UNIVERSITY OF CALIFORNIA SAN DIEGO

Safe Extremum Seeking with Applications to Particle Accