly useful when considering small-signal stability-constrained AC OPF (SSSC-OPF) problems [175, 176], or when formulating sequential matrix perturbation-based optimization methods [177]. Formally, this casts the dispatch problem as a spectral abscissa optimization, where the goal is to systematically shift the dominant system modes deeper into the stable region and increasing their damping. Finally, another potential research direction is analyzing how the sensitivities of the eigenvalues change as they are perturbed, determining methods for minimizing the overall system sensitivity in terms of the eigenvalue sensitivity to parameter uncertainties.

Regulation of Large Data Center Loads

Regarding the regulation of data center workloads discussed in Chapter 4, future work could focus on testing the proposed control approaches under highly uncertain operational conditions, considering LLM inference loads and more accurately modeling the data center as a whole. Data centers loads are determined by the computing tasks they are running, and how these tasks are scheduled and dispatched to the servers within the center impact how CPU power consumption aggregates. Further, the routing of power through these data centers passes through many power electronics devices making electrical losses an important detail to consider. To fully understand the impact of data centers on the grid modeling, we must model system dynamics. While the approaches in our work utilize steady-state distribution network models (e.g., DistFlow), subsequent studies can consider modeling the voltage dynamics of the data center power electronics and grid system to accurately capture the behavior of the system induced by abrupt LLM computing loads.

BIBLIOGRAPHY

[1] J. A. Cohn, *The grid: biography of an American technology.* MIT press, 2017.

[2] G. S. Callendar, "The artificial production of carbon dioxide and its influence on temperature," *Quarterly Journal of the Royal Meteorological Society*, vol. 64, no. 275, pp. 223–240, 1938.

[3] International Energy Agency (IEA), "Global energy review 2025," International Energy Agency (IEA), Paris, Tech. Rep., 2025. [Online]. Available: https://www.iea.org/reports/global-energy-review-2025

[4] Western Electricity Coordinating Council (WECC), "State of the interconnection 2025," Western Electricity Coordinating Council, Tech. Rep., May 2025, published May 2025. Available online: https://feature.wecc.org/soti2025/index.html. [Online]. Available: https://feature.wecc.org/soti2025/index.html

[5] International Energy Agency (IEA), "Global battery storage capacity additions, 2010–2023," https://www.iea.org/data-and-statistics/charts/global-battery-storage-capacity-additions-2010-2023, Paris, 2024, licence: CC BY 4.0.

[6] L. Inc., "2025 levelized cost of energy+ (LCOE+) report," Lazard, New York, NY, Tech. Rep., Jun. 2025, 18th edition; released June 16, 2025. [Online]. Available: https://www.lazard.com/news-announcements/lazard-releases-2025-levelized-cost-of-energyplus-report-pr/

[7] R. W. Kenyon, M. Bossart, M. Marković, K. Doubleday, R. Matsuda-Dunn, S. Mitova, S. A. Julien, E. T. Hale, and B.-M. Hodge, "Stability and control of power systems with high penetrations of inverter-based resources: An accessible review of current knowledge and open questions," *Solar Energy*, vol. 210, pp. 149–168, 2020, special Issue on Grid Integration. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0038092X20305442

[8] A. F. Hoke, A. A. Nelson, J. Tan, R. Mahmud, V. Gevorgian, M. Elkhatib, C. Antonio, D. Arakawa, and K. Fong, "The frequency-watt function: Simulation and testing for the hawaiian electric companies," National Renewable Energy Laboratory (NREL), Golden, CO (United States), Tech. Rep., 07 2017. [Online]. Available: https://www.osti.gov/biblio/1373491

[9] J. Driesen and K. Visscher, "Virtual synchronous generators," in *2008 IEEE power and energy society general meeting.* IEEE, 2008, pp. 1–3.

[10] S. V. Dhople, B. B. Johnson, and A. O. Hamadeh, "Virtual oscillator control for voltage source inverters," in *2013 51st annual allerton conference on communication, control, and computing (Allerton).* IEEE, 2013, pp. 1359–1363.

[11] M. Chandorkar, D. Divan, and R. Adapa, "Control of parallel connected inverters in standalone ac supply systems," *IEEE Trans. Ind. Appl.*, vol. 29, no. 1, pp. 136–143, Feb. 1993.

[12] P. Piagi and R. Lasseter, "Autonomous control of microgrids," in *2006 IEEE Power Engineering Society General Meeting*, 2006, pp. 8 pp.–.

[13] B. Johnson, T. G. Roberts, O. Ajala, A. Dominguez-Garcia, S. Dhople, D. Ramasubramanian, A. Tuohy, D. Divan, and B. Kroposki, "A generic primary-control model for grid-forming inverters: Towards interoperable operation & control," in *Proceedings of the 55th Hawaii International Conference on System Sciences*, 2022. [Online]. Available: https://hdl.handle.net/10125/79751

[14] North American Electric Reliability Corporation (NERC), "White paper: Grid forming functional specifications for bps-connected battery energy storage systems," North American Electric Reliability Corporation, Atlanta, GA, USA, Tech. Rep., Sep. 2023, white Paper. [Online]. Available: https://www.nerc.com/globalassets/our-work/reports/white-papers/white_paper_gfm_functional_specification.pdf

[15] R. H. Lasseter, Z. Chen, and D. Pattabiraman, "Grid-forming inverters: A critical asset for the power grid," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 8, no. 2, pp. 925–935, 2020.

[16] J. Matevosyan, B. Badrzadeh, T. Prevost, E. Quitmann, D. Ramasubramanian, H. Urdal, S. Achilles, J. MacDowell, S. H. Huang, V. Vital, J. O'Sullivan, and R. Quint, "Grid-forming inverters: Are they the key for high renewable penetration?" *IEEE Power and Energy Magazine*, vol. 17, no. 6, pp. 89–98, 2019.

[17] B. Kroposki, B. Johnson, Y. Zhang, V. Gevorgian, P. Denholm, B.-M. Hodge, and B. Hannegan, "Achieving a 100% renewable grid: Operating electric power systems with extremely high levels of variable renewable energy," *IEEE Power and Energy Magazine*, vol. 15, no. 2, p. 61–73, Mar. 2017.

[18] N. Baeckeland and G.-S. Seo, "A guide to current limiting and stability with grid-forming inverters," United States, 2025.

[19] N. Baeckeland, D. Chatterjee, M. Lu, B. Johnson, and G.-S. Seo, "Overcurrent limiting in grid-forming inverters: A comprehensive review and discussion," *IEEE Transactions on Power Electronics*, vol. 39, no. 11, pp. 14 493–14 517, 2024.

[20] B. Fan, T. Liu, F. Zhao, H. Wu, and X. Wang, "A review of current-limiting control of grid-forming inverters under symmetrical disturbances," *IEEE Open Journal of Power Electronics*, vol. 3, pp. 955–969, 2022.

[21] IEEE PES-TR77, "Power system dynamic performance: Technical report on stability definitions and characterization of dynamic behavior in systems with high penetration of power electronic interfaced technologies," IEEE Power & Energy Society (PES), Tech. Rep., 2020. [Online]. Available: https://resourcecenter.ieee-pes.org/publications/technical-reports/pes_tp_tr77_psdp_stability_051320

[22] CIGRE Working Group B4.62, "Connection of wind farms to weak ac networks," *Technical Brochure 671*, 2016.

[23] H. Liu, X. Xie, J. He, T. Xu, Z. Yu, C. Wang, and C. Zhang, "Subsynchronous interaction between direct-drive pmsg based wind farms and weak ac networks," *IEEE Trans. Power Syst.*, vol. 32, no. 6, pp. 4708–4720, 2017.

[24] S.-H. Huang, J. Schmall, J. Conto, J. Adams, Y. Zhang, and C. Carter, "Voltage control challenges on weak grids with high penetration of wind generation: Ercot experience," in *2012 IEEE Power and Energy Society General Meeting*, 2012, pp. 1–7.

[25] Y. Cheng *et al.*, "Real-world subsynchronous oscillation events in power grids with high penetrations of inverter-based resources," *IEEE Trans. Power Syst.*, vol. 38, no. 1, pp. 316–330, 2023.

[26] P. J. Hart, M. Gong, H. Liu, Z. Chen, Y. Zhang, and Y. Wang, "Provably-stable overload ride-through control for grid-forming inverters using system-wide lyapunov function analysis," *IEEE Transactions on Energy Conversion*, vol. 37, no. 4, p. 2761–2776, Dec. 2022.

[27] B. Fan, T. Liu, F. Zhao, H. Wu, and X. Wang, "A review of current-limiting control of grid-forming inverters under symmetrical disturbances," *IEEE Open J. Power Electron.*, vol. 3, pp. 955–969, 2022.

[28] G. Xing, Y. Min, L. Chen, and H. Mao, "Limit induced bifurcation of grid-connected vsc caused by current limit," *IEEE Transactions on Power Systems*, vol. 36, no. 3, p. 2717–2720, May 2021.

[29] T. Joswig-Jones and B. Zhang, "Optimal control of grid-interfacing inverters with current magnitude limits," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 4623–4628.

[30] J. Bloemink and R. Iravani, "Control of a multiple source microgrid with built-in islanding detection and current limiting," *Power Delivery, IEEE Transactions on*, vol. 27, pp. 2122–2132, 10 2012.

[31] L. Zhou, S. Liu, Y. Chen, W. Yi, S. Wang, X. Zhou, W. Wu, J. Zhou, C. Xiao, and A. Liu, "Harmonic current and inrush fault current coordinated suppression method for vsg under non-ideal grid condition," *IEEE Transactions on Power Electronics*, vol. 36, no. 1, pp. 1030–1042, 2021.

[32] A. D. Paquette and D. M. Divan, "Virtual impedance current limiting for inverters in microgrids with synchronous generators," *IEEE Transactions on Industry Applications*, vol. 51, no. 2, pp. 1630–1638, 2015.

[33] H. Wu and X. Wang, "Small-signal modeling and controller parameters tuning of grid-forming vscs with adaptive virtual impedance-based current limitation," *IEEE Transactions on Power Electronics*, vol. 37, no. 6, pp. 7185–7199, 2022.

[34] P. Park, "A revisited popov criterion for nonlinear lur'e systems with sector-restrictions," *International Journal of Control*, vol. 68, no. 3, pp. 461–470, 1997.

[35] C. R. D. Osório, G. G. Koch, H. Pinheiro, R. C. Oliveira, and V. F. Montagner, "Robust current control of grid-tied inverters affected by lcl filter soft-saturation," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 8, pp. 6550–6561, 2019.

[36] O. Ajala, N. Baeckeland, B. Johnson, S. Dhople, and A. Domínguez-García, "Model reduction and dynamic aggregation of grid-forming inverter networks," *IEEE Transactions on Power Systems*, 2022.

[37] H. Wu, X. Wang, and L. Zhao, "Design considerations of current-limiting control for grid-forming capability enhancement of vscs under large grid disturbances," *IEEE Transactions on Power Electronics*, p. 1–5, 2024.

[38] M. Schneeberger, S. Mastellone, and F. Dörfler, "Advanced safety filter based on sos control barrier and lyapunov functions," 2025. [Online]. Available: https://arxiv.org/abs/2401.06901

[39] D. Groß and F. Dorfler, "Projected grid-forming control for current-limiting of power converters," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. Monticello, IL, USA: IEEE, Sep. 2019, p. 326–333.

[40] D. Groß and X. Lyu, "Towards constrained grid-forming control," in *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. Monticello, IL, USA: IEEE, Sep. 2023, pp. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/10313424/

[41] A. Yazdani and R. Iravani, *Voltage-Sourced Converters in Power Systems: Modeling, Control, and Applications*, ser. IEEE Press. Wiley, 2010. [Online]. Available: https://ieeexplore.ieee.org/book/6739364

[42] A. Nicolini, H. Pinheiro, F. Carnielutti, and J. Massing, "PLL parameters tuning guidelines to increase stability margins in multiple three-phase converters connected to weak grids," *IET Renewable Power Generation*, vol. 14, no. 12, Sep. 2020. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1049/iet-rpg.2020.0028

[43] N. Pogaku, M. Prodanovic, and T. C. Green, "Modeling, analysis and testing of autonomous operation of an inverter-based microgrid," *IEEE Transactions on Power Electronics*, vol. 22, no. 2, p. 613–625, Mar. 2007.

[44] S. Zhou, J. Liu, L. Zhou, and H. She, "Cross-coupling and decoupling techniques in the current control of grid-connected voltage source converter," in *2015 IEEE Applied Power Electronics Conference and Exposition (APEC)*. Charlotte, NC, USA: IEEE, Mar. 2015, p. 2821–2827.

[45] F. L. Luo and H. Ye, *Advanced DC/AC inverters: applications in renewable energy*. Crc Press, 2017.

[46] N. Bottrell and T. C. Green, "Comparison of current-limiting strategies during fault ride-through of inverters to prevent latch-up and wind-up," *IEEE Trans. Power Electron.*, vol. 29, no. 7, pp. 3786–3797, 2013.

[47] J. G. Kassakian, D. J. Perreault, G. C. Verghese, and M. F. Schlecht, *Principles of power electronics*. Cambridge University Press, 2023.

[48] N. R. Sorensen, E. V. Thomas, M. A. Quintana, S. Barkaszi, A. Rosenthal, Z. Zhang, and S. Kurtz, "Thermal study of inverter components," *IEEE journal of photovoltaics*, vol. 3, no. 2, pp. 807–813, 2013.

[49] J. J. Grainger and W. D. Stevenson, *Power System Analysis*. McGraw-Hill, 1994.

[50] T. Qoria, F. Gruson, F. Colas, X. Kestelyn, and X. Guillaud, "Current limiting algorithms and transient stability analysis of grid-forming vscs," *Electric Power Systems Research*, vol. 189, p. 106726, 2020.

[51] J. P. Hespanha, *Linear systems theory*. Princeton: Princeton University Press, 2009, ch. 8.6 Discrete-Time Case, pp. 65–66.

[52] S. P. Boyd, *Linear matrix inequalities in system and control theory*, ser. SIAM studies in applied mathematics. Philadelphia: Society for Industrial and Applied Mathematics, 1994, no. vol. 15, ch. 2.1 Linear Matrix Inequalities, pp. 7–9.

[53] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: https://www.gurobi.com

[54] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.

[55] T. Joswig-Jones and B. Zhang, "Safe control of grid-interfacing inverters with current magnitude limits," in *2025 Hawaii International Conference on System Sciences*. Hawaii International Conference on System Sciences, 2025. [Online]. Available: http://dx.doi.org/10.24251/HICSS.2025.363

[56] A. Yazdani and R. Iravani, *Voltage-Sourced Converters in Power Systems*. Hoboken, NJ: John Wiley & Sons, Inc., 2010.

[57] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, p. 3861–3876, Aug. 2017.

[58] E. D. Sontag, "A lyapunov-like characterization of asymptotic controllability," *SIAM Journal on Control and Optimization*, vol. 21, no. 3, pp. 462–471, 1983. [Online]. Available: https://doi.org/10.1137/0321028

[59] Z. Artstein, "Stabilization with relaxed controls," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 7, no. 11, pp. 1163–1173, 1983.

[60] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems*, vol. 43, no. 5, p. 137–177, Oct. 2023.

[61] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *53rd IEEE Conference on Decision and Control*. Los Angeles, CA, USA: IEEE, Dec. 2014, p. 6271–6278.

[62] A. Singletary, Y. Chen, and A. D. Ames, "Control barrier functions for sampled-data systems with input delays," 2020. [Online]. Available: https://arxiv.org/abs/2005.06418

[63] S. Buso and P. Mattavelli, *Digital Control in Power Electronics*, ser. Synthesis Lectures on Power Electronics. Cham: Springer International Publishing, 2015.

[64] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," *Journal of Control and Decision*, vol. 5, no. 1, pp. 42–60, 2018.

[65] L. Streitmatter, T. Joswig-Jones, and B. Zhang, "Geometry of the feasible output regions of grid-interfacing inverters with current limits," *IEEE Control Systems Letters*, vol. 9, pp. 685–690, 2025.

[66] S. Jayalath and M. Hanif, "Generalized lcl-filter design algorithm for grid-connected voltage-source inverter," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 3, pp. 1905–1915, 2016.

[67] S. B. Kjaer, J. K. Pedersen, and F. Blaabjerg, "A review of single-phase grid-connected inverters for photovoltaic modules," *IEEE transactions on industry applications*, vol. 41, no. 5, pp. 1292–1306, 2005.

[68] Y. Gu and T. C. Green, "Power system stability with a high penetration of inverter-based resources," *Proceedings of the IEEE*, vol. 111, no. 7, pp. 832–853, 2022.

[69] L. Huang, H. Xin, Z. Wang, L. Zhang, K. Wu, and J. Hu, "Transient stability analysis and control design of droop-controlled voltage source converters considering current limitation," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 578–591, 2017.

[70] M. Lu, V. Purba, S. Dhople, and B. Johnson, "Comparison of droop control and virtual oscillator control realized by andronov-hopf dynamics," in *IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society*, 2020, pp. 4051–4056.

[71] Y. Firouz, M. Tavakoli Bina, and B. Eskandari, "Efficiency of three-level neutral-point clamped converters: analysis and experimental validation of power losses, thermal modelling and lifetime prediction," *IET Power Electronics*, vol. 7, no. 1, pp. 209–219, 2014.

[72] L. Streitmatter, T. Joswig-Jones, and B. Zhang, "Convexity and optimal online control of grid-interfacing converters with current limits," 2026. [Online]. Available: https://arxiv.org/abs/2603.17135

[73] T. Joswig-Jones and B. Zhang, "Safe trajectory gradient flow control of a grid-interfacing inverter," in *2026 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2026, in press.

[74] A. Hauswirth, Z. He, S. Bolognani, G. Hug, and F. Dörfler, "Optimization algorithms as robust feedback controllers," *Annual Reviews in Control*, vol. 57, p. 100941, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1367578824000105

[75] Y. Chen, L. Cothren, J. Cortés, and E. Dall'Anese, "Online regulation of dynamical systems to solutions of constrained optimization problems," *IEEE Control Systems Letters*, vol. 7, pp. 3789–3794, 2023.

[76] M. Colombino, E. Dall'Anese, and A. Bernstein, "Online optimization as a feedback controller: Stability and tracking," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 1, pp. 422–432, 2020.

[77] G. Delimpaltadakis, P. Mestres, J. Cortés, and W. P. M. H. Heemels, "Feedback optimization with state constraints through control barrier functions," 2025. [Online]. Available: https://arxiv.org/abs/2504.00813

[78] A. Allibhoy and J. Cortés, "Anytime solution of constrained nonlinear programs via control barrier functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 6527–6532.

[79] ——, "Control-barrier-function-based design of gradient flows for constrained nonlinear programming," *IEEE Trans. Autom. Control*, vol. 69, no. 6, pp. 3499–3514, 2024.

[80] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020. [Online]. Available: https://doi.org/10.1007/s12532-020-00179-2

[81] P. Virtanen *et al.*, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.

[82] F. Fiedler *et al.*, "do-mpc: Towards fair nonlinear and robust model predictive control," *Control Engineering Practice*, vol. 140, p. 105676, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0967066123002459

[83] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[84] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.

[85] N. Baeckeland and G.-S. Seo, "Enhanced large-signal stability method for grid-forming inverters during current limiting," in *2023 IEEE 24th Workshop on Control and Modeling for Power Electronics (COMPEL)*, 2023, pp. 1–8.

[86] T. Joswig-Jones, S. Dong, J. Tan, D. Wu, and B. Zhang, “A review of system strength metrics for inverter-based power systems,” in *2026 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2026, in press.

[87] Australian Energy Market Operator (AEMO), “Integrated final report: Sa black system 28 september 2016,” AEMO, Tech. Rep., 2017. [Online]. Available: https://www.aemo.com.au/-/media/files/electricity/nem/market_notices_and_events/power_system_incident_reports/2017/integrated-final-report-sa-black-system-28-september-2016.pdf

[88] EirGrid and SONI, “System strength requirements: Secure operation of the all-island power system in the presence of high shares of inverter based resources,” EirGrid, Tech. Rep., 2024.

[89] H. H. Alhelou, N. Hosseinzadeh, and B. Bahrani, *Power System Strength*. The Institution of Engineering and Technology, 2023. [Online]. Available: https://digital-library.theiet.org/doi/abs/10.1049/PBPO247E

[90] North American Electric Reliability Corporation (NERC), “Integrating inverter-based resources into low short circuit strength systems reliability guidelines,” Dec. 2017. [Online]. Available: https://www.nerc.com/comm/RSTC_Reliability_Guidelines/Item_4a._Integrating%20_Inverter-Based_Resources_into_Low_Short_Circuit_Strength_Systems_-_2017-11-08-FINAL.pdf

[91] Australian Energy Market Operator (AEMO), “System strength requirements methodology,” AEMO, Tech. Rep., Dec. 2022.

[92] P. Kundur, N. J. Balu, and M. G. Lauby, *Power system stability and control*. McGraw-hill New York, 1994, vol. 7.

[93] S. Gordon, K. Bell, and Q. Hong, “Implications of reduced fault level and its relationship to system strength: A scotland case study,” in *CIGRÉ Session 2022*, Paris, France, Aug. 28–Sep. 2 2022, paper C4-10204_2022.

[94] M. G. Dozein, B. Berry, J. V. Milanović, and P. Mancarella, “System strength beyond fault level,” *IEEE Access*, vol. 13, pp. 104 184–104 200, 2025.

[95] A. Boričić, J. L. R. Torres, and M. Popov, “System strength: Classification, evaluation methods, and emerging challenges in IBR-dominated grids,” in *2022 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia)*, 2022, pp. 185–189.

[96] W. Wang, G. M. Huang, D. Ramasubramanian, and E. Farantatos, “Transient stability analysis and stability margin evaluation of phase-locked loop synchronised converter-based generators,” *IET Generation, Transmission & Distribution*, vol. 14,

no. 22, pp. 5000–5010, 2020. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-gtd.2020.0620

[97] D. Bowman, D. Ramasubramanian, R. McCann, E. Farantatos, A. Gaikwad, and J. Caspary, "Spp grid strength study with high inverter-based resource penetration," in *2019 North American Power Symposium (NAPS)*, 2019, pp. 1–6.

[98] Southwest Power Pool (SPP), "Spp 2019 inverter-based generation integration study," Southwest Power Pool, Inc., Tech. Rep., 2019, accessed: 2025-09-12. [Online]. Available: https://spp.org/documents/64834/20190828%20-%20spp%202019%20inverter%20based%20generation%20integration%20study.pdf

[99] W. W. Baker, D. Ramasubramanian, M. Val Escudero, E. Farantatos, and A. Gaikwad, "Application of an advanced short circuit strength metric to evaluate ireland's high renewable penetration scenarios," *IET Conference Proceedings*, vol. 2021, no. 2, p. 167–172, Jul. 2021.

[100] C. Henderson, A. Egea-Alvarez, T. Kneuppel, G. Yang, and L. Xu, "Grid strength impedance metric: An alternative to SCR for evaluating system strength in converter dominated systems," *IEEE Trans. Power Del.*, vol. 39, no. 1, p. 386–396, Feb. 2024.

[101] C. L. Fortescue, "Method of symmetrical co-ordinates applied to the solution of polyphase networks," *Transactions of the American Institute of Electrical Engineers*, vol. XXXVII, no. 2, pp. 1027–1140, 1918.

[102] N. Hosseinzadeh, L. Liyanarachchi, A. Gargoom, H. H. Alhelou, B. Bahrani, F. Fahiman, and E. Farahani, *Power system strength assessment with high penetration of inverter-based resources - a conceptual approach.* The Institution of Engineering and Technology, 2024, ch. Chapter 1, pp. 1–23. [Online]. Available: https://digital-library.theiet.org/doi/abs/10.1049/PBPO247E_ch1

[103] J. Z. Zhou, H. Ding, S. Fan, Y. Zhang, and A. M. Gole, "Impact of short-circuit ratio and phase-locked-loop parameters on the small-signal behavior of a vsc-hvdc converter," *IEEE Trans. Power Del.*, vol. 29, no. 5, pp. 2287–2296, 2014.

[104] IEEE, "IEEE guide for planning dc links terminating at ac locations having low short-circuit capacities," *IEEE Std 1204-1997*, pp. 1–216, 1997.

[105] Y. Li, Y. Gu, and T. C. Green, "Revisiting grid-forming and grid-following inverters: A duality theory," *IEEE Trans. Power Syst.*, vol. 37, no. 6, pp. 4541–4554, 2022.

[106] P. M. Anderson, *Analysis of Faulted Power Systems.* IEEE Press, 1995.

[107] International Electrotechnical Commission, *IEC 60909-0: Short-circuit currents in three-phase a.c. systems - Part 0: Calculation of currents*, IEC Std., 2016.

[108] Y. Zhang, S.-H. F. Huang, J. Schmall, J. Conto, J. Billo, and E. Rehman, "Evaluating system strength for large-scale wind plant integration," in *2014 IEEE PES General Meeting — Conference & Exposition.* National Harbor, MD, USA: IEEE, Jul. 2014, p. 1–5. [Online]. Available: http://ieeexplore.ieee.org/document/6939043/

[109] R. Fernandes, S. Achilles, and J. MacDowell, "Report to NERC ERSTF for composite short circuit ratio (CSCR) estimation guideline," *GE Energy Consulting*, 2015.

[110] CIGRE WG B4.41, "Systems with multiple dc infeed," CIGRE, Tech. Rep., 2008.

[111] D. Wu, G. Li, M. Javadi, A. M. Malyscheff, M. Hong, and J. N. Jiang, "Assessing impact of renewable energy integration on system strength using site-dependent short circuit ratio," *IEEE Trans. Sustain. Energy*, vol. 9, no. 3, p. 1072–1080, Jul. 2018.

[112] L. Unruh and A. Isaacs, "Description of Inverter Penetration SCR (IPSCR) metric for quantifying system strength in large networks," Electranix, Tech. Rep., 2021.

[113] S. Dong, Z. Guo, J. Tan, A. Hoke, and B. Kroposki, "A normalized short-circuit ratio for accessing impact of grid-following inverters on grid stability," 2025, in review.

[114] J. Sun, "Impedance-based stability criterion for grid-connected inverters," *IEEE Transactions on Power Electronics*, vol. 26, no. 11, pp. 3075–3078, 2011.

[115] Y. Gu, Y. Li, Y. Zhu, and T. C. Green, "Impedance-based whole-system modeling for a composite grid via embedding of frame dynamics," *IEEE Trans. Power Syst.*, vol. 36, no. 1, p. 336–345, Jan. 2021.

[116] Y. Zhu, Y. Gu, Y. Li, and T. Green, "Participation analysis in impedance models: The grey-box approach for power system stability," *IEEE Trans. Power Syst.*, vol. 37, no. 1, p. 343–353, Jan. 2022.

[117] E. Stein and R. Shakarchi, *Complex Analysis*, ser. Princeton lectures in analysis. Princeton University Press, 2010. [Online]. Available: https://books.google.com/books?id=0ECHh9tjPUAC

[118] H. Zong, J. Lyu, C. Zhang, X. Cai, M. Molinas, and F. Rao, "Mimo impedance based stability analysis of dfig-based wind farm with mmc-hvdc in modifed sequence domain," in *8th Renewable Power Generation Conference (RPG 2019)*, 2019, pp. 1–7.

[119] H. Zou, Z. Tang, T. Hu, and G. Todeschini, "Critical short circuit ratio calculation for mimo systems using disk margins," in *Proceedings of PowerTech*, Kiel, Germany, 2025.

[120] F. Zhang, H. Xin, D. Wu, Z. Wang, and D. Gan, "Assessing strength of multi-infeed LCC-HVDC systems using generalized short-circuit ratio," *IEEE Trans. Power Syst.*, vol. 34, no. 1, p. 467–480, Jan. 2019.

[121] W. Dong, H. Xin, D. Wu, and L. Huang, "Small signal stability analysis of multi-infeed power electronic systems based on grid strength assessment," *IEEE Trans. Power Syst.*, vol. 34, no. 2, p. 1393–1403, Mar. 2019.

[122] Y. Zhou *et al.*, "Small-signal stability assessment of heterogeneous grid-following converter power systems based on grid strength analysis," *IEEE Trans. Power Syst.*, vol. 38, no. 3, p. 2566–2579, May 2023.

[123] F. Ma, H. Xin, D. Wu, Y. Liu, and X. Chen, "Assessing small-signal grid strength of 100% inverter-based power systems," *IEEE Trans. Power Del.*, vol. 39, no. 5, p. 2784–2796, Oct. 2024.

[124] C. Yang, L. Huang, H. Xin, and P. Ju, "Placing grid-forming converters to enhance small signal stability of pll-integrated power systems," *IEEE Trans. Power Syst.*, vol. 36, no. 4, p. 3563–3573, Jul. 2021.

[125] Y. Liu, Y. Chen, H. Xin, J. Tu, L. Zhang, M. Song, and J. Zhu, "System strength constrained grid-forming energy storage planning in renewable power systems," *IEEE Trans. Sustain. Energy*, vol. 16, no. 2, p. 981–994, Apr. 2025.

[126] H. Yuan, H. Xin, D. Wu, Z. Li, X. Qin, Y. Zhou, and L. Huang, "Assessing maximal capacity of grid-following converters with grid strength constraints," *IEEE Trans. Sustain. Energy*, vol. 13, no. 4, p. 2119–2132, Oct. 2022.

[127] X. Lai, Y. Liu, X. Zhang, H. Xin, J. Zhu, and T. Dai, "Grid strength constrained cooperative control of renewable power plants for guaranteed small-signal stability," *IEEE Trans. Power Syst.*, vol. 39, no. 6, p. 7425–7428, Nov. 2024.

[128] Z. Hou, B. Yin, F. Ma, C. Cong, R. Leng, L. Huang, and H. Xin, "System strength assessment of renewable energy sending systems with local loads," in *2025 IEEE Power & Energy Society General Meeting (PESGM)*, 2025.

[129] R. Leng, C. Cong, H. Chen, B. Yin, H. Xin, and L. Huang, "Small-signal stability analysis of grid-following/ grid-forming hybrid systems based on voltage and current strength," in *2025 IEEE Power & Energy Society General Meeting (PESGM)*, 2025, pp. 1–5.

[130] D. You, S. Jeong, Y. Yoo, S. Jung, and G. Jang, "Strength assessment procedure using grid strength impedance metric with pscad," in *2025 IEEE International Conference on Consumer Electronics (ICCE).* Las Vegas, NV, USA: IEEE, Jan. 2025, p. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/10930058/

[131] Y. Zhu, T. C. Green, X. Zhou, Y. Li, D. Kong, and Y. Gu, "Impedance margin ratio: A new metric for small-signal system strength," *IEEE Trans. Power Syst.*, vol. 39, no. 6, p. 7291–7303, Nov. 2024.

[132] A. Hadjileonidas, Y. Zhu, and T. C. Green, "Admittance margin: A guideline for placement of grid-following and grid-forming inverters regarding small-signal stability," in *IECON 2024 - 50th Annual Conference of the IEEE Industrial Electronics Society*. Chicago, IL, USA: IEEE, Nov. 2024, p. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/10905868/

[133] M. P. Richwine, N. W. Miller, A. J. Siler, H. Jung, and P. Dalton, "Power system stability analysis & planning using impedance-based methods," *IET Conference Proceedings*, vol. 2023, no. 20, p. 355–363, Nov. 2023.

[134] Y. Wang, X. Wang, F. Blaabjerg, and Z. Chen, "Frequency scanning-based stability analysis method for grid-connected inverter system," in *2017 IEEE 3rd International Future Energy Electronics Conference and ECCE Asia (IFEEC 2017 - ECCE Asia)*, 2017, pp. 1575–1580.

[135] T. Joswig-Jones, B. Zhang, S. Dong, and J. Tan, "System strength sensitivity to power flow perturbations in ac power systems with inverter-based resources," in *2026 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2026, in press.

[136] J. Ma, Z. Dong, and P. Zhang, "Eigenvalue sensitivity analysis for dynamic power system," in *2006 International Conference on Power System Technology*. Chongqing, China: IEEE, Oct. 2006, p. 1–7. [Online]. Available: http://ieeexplore.ieee.org/document/4116349/

[137] J. W. Demmel, *Applied Numerical Linear Algebra*. Philadelphia, PA: SIAM, 1997, ch. 7: Iterative Methods for Linear Systems, pp. 283–360.

[138] B. Gustavsen and A. Semlyen, "Rational approximation of frequency domain responses by vector fitting," *IEEE Trans. Power Del.*, vol. 14, no. 3, pp. 1052–1061, 1999.

[139] G. H. Golub and C. F. Van Loan, "Unsymmetric eigenvalue problems," in *Matrix Computations*, 4th ed. Baltimore, MD, USA: Johns Hopkins University Press, 2013, ch. 7.

[140] P. Kundur, "Small-signal stability," in *Power System Stability and Control*. New York, NY, USA: McGraw-Hill, 1994, ch. 12.

[141] J. Rommes and N. Martins, "Computing large-scale system eigenvalues most sensitive to parameter changes, with applications to power system small-signal stability," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 434–442, 2008.

[142] Y. Li, "Future-power-networks/simplus-grid-tool," 2021. [Online]. Available: https://github.com/Future-Power-Networks/Simplus-Grid-Tool

[143] R. D. Christie, "Power systems test case archive: 118 bus power flow test case," University of Washington, Department of Electrical Engineering. [Online]. Available: https://labs.ece.uw.edu/pstca/pf118/pg_tca118bus.htm, Aug. 1993.

[144] OPAL-RT Technologies, "PSS/E import – IEEE 118-bus system: Islanding of power grid in presence of large disturbances (phasor06_ieee118)," ePHASORsim Documentation. [Online]. Available: https://opal-rt.atlassian.net/wiki/spaces/PDOCHS/pages/416711042.

[145] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, "Scaling laws for neural language models," *arXiv preprint arXiv:2001.08361*, 2020.

[146] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. Casas, L. A. Hendricks, J. Welbl, A. Clark *et al.*, "Training compute-optimal large language models," *arXiv preprint arXiv:2203.15556*, vol. 10, 2022.

[147] I. E. Agency, "Energy and ai," https://www.iea.org/reports/energy-and-ai, 2025.

[148] Y. Chen and B. Zhang, "Voltage regulation in distribution systems with data center loads," Jul. 2025, arXiv:2507.06416 [eess]. [Online]. Available: http://arxiv.org/abs/2507.06416

[149] "IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," IEEE Std 1547-2018 (Revision of IEEE Std 1547-2003), IEEE Standards Association, 2018.

[150] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3794–3803, 2015.

[151] K. Baker, A. Bernstein, E. Dall'Anese, and C. Zhao, "Network-cognizant voltage droop control for distribution grids," *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 2098–2108, 2018.

[152] W. Cui, Y. Xie, S. Low, A. Wierman, and B. Zhang, "Leveraging predictions in power system voltage control: An adaptive approach," *arXiv preprint arXiv:2509.09937*, 2025.

[153] V. Sakalkar, V. Kontorinis, D. Landhuis, S. Li, D. De Ronde, T. Blooming, A. Ramesh, J. Kennedy, C. Malone, J. Clidaras *et al.*, "Data center power oversubscription with a medium voltage power plane and priority-aware capping," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 497–511.

[154] S. Hou, W. Ni, S. Zhao, B. Cheng, S. Chen, and J. Chen, "Decentralized real-time optimization of voltage reconfigurable cloud computing data center," *IEEE Transactions on Green Communications and Networking*, vol. 4, no. 2, pp. 577–592, 2020.

[155] Y. Xie, W. Cui, and A. Wierman, "Enhancing data center low-voltage ride-through," *arXiv preprint arXiv:2510.03867*, 2025.

[156] Z. Liang, J.-W. Chung, M. Chowdhury, J. Chen, and V. Dvorkin, "Gpu-to-grid: Voltage regulation via gpu utilization control," *arXiv preprint arXiv:2602.05116*, 2026.

[157] A. Azizi, S. Morovati, A. Zamani, J. Piruzza, D. Guo, Z. Liu, and P. Taimela, "Strengthening data center operations using grid-forming battery energy storage as a line-interactive uninterruptible power supply," *International Journal of Electrical Power & Energy Systems*, vol. 175, p. 111638, 2026.

[158] B. A. Ross, X. Lyu, U. C. Nwaneto, S. M. Mohiuddin, and A. B. Nassif, "Using grid-forming energy storage systems to provide dynamic active power support for hyperscale data center," *IEEE Access*, 2026.

[159] Y. Fan, X. Tao, Z. Ye, K. Kumar, and R. Gadh, "From co-location to coordination: Ev aggregator dispatch for voltage regulation of ai data center–integrated distribution systems," *Available at SSRN 6346298*, 2026.

[160] E. Choukse, B. Warrier, S. Heath, L. Belmont, A. Zhao, H. A. Khan, B. Harry, M. Kappel, R. J. Hewett, K. Datta *et al.*, "Power stabilization for ai training datacenters," *arXiv preprint arXiv:2508.14318*, 2025.

[161] M. Yan, T. Joswig-Jones, B. Zhang, Y. Chen, and W. Cui, "Switching-reference voltage control for distribution systems with ai-training data centers," 2026. [Online]. Available: https://arxiv.org/abs/2603.15588

[162] M. Baran and F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401–1407, 1989.

[163] M. Farivar and S. H. Low, "Branch flow model: Relaxations and convexification—part i," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 2554–2564, 2013.

[164] ANSI/NEMA, *American National Standard for Electric Power Systems and Equipment-Voltage Ratings (60 Hertz).* National Electrical Manufacturers Association, 2025.

[165] W. Cui, J. Li, and B. Zhang, "Decentralized safe reinforcement learning for inverter-based voltage control," *Electric Power Systems Research*, vol. 211, p. 108609, 2022.

[166] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE transactions on power systems*, vol. 33, no. 3, pp. 2882–2894, 2017.

[167] X. Yuan, J. Hu, and S. Cheng, "Multi-time scale dynamics in power electronics-dominated power systems," *Frontiers of Mechanical Engineering*, vol. 12, no. 3, pp. 303–311, 2017.

[168] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.

[169] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "Qlora: Efficient fine-tuning of quantized llms," *Advances in neural information processing systems*, vol. 36, pp. 10 088–10 115, 2023.

[170] L. A. Barroso, J. Clidaras, and U. Hölzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Second Edition*, 2013. [Online]. Available: http://dx.doi.org/10.2200/S00516ED2V01Y201306CAC024

[171] B. Zhang and E. Nasr, "Distributed redundant integration of data center battery storage with the grid for regulation services," in *2021 IEEE Power & Energy Society General Meeting (PESGM).* Washington, DC, USA: IEEE, Jul. 2021, p. 1–5. [Online]. Available: https://ieeexplore.ieee.org/document/9638219/

[172] R. Heydari, H. Young, F. Flores-Bahamonde, S. Vaez-Zadeh, C. González-Castaño, S. Sabzevari, and J. Rodríguez, "Model-free predictive control of grid-forming inverters with *lcl* filters," *IEEE Transactions on Power Electronics*, vol. 37, no. 8, pp. 9200–9211, 2022.

[173] L. Yang, Z. Xu, J. Østergaard, Z. Y. Dong, K. P. Wong, and X. Ma, "Oscillatory stability and eigenvalue sensitivity analysis of a dfig wind turbine system," *IEEE Transactions on Energy Conversion*, vol. 26, no. 1, pp. 328–339, 2011.

[174] J. Fang, S. Yang, Z. Xie, X. Zhang, and L. Chang, "Gradient-based black-box method for improving the stability margin of power station constructed by inverters without exposing controller details," *IEEE Transactions on Power Electronics*, vol. 40, no. 4, pp. 4830–4847, 2025.

[175] D. Pullaguram, R. Madani, T. Altun, and A. Davoudi, "Small-signal stability-constrained optimal power flow for inverter dominant autonomous microgrids," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 7, pp. 7318–7328, 2022.

[176] Y. Chen, Q. Yang, H. Xiong, X. Chen, and J. Wen, "Optimal operation for energy storage with oscillation stability constraints considering spatial differences," *IEEE Transactions on Power Systems*, vol. 41, no. 2, p. 868–883, Mar. 2026.

[177] Y. Qin, C. Duan, B. Zhao, C. Wang, and Z. Peng, "A matrix perturbation-based sequential optimization method for virtual inertia in grid-following and grid-forming hybrid systems," in *2025 9th International Conference on Smart Grid and Smart Cities (ICSGSC)*, 2025, pp. 338–344.

Appendix A

OPTIMAL & SAFE CONTROL OF GRID-INTERFACING INVERTERS WITH CURRENT MAGNITUDE LIMITS

A.1 Full-order Inverter with LCL *Filter Model*

We consider an inverter system connected to the grid via an *LCL* filter [41], with inner control loops [36], a phase-locked loop (PLL) [42] and a current limiter. This system has the dynamics

$$\dot{x} = \begin{bmatrix} \dot{\delta}_{\mathrm{pll}} \\ \dot{\Pi}_{\mathrm{pll}} \\ \dot{\Phi}_{\mathrm{d}} \\ \dot{\Phi}_{\mathrm{q}} \\ \dot{\Gamma}_{\mathrm{d}} \\ \dot{\Gamma}_{\mathrm{q}} \\ \dot{I}_{i\mathrm{d}} \\ \dot{I}_{i\mathrm{q}} \\ \dot{W}_{\mathrm{d}} \\ \dot{W}_{\mathrm{q}} \\ \dot{I}_{g\mathrm{d}} \\ \dot{I}_{g\mathrm{q}} \end{bmatrix} = \begin{bmatrix} k_{i_{pll}}\Pi_{\mathrm{pll}} + k_{p_{pll}}W_{\mathrm{q}} \\ W_{\mathrm{q}} \\ \sqrt{2}V\cos(\delta_i) - W_{\mathrm{d}} \\ \sqrt{2}V\sin(\delta_i) - W_{\mathrm{q}} \\ -I_{i\mathrm{d}} + \mathrm{sat}(I_{i\mathrm{d}}^*) \\ -I_{i\mathrm{q}} + \mathrm{sat}\left(I_{i\mathrm{q}}^*\right) \\ \omega_{\mathrm{ref}}I_{i\mathrm{q}} - \frac{R_iI_{i\mathrm{d}}}{L_i} + \frac{V_{\mathrm{d}}-W_{\mathrm{d}}}{L_i} \\ -\omega_{\mathrm{ref}}I_{i\mathrm{d}} - \frac{R_iI_{i\mathrm{q}}}{L_i} + \frac{V_{\mathrm{q}}-W_{\mathrm{q}}}{L_i} \\ \omega_{\mathrm{ref}}W_{\mathrm{q}} + \frac{-I_{g\mathrm{d}}+I_{i\mathrm{d}}}{C} \\ -\omega_{\mathrm{ref}}W_{\mathrm{d}} + \frac{-I_{g\mathrm{q}}+I_{i\mathrm{q}}}{C} \\ \omega_{\mathrm{ref}}I_{g\mathrm{q}} - \frac{R_gI_{g\mathrm{d}}}{L_g} + \frac{-\sqrt{2}E\cos\left(\delta_{\mathrm{pll}}\right)+W_{\mathrm{d}}}{L_g} \\ -\omega_{\mathrm{ref}}I_{g\mathrm{d}} - \frac{R_gI_{g\mathrm{q}}}{L_g} + \frac{\sqrt{2}E\sin\left(\delta_{\mathrm{pll}}\right)+W_{\mathrm{q}}}{L_g} \end{bmatrix} \tag{A.1}$$

where

$$u = \begin{bmatrix} V \\ \delta_i \end{bmatrix} = -K\begin{bmatrix} -I_{\mathrm{d}}^* + I_{g\mathrm{d}} \\ -I_{\mathrm{q}}^* + I_{g\mathrm{q}} \end{bmatrix} + \begin{bmatrix} V^* \\ \delta^* \end{bmatrix} \tag{A.2}$$

and I_{id}^*, I_{iq}^*, V_{d}, V_{q}, ω_{ref} are defined as

$$\begin{bmatrix} I_{id}^* \\ I_{iq}^* \\ V_{\mathrm{d}} \\ V_{\mathrm{q}} \\ \omega_{\mathrm{ref}} \end{bmatrix} = \begin{bmatrix} -C\omega_{\mathrm{nom}}W_{\mathrm{q}} + k_{iv}\Phi_{\mathrm{d}} + k_{pv}(\sqrt{2}V\cos(\delta_i) - W_{\mathrm{d}}) + I_{gd} \\ C\omega_{\mathrm{nom}}W_{\mathrm{d}} + k_{iv}\Phi_{\mathrm{q}} + k_{pv}(\sqrt{2}V\sin(\delta_i) - W_{\mathrm{q}}) + I_{gq} \\ -L_i\omega_{\mathrm{nom}}I_{iq} + k_{ii}\Gamma_{\mathrm{d}} + k_{pi}I_{id}^* - I_{id} + W_{\mathrm{d}} \\ L_i\omega_{\mathrm{nom}}I_{id} + k_{ii}\Gamma_{\mathrm{q}} + k_{pi}I_{iq}^* - I_{iq} + W_{\mathrm{q}} \\ \dot{\delta}_{\mathrm{pll}} \end{bmatrix} \tag{A.3}$$

where the states x contains the PLL angle difference from the grid voltage (δ_{pll}), the PLL integral term (Π_{pll}), the voltage controller integral terms (Φ_{dq}), the current controller integral terms (Γ_{dq}), the inverter-side inductor currents ($\mathbf{I}_{i\mathrm{dq}}$), the capacitor voltages ($\mathbf{W}_{\mathrm{dq}}$), and the grid-side inductor currents ($\mathbf{I}_{g\mathrm{dq}}$). In the dq-reference frame the power can be calculated as

$$\begin{bmatrix} P \\ Q \end{bmatrix} = \frac{3}{2}\begin{bmatrix} V_{\mathrm{d}} & V_{\mathrm{q}} \\ V_{\mathrm{q}} & -V_{\mathrm{d}} \end{bmatrix}\begin{bmatrix} I_{\mathrm{d}} \\ I_{\mathrm{q}} \end{bmatrix}, \tag{A.4}$$

with V_{dq}, I_{dq} being the voltage and current at the node of interest.

A.2 Proof of Theorem 1

We select the Lyapunov function $V(z) = z^\top z$. It is clear that $V(0) = 0$. Then it suffices to show that $V(\Delta\mathbf{x}_{t+1}) < V(\Delta\mathbf{x}_t)\ \forall\ \Delta\mathbf{x}_t \in \mathbb{R}^2,\ \Delta\mathbf{x}_t \neq 0$.

If $\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2 \leq 1$, then $\Delta\mathbf{x}_{t+1} = \mathbf{A}(\Delta\mathbf{x}_t) + Bu(\Delta\mathbf{x}_t)$. Then $V(\Delta\mathbf{x}_{t+1}) < V(\Delta\mathbf{x}_t)$ by (2.8). Therefore, we focus on when $\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2 > 1$, where $\hat{\mathbf{x}}_t = \frac{\Delta\mathbf{x}_t + \mathbf{x}^*}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2}$ from (2.5) and (2.7b) such that we have the step dynamics

$$\Delta\mathbf{x}_{t+1} = (\mathbf{A} - BK)\left(\frac{\Delta\mathbf{x}_t + \mathbf{x}^*}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2} - \mathbf{x}^*\right).$$

By (2.8), we have

$$V(\Delta\mathbf{x}_{t+1}) < \left\|\frac{\Delta\mathbf{x}_t + \mathbf{x}^*}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2} - \mathbf{x}^*\right\|_2^2$$

Now we show that

$$V(\Delta\mathbf{x}_{t+1}) < \left\|\frac{\Delta\mathbf{x}_t + \mathbf{x}^*}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2} - \mathbf{x}^*\right\|_2^2 < V(\Delta\mathbf{x}_t) = \|\Delta\mathbf{x}_t\|_2^2. \tag{A.5}$$

Given the vector $\mathbf{x}_t$, let $\mathbf{v}$ to be a vector orthogonal to it. Then $\{\mathbf{x}_t, \mathbf{v}\}$ forms a basis in $\mathbb{R}^2$, and we can write $\mathbf{x}^*$ as

$$\mathbf{x}^* = c_1\mathbf{x}_t + c_2\mathbf{v},$$

where $c_1 = \frac{\langle \mathbf{x}_t, \mathbf{x}^* \rangle}{\|\mathbf{x}_t\|_2^2}$ and the term $c_1\mathbf{x}_t$ represents the projection of $\mathbf{x}^*$ onto $\mathbf{x}_t$. Then

$$V(\Delta\mathbf{x}_t) = \|\mathbf{x}_t - \mathbf{x}^*\|_2^2 = \left\| \mathbf{x}_t - \frac{\langle \mathbf{x}_t, \mathbf{x}^* \rangle}{\|\mathbf{x}_t\|_2^2}\mathbf{x}_t \right\|_2^2 + \|c_2\mathbf{v}\|_2^2. \tag{A.6}$$

To show (A.5), we have

$$\begin{aligned}
V(\Delta\mathbf{x}_{t+1}) &< \left\| \frac{\Delta\mathbf{x}_t + \mathbf{x}^*}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2} - \mathbf{x}^* \right\|_2^2 \\
&\overset{(a)}{=} \left\| \frac{1 - \langle \mathbf{x}_t, \mathbf{x}^* \rangle}{\|\Delta\mathbf{x}_t + \mathbf{x}^*\|_2}\mathbf{x}_t \right\|_2^2 + \|c_2\mathbf{v}\|_2^2 \\
&= \left(\frac{1}{\|\mathbf{x}_t\|_2} - \frac{\langle \mathbf{x}^*, \mathbf{x}_t \rangle}{\|\mathbf{x}_t\|_2^2} \right)^2 \|\mathbf{x}_t\|_2^2 + \|c_2 v\|_2^2 \\
&\overset{(b)}{<} \left(1 - \frac{\langle \mathbf{x}^*, \mathbf{x}_t \rangle}{\|\mathbf{x}_t\|_2^2} \right)^2 \|\mathbf{x}_t\|_2^2 + \|c_2\mathbf{v}\|_2^2 \\
&= \left\| \mathbf{x}_t - \frac{\langle \mathbf{x}_t, \mathbf{x}^* \rangle}{\|\mathbf{x}_t\|_2^2}\mathbf{x}_t \right\|_2^2 + \|c_2\mathbf{v}\|_2^2 \\
&\overset{(c)}{=} \|\mathbf{x}_t - \mathbf{x}^*\|_2^2,
\end{aligned}$$

where (a) follows from $\mathbf{x}_t$ and $\mathbf{v}$ being orthogonal to each other and (c) follows (A.6). The step (b) follows from the fact that

$$\frac{\langle \mathbf{x}^*, \mathbf{x}_t \rangle}{\|\mathbf{x}_t\|_2^2} < \frac{1}{\|\mathbf{x}_t\|_2} < 1.$$

To see this, note that by assumption, we assume $\mathbf{x}_t$ saturates, so $\|\mathbf{x}_t\|_2^2 > 1$ and $\frac{1}{\|\mathbf{x}_t\|_2} < 1$. Next, we manipulate the inequality $\frac{\langle \mathbf{x}^*, \mathbf{x}_t \rangle}{\|\mathbf{x}_t\|_2^2} < \frac{1}{\|\mathbf{x}_t\|_2}$ by applying the Cauchy-Schwarz inequality ($|\langle \mathbf{x}^*, \mathbf{x}_t \rangle| \leq \|\mathbf{x}^*\|_2 \cdot \|\mathbf{x}_t\|_2$) to get

$$\frac{\langle \mathbf{x}^*, \mathbf{x}_t \rangle}{\|\mathbf{x}_t\|_2^2} \leq \frac{\|\mathbf{x}^*\|_2 \cdot \|\mathbf{x}_t\|_2}{\|\mathbf{x}_t\|_2^2} = \frac{\|\mathbf{x}^*\|_2}{\|\mathbf{x}_t\|_2} < \frac{1}{\|\mathbf{x}_t\|_2},$$

which is true since we assume $\|\mathbf{x}^*\|_2 < 1$.

A.3 Proof of Theorem 2

Proof of Theorem 2: We have

$$\hat{\mathbf{A}} = \mathbf{I} - \Delta t \begin{bmatrix} \frac{R}{L} & \omega_{\text{nom}} \\ -\omega_{\text{nom}} & \frac{R}{L} \end{bmatrix} = \mathbf{I} - \Delta t \mathbf{A}$$

$$\hat{\mathbf{B}} = \Delta t \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix} = \Delta t \mathbf{B}$$

where $\mathbf{A} = \begin{bmatrix} \frac{R}{L} & \omega_{\text{nom}} \\ -\omega_{\text{nom}} & \frac{R}{L} \end{bmatrix}$ and $\mathbf{B} = \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix}$. Suppose $\mathbf{K}$ satisfies $(\mathbf{A} - \mathbf{BK})^\top(\mathbf{A} - \mathbf{BK}) - \mathbf{I} \prec 0$. Expanding this inequality and neglecting the terms that contain Δt^2 (by the assumption that Δt is sufficient small), we have

$$-\Delta t(\mathbf{A}^\top + \mathbf{K}^\top \mathbf{B}^\top + \mathbf{A} + \mathbf{BK}) \prec 0,$$

or

$$\mathbf{A}^\top + \mathbf{K}^\top \mathbf{B}^\top + \mathbf{A} + \mathbf{BK} \succ 0.$$

Writing the matrices out, we have

$$\begin{bmatrix} 2\frac{R}{L} & 0 \\ 0 & 2\frac{R}{L} \end{bmatrix}^\top + \mathbf{K}^\top \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix}^\top + \begin{bmatrix} \frac{\sqrt{2}}{L} & 0 \\ 0 & \frac{\sqrt{2}E}{L} \end{bmatrix} \mathbf{K} \succ 0$$

Since $L > 0$, we obtain

$$\begin{bmatrix} 2R & 0 \\ 0 & 2R \end{bmatrix}^\top + \mathbf{K}^\top \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2}E \end{bmatrix}^\top + \begin{bmatrix} \sqrt{2} & 0 \\ 0 & \sqrt{2}E \end{bmatrix} \mathbf{K} \succ 0.$$

Now we note that the left hand side expression is a monotonic in terms of R. Namely, if $\mathbf{K}$ satisfies this inequality for R_1, then $\mathbf{K}$ also satisfies this inequality for any $R_2 > R_1$. ■

A.4 Proof of Theorem 3

We have two major lemmas that establish the existence of a safe and stable K.

Lemma 2 (Conditions for a safe & stable system)**.** *Given* $\mathbf{M} \in \mathbb{R}^{p\times p}$*,* $\mathbf{y} \in \mathbb{R}^p$*, and* $\mathbf{z} \in \mathbb{R}^p$*, where* $\|\mathbf{y}\|_2 \leq \|\mathbf{z}\|_2$*, if* $\mathbf{M}^\top \mathbf{y} = \lambda \mathbf{y}$*,* $\lambda_{\min}(\mathbf{M} + \mathbf{M}^\top) \geq \lambda$*, and* $\mathbf{M} + \mathbf{M}^\top \succeq 0$*, then*

$$\mathbf{z}^\top \mathbf{M}\mathbf{z} - \mathbf{z}^\top \mathbf{M}\mathbf{y} \geq 0.$$

Lemma 3 (Existence of a safe & stable $\mathbf{K}$)**.** *Given* $\mathbf{A} \in \mathbb{R}^{p\times p}, \mathbf{B} \in \mathbb{R}^{p\times 1}$*. Given* $\mathbf{x}^*$ *and* $\mathbf{u}^*$*, satisfying* $\mathbf{A}\mathbf{x}^* + \mathbf{B}\mathbf{u}^* = 0$*, if* $\mathbf{A} + \mathbf{A}^\top \prec 0$*, and* $\mathbf{A}^{-1}\mathbf{B} \neq 0$*, then there exists a* $\mathbf{K} \in \mathbb{R}^{1\times p}$*, such that the closed-loop system* $\mathbf{N} \coloneqq \mathbf{A} - \mathbf{B}\mathbf{K}$ *has the properties* $\mathbf{N}^\top \mathbf{x}^* = \lambda \mathbf{x}^*$*,* $\lambda_{\max}(\mathbf{N} + N^\top) \leq \lambda$*, and* $\mathbf{N} + \mathbf{N}^\top \preceq 0$*.*

We can apply these lemmas to prove **Theorem** 3 as follows. Since we have $\mathbf{x}^*, \mathbf{u}^*$ satisfying (2.16), then the closed-loop dynamics of (2.2) with control $\mathbf{u} = \mathbf{u}^* - \mathbf{K}(\mathbf{x} - \mathbf{x}^*)$ can be equivalently represented as $\dot{\mathbf{x}} = -\bar{\mathbf{A}}(\mathbf{x} - \mathbf{x}^*)$, where $\bar{\mathbf{A}} = -(\mathbf{A} - \mathbf{B}\mathbf{K})$. For safety of this equivalent system we require (2.14) to hold for all $\mathbf{x} \in \partial\mathcal{S}$, with $f(\mathbf{x}) = -\bar{\mathbf{A}}(\mathbf{x} - \mathbf{x}^*), g(\mathbf{x}) = 0$. This is equivalent to the inequality

$$\mathbf{x}^\top \bar{\mathbf{A}}\mathbf{x} - \mathbf{x}^\top \bar{\mathbf{A}}\mathbf{x}^* \geq 0 \;\; \forall \;\; \mathbf{x} \in \partial\mathcal{S}.$$

We can show this inequality holds using **Lemma** 2, where $\mathbf{M} = \bar{\mathbf{A}}$, $\mathbf{z} = \mathbf{x}$ and $\mathbf{y} = \mathbf{x}^*$, if we can find a $\mathbf{K}$ that satisfies the following criteria:

$$(\mathbf{x}^*)^\top (\mathbf{A} - \mathbf{B}\mathbf{K}) = (\mathbf{x}^*)^\top \lambda, \tag{A.7a}$$

$$\lambda_{\max}((\mathbf{A} - \mathbf{B}\mathbf{K}) + (\mathbf{A} - \mathbf{B}\mathbf{K})^\top) \leq \lambda, \tag{A.7b}$$

$$(\mathbf{A} - \mathbf{B}\mathbf{K}) + (\mathbf{A} - \mathbf{B}\mathbf{K})^\top \prec 0, \tag{A.7c}$$

where $\lambda \in \mathbb{R}$ is the eigenvalue associated with $\mathbf{x}^*$. With **Lemma** 3, we can see that for our system a $\mathbf{K}$ satisfying (A.7) always exists, as the $\mathbf{A}$ matrix of (2.2) meets the SDP condition, $\mathbf{A} + \mathbf{A}^\top \prec 0$, and the system is controllable.

Once these two lemmas are established we can conclude that there exists a safe linear feedback controller such that there always exists at least one feasible solution for the safety filter (QP) with our system (2.2).

Proof of Lemma 2. We start by noting that $\mathbf{z}^\top \mathbf{M}\mathbf{z} - \mathbf{z}^\top \mathbf{M}\mathbf{y} \geq 0$ is equivalent to $(\mathbf{z} - \mathbf{y})^\top \mathbf{M}(\mathbf{z} - \mathbf{y}) \geq \mathbf{y}^\top \mathbf{M}\mathbf{y} - \mathbf{y}^\top \mathbf{M}\mathbf{z}$. Since $\mathbf{y}$ is a left eigenvector of $\mathbf{M}$ with eigenvalue λ this inequality is equivalent to

$$(\mathbf{z}-\mathbf{y})^\top \mathbf{M}(\mathbf{z}-\mathbf{y}) \geq \lambda \mathbf{y}^\top \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{z}.$$

Noting that $(\mathbf{z}-\mathbf{y})^\top \mathbf{M}(\mathbf{z}-\mathbf{y}) = (\mathbf{z}-\mathbf{y})^\top \frac{\mathbf{M}+\mathbf{M}^\top}{2}(\mathbf{z}-\mathbf{y})$, we have

$$(\mathbf{z}-\mathbf{y})^\top \frac{\mathbf{M}+\mathbf{M}^\top}{2}(\mathbf{z}-\mathbf{y}) \geq \lambda \mathbf{y}^\top \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{z}.$$

We define $\mathbf{v}_1, \mathbf{v}_2$ and λ_1, λ_2 to be the eigenvectors and eigenvalues of $\frac{\mathbf{M}+\mathbf{M}^\top}{2}$, and note that we can pick $\mathbf{v}_1, \mathbf{v}_2$ to form an orthonormal basis as $\frac{\mathbf{M}+\mathbf{M}^\top}{2}$ is real and symmetric. We define $c_1 = (\mathbf{z}-\mathbf{y})^\top \mathbf{v}_1$, $c_2 = (\mathbf{z}-\mathbf{y})^\top \mathbf{v}_2$ such that $(\mathbf{z}-\mathbf{y}) = c_1\mathbf{v}_1 + c_2\mathbf{v}_2$ and our inequality is equal to

$$\lambda_1 c_1^2 + \lambda_2 c_2^2 \geq \lambda \mathbf{y}^\top \mathbf{y} - \lambda \mathbf{y}^\top \mathbf{z}$$

Rearranging $(\mathbf{z}-\mathbf{y}) = c_1\mathbf{v}_1 + c_2\mathbf{v}_2$ we have $\mathbf{z} = \mathbf{y} + c_1\mathbf{v}_1 + c_2\mathbf{v}_2$ which we substitute in for the remaining $\mathbf{z}$ to get

$$\lambda_1 c_1^2 + \lambda_2 c_2^2 \geq -\lambda(\mathbf{y}^\top(c_1\mathbf{v}_1 + c_2\mathbf{v}_2)).$$

Using our assumption that $\lambda_1, \lambda_2 \geq \frac{\lambda}{2}$ we can see that the left-hand side of the inequality is lower bounded as

$$\lambda_1 c_1^2 + \lambda_2 c_2^2 \geq \frac{\lambda}{2}(c_1^2 + c_2^2) \geq -\lambda(\mathbf{y}^\top(c_1\mathbf{v}_1 + c_2\mathbf{v}_2)),$$

which is equivalent to showing that

$$c_1^2 + c_2^2 \geq -2(c_1\mathbf{y}^\top\mathbf{v}_1 + c_2\mathbf{y}^\top\mathbf{v}_2),$$

as $\lambda_1 c_1^2, \lambda_2 c_2^2 \geq 0$ ($\mathbf{M}+\mathbf{M}^\top \succ 0 \implies \lambda_1, \lambda_2 > 0$). This can be shown using **Lemma** 4, with $\zeta = \mathbf{z}, \gamma = \mathbf{y}$, and $\{\mathbf{w}_1, \mathbf{w}_2\} = \{\mathbf{v}_1, \mathbf{v}_2\}$, as by assumption $\|\mathbf{z}\|_2 \geq \|\mathbf{y}\|_2$. ∎

Lemma 4. *Given an orthonormal basis* $\{\mathbf{w}_1, \mathbf{w}_2\}$*, and* $\gamma, \zeta \in \mathbb{R}^2$ *such that* $\|\zeta\|_2 \geq \|\gamma\|_2$*,*

$$c_1^2 + c_2^2 \geq -2(c_1 b + c_2 d),$$

where $c_1 = a - b$*,* $c_2 = c - d$*, and* $a = \zeta^\top\mathbf{w}_1, b = \gamma^\top\mathbf{w}_1, c = \zeta^\top\mathbf{w}_2, d = \gamma^\top\mathbf{w}_2$*.*

Proof. Plugging the definitions for c_1 and c_2 into the inequality and expanding we have

$$(a^2 - 2ab + b^2) + (c^2 - 2cd + d^2) \geq -2(ab - b^2 + cd - d^2).$$

Noting that $a^2 + c^2 = \|\zeta\|_2^2$ and $b^2 + d^2 = \|\gamma\|_2^2$, as $\mathbf{w}_1 \perp \mathbf{w}_2$, this reduces to

$$\|\zeta\|_2^2 + \|\gamma\|_2^2 \geq 2\|\gamma\|_2^2,$$

which is true as we assume $\|\zeta\|_2 \geq \|\gamma\|_2$. ∎

Because of space constraints, we only show the proof for Lemma 3 when the control input is scalar (i.e. $u \in \mathbb{R}^1$). The proof for the vector case is considerably more involved and not needed for the results we present.

Proof of Lemma 3. First we note that $\mathbf{x}^* = \mathbf{A}^{-1}\mathbf{B}u^*$ such that $\mathbf{N}^\top \mathbf{x}^* = \lambda \mathbf{x}^* \iff u^*\mathbf{B}^\top\mathbf{A}^{-\top}(\mathbf{A} - \mathbf{B}\mathbf{K}) = \lambda u^*\mathbf{B}^\top\mathbf{A}^{-\top}$. Assuming $u^* \neq 0$ and dividing it out we get

$$\mathbf{B}^\top\mathbf{A}^{-\top}(\mathbf{A} - \mathbf{B}\mathbf{K}) = \lambda\mathbf{B}^\top\mathbf{A}^{-\top}.$$

Solving for $\mathbf{K}$, noting that $\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}$ is a scalar, we get

$$\mathbf{K} = \frac{1}{\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}}\left(\mathbf{B}^\top\mathbf{A}^{-\top}(\mathbf{A} - \lambda\mathbf{I})\right).$$

Plugging this into our expression for $\mathbf{N}$ we have

$$\mathbf{A} - \mathbf{B}\mathbf{K} = \mathbf{A} - \frac{1}{\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}}\left(\mathbf{B}\mathbf{B}^\top\mathbf{A}^{-\top}(\mathbf{A} - \lambda\mathbf{I})\right),$$

which can be substituted into the inequality $\mathbf{A} - \mathbf{B}\mathbf{K} + (\mathbf{A} - \mathbf{B}\mathbf{K})^\top \prec \lambda\mathbf{I}$ to get

$$\mathbf{M}_1 - \lambda\mathbf{M}_2 \prec 0,$$

where

$$\begin{aligned}
\mathbf{M}_1 =& \mathbf{A} + \mathbf{A}^\top \\
&- \frac{1}{\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}}\left(\mathbf{B}\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A} + \mathbf{A}^\top\mathbf{A}^{-1}\mathbf{B}\mathbf{B}^\top\right), \\
\mathbf{M}_2 =& \mathbf{I} - \frac{1}{\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}}\left(\mathbf{B}\mathbf{B}^\top\mathbf{A}^{-\top} + \mathbf{A}^{-1}\mathbf{B}\mathbf{B}^\top\right).
\end{aligned}$$

We will show the following such that we can always choose $\lambda < 0$, sufficiently small, to have $\mathbf{M}_1 - \lambda \mathbf{M}_2 \prec 0$: $\mathbf{M}_1 \preceq 0$, with a single zero eigenvalue associated with the eigenvector $\mathbf{z}$ (i.e. $\mathbf{M}_1\mathbf{z} = 0$), and $\mathbf{z}^\top \mathbf{M}_2 \mathbf{z} < 0$.

First, we show that $\mathbf{M}_1 \preceq 0$, using the fact that if $\mathbf{X} \prec 0$, then $\mathbf{P}^\top \mathbf{X} \mathbf{P} \preceq 0 \ \forall \ \mathbf{P}$. To construct this form consider $\mathbf{P} = \mathbf{I} - \frac{\mathbf{A}^{-1}\mathbf{B}\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A}}{\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B}}$. By direct computation, $\mathbf{M}_1 = \mathbf{P}^\top(\mathbf{A} + \mathbf{A}^\top)\mathbf{P}$. This implies that $\mathbf{M}_1 \preceq 0$ as $(\mathbf{A} + \mathbf{A}^\top) \prec 0$.

Now, we can show that $\mathbf{M}_1$ has only one zero eigenvalue using this same $\mathbf{P}^\top\mathbf{X}\mathbf{P}$ form. If $\mathbf{v}$ is a zero eigenvector of $\mathbf{M}_1$ then we have that $\mathbf{v}^\top\mathbf{P}^\top(\mathbf{A}+\mathbf{A}^\top)\mathbf{P}\mathbf{v} = 0$. Since we assume $\mathbf{A} + \mathbf{A}^\top \prec 0$, then this implies that $\mathbf{P}\mathbf{v} = 0$. Note that $\mathbf{P}$ is the sum of a rank n matrix, $\mathbf{I}$, and a rank 1 matrix, which implies that $\mathbf{P}$ is at least rank $n-1$ and can have at most one zero eigenvalue. We see that this zero eigenvalue is associated with the eigenvector $\mathbf{z} = \mathbf{A}^{-1}\mathbf{B}$ by direct computation; $\mathbf{M}_1\mathbf{z} = 0$ using the equality $\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B} = \mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}$, which holds as $\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B}$ is a scalar.

Lastly, we can show that $\mathbf{z}^\top\mathbf{M}_2\mathbf{z} < 0$ as

$$
\begin{aligned}
\mathbf{z}^\top\mathbf{M}_2\mathbf{z} =& (\mathbf{A}^{-1}\mathbf{B})^\top\mathbf{M}_2(\mathbf{A}^{-1}\mathbf{B}) \\
=& \mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A}^{-1}\mathbf{B} \\
& - \frac{1}{\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B}}\big((\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B})\mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A}^{-1}\mathbf{B} \\
& \qquad\qquad + \mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A}^{-1}\mathbf{B}(\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B})\big) \\
\stackrel{(a)}{=}& - \mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{A}^{-1}\mathbf{B} \\
=& - \|\mathbf{A}^{-1}\mathbf{B}\|_2^2 < 0
\end{aligned}
$$

where (a) follows from $\mathbf{B}^\top\mathbf{A}^{-1}\mathbf{B} = \mathbf{B}^\top\mathbf{A}^{-\top}\mathbf{B}$. ■

A.5 Alternate Proof of Theorem 4

We can write the quantities in (2.19) just as a function of the current using the fact that at equilibrium the steady state voltage is

$$\overline{\mathbf{V}}_{\mathrm{dq}} = \mathbf{E}_{\mathrm{dq}} - L\mathbf{A}\bar{\mathbf{I}}_{\mathrm{dq}}.$$

Substituting this into (2.19), we have

$$\overline{P} = \frac{3}{2}\bar{\mathbf{I}}_{\mathrm{dq}}^{\top}(\mathbf{E}_{\mathrm{dq}} - L\mathbf{A}\bar{\mathbf{I}}_{\mathrm{dq}}) \tag{A.8a}$$

$$\overline{Q} = \frac{3}{2}\bar{\mathbf{I}}_{\mathrm{dq}}^{\top}\mathbf{J}(\mathbf{E}_{\mathrm{dq}} - L\mathbf{A}\bar{\mathbf{I}}_{\mathrm{dq}}) \tag{A.8b}$$

$$V_{\mathrm{dq}}^2 = (\mathbf{E}_{\mathrm{dq}} - L\mathbf{A}\bar{\mathbf{I}}_{\mathrm{dq}})^{\top}(\mathbf{E}_{\mathrm{dq}} - L\mathbf{A}\bar{\mathbf{I}}_{\mathrm{dq}}). \tag{A.8c}$$

We note that these equations can be represented in the form of a quadratic expression (i.e. $x^{\top}Nx$) as

$$\overline{P} = \frac{3}{2}\begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix}^{\top}\begin{bmatrix}-L\mathbf{A}^{\top} & 0 \\ \mathbf{E}_{\mathrm{dq}}^{\top} & 0\end{bmatrix}\begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix} \tag{A.9}$$

$$\overline{Q} = \frac{3}{2}\begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix}^{\top}\begin{bmatrix}-L\mathbf{A}^{\top}\mathbf{J} & 0 \\ \mathbf{E}_{\mathrm{dq}}^{\top}\mathbf{J} & 0\end{bmatrix}\begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix} \tag{A.10}$$

$$V_{\mathrm{dq}}^2 = \begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix}^{\top}\begin{bmatrix}L^2\mathbf{A}^{\top}\mathbf{A} & -L\mathbf{A}^{\top}\mathbf{E}_{\mathrm{dq}} \\ -L\mathbf{E}_{\mathrm{dq}}^{\top}\mathbf{A} & \mathbf{E}_{\mathrm{dq}}^{\top}\mathbf{E}_{\mathrm{dq}}\end{bmatrix}\begin{bmatrix}\bar{\mathbf{I}}_{\mathrm{dq}} \\ 1\end{bmatrix}. \tag{A.11}$$

Because of the particular form of these equations, Theorem 1 follows directly from the following Lemma.

Lemma 5. *Given a function $f(x) = (\hat{x}^{\top}N\hat{x}, \hat{x}^{\top}M\hat{x})$, with $\hat{x} = (x, 1)$ and $N, M \in \{L \mid L \in \mathbb{R}^3, L_{1,1} = L_{2,2}, L_{1,2} = -L_{2,1}\}$, and the unit ball $\mathcal{X} := \{x \mid \|x\|_2^2 \leq 1\}$, the set $\mathcal{Z} := \{f(x) \mid x \in \mathcal{X}\}$ is convex.*

Proof. We note that we can take the trace of both terms of $f(x)$ and apply the cyclic property of trace such that we equivalently have

$$f(x) = (\mathrm{tr}(N\hat{x}\hat{x}^{\top}), \mathrm{tr}(M\hat{x}\hat{x}^{\top})).$$

Now let W be the set of all symmetric 3×3 matrices $w \in \mathbb{S}^3$ where: $w \succeq 0$, $w_{1,1}+w_{2,2} \leq 1$, and $w_{3,3} = 1$. We can relax the $\hat{x}\hat{x}^{\top}$ terms of $\mathcal{Z}$ into w such that we have a new set

$$\hat{\mathcal{Z}} = \{(\mathrm{tr}(Nw), \mathrm{tr}(Mw) \mid w \in W\}.$$

It clear that $\mathcal{Z} \subset \hat{\mathcal{Z}}$. $\hat{\mathcal{Z}}$ is clearly convex in w as trace and matrix multiplication are linear operations and the set W is convex.

Expanding the trace terms, we note the constraints on N, M and the symmetry of w ($w_{1,2} = w_{2,1}$). Because the upper-left 2×2 blocks of N and M are defined to be skew-symmetric in their off-diagonals (e.g. $N_{1,2} = -N_{2,1}$), the cross terms cancel: $N_{1,2}w_{2,1} + N_{2,1}w_{1,2} = 0$. Therefore, the traces simplify entirely to:

$$(\operatorname{tr}(Nw), \operatorname{tr}(Mw)) = \begin{bmatrix} a_1 & a_1 & b_1^\top \\ a_2 & a_2 & b_2^\top \end{bmatrix} \begin{bmatrix} w_{1,1} \\ w_{2,2} \\ w_{1,3} \\ w_{2,3} \end{bmatrix} + \begin{bmatrix} c_1 \\ c_2 \end{bmatrix},$$

where $a_1 = N_{1,1} = N_{2,2}, a_2 = M_{1,1} = M_{2,2}, b_1 = (N_{1,3} + N_{3,1}, N_{2,3} + N_{3,2}), b_2 = (M_{1,3} + M_{3,1}, M_{2,3} + M_{3,2}), c_1 = N_{3,3}, c_2 = M_{3,3}$. Similarly we have

$$(\operatorname{tr}(N\hat{x}\hat{x}^\top), \operatorname{tr}(M\hat{x}\hat{x}^\top)) = \begin{bmatrix} a_1 & a_1 & b_1^\top \\ a_2 & a_2 & b_2^\top \end{bmatrix} \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} c_1 \\ c_2 \end{bmatrix},$$

Now we will show that for any $w \in W$ we can always pick a $x \in \mathcal{X}$ such that $(\operatorname{tr}(N\hat{x}\hat{x}^\top), \operatorname{tr}(M\hat{x}\hat{x}^\top)) = (\operatorname{tr}(Nw), \operatorname{tr}(Mw))$. Noting our constraints on the set W, and that in general for an SDP matrix, $A \succeq 0$, the off-diagonals are bound such that $A_{ij}^2 \leq (A_{ii}A_{jj})$ we have that $w_{1,3}^2 \leq w_{1,1}w_{3,3}$ and $w_{2,3}^2 \leq w_{2,2}w_{3,3}$. This provides the bound $w_{1,3}^2 + w_{2,3}^2 \leq w_{11}w_{3,3} + w_{2,2}w_{3,3}$. We have $w_{3,3} = 1$, so this becomes $w_{1,3}^2 + w_{2,3}^2 \leq w_{11} + w_{2,2} \leq 1$. In the following, for brevity we define $a := (a_1, a_2)$, $B := \begin{bmatrix} b_1^\top \\ b_2^\top \end{bmatrix}$, $c := (c_1, c_2)$, such that

$$(\operatorname{tr}(Nw), \operatorname{tr}(Mw)) = a \cdot (w_{1,1} + w_{2,2}) + B \begin{bmatrix} w_{1,3} \\ w_{2,3} \end{bmatrix} + c, \tag{A.12}$$

$$(\operatorname{tr}(N\hat{x}\hat{x}^\top), \operatorname{tr}(M\hat{x}\hat{x}^\top)) = a \cdot (x_1^2 + x_2^2) + B \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + c. \tag{A.13}$$

For the last step we want to show that for a given w there exists an $x \in \mathcal{X}$ such that

$$a \cdot (w_{1,1} + w_{2,2}) + B \cdot (w_{1,3}, w_{2,3}) = a \cdot (x_1^2 + x_2^2) + B \cdot (x_1, x_2) \tag{A.14}$$

To prove (A.14), we seek a vector, denoted by $\tilde{w}$, shifted from w such that the output of (A.12) remains invariant, but the relation $\tilde{w}_{1,1} + \tilde{w}_{2,2} = \tilde{w}_{1,3}^2 + \tilde{w}_{2,3}^2$ is satisfied. We can isolate two trivial edge cases:

1. B **is singular (not full rank):** We can set $\tilde{w}_{1,1} = w_{1,1}$ and $\tilde{w}_{2,2} = w_{2,2}$, and shift $(\tilde{w}_{1,3}, \tilde{w}_{2,3})$ along the null space of B without changing the trace output. By scaling along this null space, we can easily satisfy $\tilde{w}_{1,3}^2 + \tilde{w}_{2,3}^2 = w_{1,1} + w_{2,2}$.

2. $a = \mathbf{0}$**:** The trace output becomes completely independent of the diagonal terms. We can therefore arbitrarily decrease $\tilde{w}_{1,1} + \tilde{w}_{2,2}$ until $\tilde{w}_{1,1} + \tilde{w}_{2,2} = w_{1,3}^2 + w_{2,3}^2$ without impacting (A.12).

Therefore, the remainder of the proof focuses on the non-trivial case where B is full rank (i.e. invertible) and $a \neq \mathbf{0}$.

We start by determining how we can move the terms of w such that the value of (A.12) does not change. Solving for $(w_{1,3}, w_{2,3})$ from

$$a \cdot (w_{1,1} + w_{2,2}) + B \cdot (w_{1,3}, w_{2,3}) = 0. \tag{A.15}$$

we find $(w_{1,3}, w_{2,3}) = -B^{-1}a \cdot (w_{1,1} + w_{2,2})$. This implies that

$$a \cdot (w_{1,1} + w_{2,2}) + B \cdot (w_{1,3}, w_{2,3}) = a(w_{1,1} + w_{2,2} + \Delta) + B \cdot \big((w_{1,3}, w_{2,3}) + -B^{-1}a \cdot \Delta\big)$$

for any $\Delta \in \mathbb{R}$. We want to find a Δ such that

$$w_{1,1} + w_{2,2} + \Delta = \|(w_{1,3}, w_{2,3}) + -B^{-1}a \cdot \Delta\|_2^2 \tag{A.16}$$

and $-(w_{1,1} + w_{2,2}) \leq \Delta \leq 1 - (w_{1,1} + w_{2,2})$.

We can show this using our constraints $w_{1,3}^2 \leq w_{1,1}, w_{2,3}^2 \leq w_{2,2}$ that come from $w \succeq 0, w_{3,3} = 1$. We start by expanding and rearranging (A.16) to get

$$\|B^{-1}a\|_2^2 \cdot \Delta^2 - (1 + 2(w_{1,3}, w_{2,3})^\top B^{-1}a) \cdot \Delta - (w_{1,1} + w_{2,2}) - \|(w_{1,3}, w_{2,3})\|_2^2 = 0$$

which is a quadratic in terms of Δ. The solutions to this quadratic are of the form

$$\frac{-g \pm \sqrt{g^2 - 4fh}}{2f},$$

where $f = \|B^{-1}a\|_2^2$, $g = -(1+2(w_{1,3}, w_{2,3})^\top B^{-1}a)$, and $h = -(w_{1,1}+w_{2,2})-\|(w_{1,3}, w_{2,3})\|_2^2$. We note that $f \geq 0$ and $-1 \leq h \leq 0$. This implies that $-f \leq fh \leq 0$ and the term within the square root, $g^2 - 4fh$, is always positive and greater than or equal to g^2 so it has a real valued solution. Now we look at the case where the $\pm$ is negative. With this we know that the numerator $-g - \sqrt{g^2 - 4fh} \leq 0$ for any g, and with the denominator $2f \geq 0$ we can conclude that for one solution $\Delta \leq 0$. We call this the negative solution Δ^-. Now to show the lower bound on Δ, we assume that $\Delta^- < -(w_{1,1} - w_{2,2})$. This implies that

$$w_{1,1} + w_{2,2} + \Delta^- = \|(w_{1,3}, w_{2,3}) + -B^{-1}a \cdot \Delta^-\|_2^2 < 0,$$

which is a contradiction. Therefore, $\Delta \geq -(w_{1,1} + w_{2,2})$. This implies that for any given w, there exists an Δ^-, and $\tilde{w}$ where $\tilde{w}_{1,1} + \tilde{w}_{2,2} = w_{1,1} + w_{2,2} + \Delta^-$, and $(\tilde{w}_{1,3}, \tilde{w}_{2,3}) = (w_{1,3}, w_{2,3}) - B^{-1}a \cdot \Delta^-$, such that $\tilde{w}_{1,1} + \tilde{w}_{2,2} = \tilde{w}_{1,3}^2 + \tilde{w}_{2,3}^2$. This implies that we can satisfy (A.14) by selecting $x_1 = \tilde{w}_{1,3}$ and $x_2 = \tilde{w}_{2,3}$, such that $x_1^2 + x_2^2 = \tilde{w}_{1,1} + \tilde{w}_{2,2} = w_{1,1} + w_{2,2} + \Delta^-$. Finally, because we established that $\Delta^- \leq 0$ and we are given $w_{1,1} + w_{2,2} \leq 1$, it strictly follows that $x_1^2 + x_2^2 \leq 1$. Thus, the constructed x is guaranteed to lie within the unit ball $\mathcal{X}$. Being able to construct such an $x \in \mathcal{X}$ for any $w \in W$ establishes that $\mathcal{Z} = \hat{\mathcal{Z}}$ and we can conclude that the set $\mathcal{Z}$ is convex. ∎

Appendix B

SYSTEM EIGENVALUE SENSITIVITIES

B.1 Proof of Device Impedance Sensitivity

The closed-loop whole-system admittance matrix $\hat{\mathbf{Y}}(s)$ is defined as the relationship between the grid network admittance $\mathbf{Y}_{\text{net}}(s)$ and the aggregated device impedance $\mathbf{Z}_{\text{dev}}(s)$:

$$\hat{\mathbf{Y}} = (I + \mathbf{Y}_{\text{net}}\mathbf{Z}_{\text{dev}})^{-1}\mathbf{Y}_{\text{net}} \tag{B.1}$$

To find the first-order sensitivity of the system with respect to a local power flow parameter x that only affects the device control state, we first invert the whole-system admittance matrix to find the whole-system impedance $\hat{\mathbf{Z}}$:

$$\hat{\mathbf{Y}}^{-1} = \mathbf{Y}_{\text{net}}^{-1}(I + \mathbf{Y}_{\text{net}}\mathbf{Z}_{\text{dev}}) \tag{B.2}$$

Distributing the inverse network admittance yields:

$$\hat{\mathbf{Y}}^{-1} = \mathbf{Y}_{\text{net}}^{-1} + \mathbf{Y}_{\text{net}}^{-1}\mathbf{Y}_{\text{net}}\mathbf{Z}_{\text{dev}} = \mathbf{Y}_{\text{net}}^{-1} + \mathbf{Z}_{\text{dev}} \tag{B.3}$$

Taking the partial derivative of both sides with respect to the power flow parameter x:

$$\frac{\partial \hat{\mathbf{Y}}^{-1}}{\partial x} = \frac{\partial \mathbf{Y}_{\text{net}}^{-1}}{\partial x} + \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial x} \tag{B.4}$$

Because the physical network topology $\mathbf{Y}_{\text{net}}$ is independent of the device power flow set-points, its partial derivative with respect to x is zero. Therefore, the sensitivity of the inverse whole-system admittance reduces exactly to the sensitivity of the device impedance:

$$\frac{\partial \hat{\mathbf{Y}}^{-1}}{\partial x} = \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial x} \tag{B.5}$$

B.2 Proof of Network Admittance Sensitivity

Following the derivation in Appendix B.1, the inverse of the whole-system admittance matrix is analytically defined as the sum of the inverse network admittance and the device impedance:

$$\hat{\mathbf{Y}}^{-1} = \mathbf{Y}_{\text{net}}^{-1} + \mathbf{Z}_{\text{dev}} \tag{B.6}$$

To find the sensitivity of the system with respect to a change in the physical network topology, we take the partial derivative of both sides with respect to a specific line susceptance parameter B_{ij}:

$$\frac{\partial \hat{\mathbf{Y}}^{-1}}{\partial B_{ij}} = \frac{\partial \mathbf{Y}_{\text{net}}^{-1}}{\partial B_{ij}} + \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial B_{ij}} \tag{B.7}$$

Applying the standard matrix derivative identity for an inverse matrix ($\frac{\partial \mathbf{A}^{-1}}{\partial x} = -\mathbf{A}^{-1}\frac{\partial \mathbf{A}}{\partial x}\mathbf{A}^{-1}$), the sensitivity of the inverse whole-system admittance evaluates to:

$$\frac{\partial \hat{\mathbf{Y}}^{-1}}{\partial B_{ij}} = -\mathbf{Y}_{\text{net}}^{-1}\frac{\partial \mathbf{Y}_{\text{net}}}{\partial B_{ij}}\mathbf{Y}_{\text{net}}^{-1} + \frac{\partial \mathbf{Z}_{\text{dev}}}{\partial B_{ij}}. \tag{B.8}$$

VITA

Trager Joswig-Jones was born and raised in Pullman, Washington. He received his PhD, Master's, and Bachelor's degrees in Electrical Engineering from the University of Washington in 2026, 2023, and 2021, respectively. Throughout his studies he interned at Schweitzer Engineering Laboratories (Pullman, WA), Micron Technology (Boise, ID), General Motors (Milford, MI), and the National Renewable Energy Laboratory (Golden, CO). During his studies he was a recipient of and funded in part by the Clean Energy Institute Graduate Fellowship, the Grainger Fellowship, and the Dave and Debbie Galloway Term Fellowship. Outside of power systems research, he enjoys climbing and setting routes at the UW climbing gym, as well as exploring and enjoying nature as time and bodily injury permits.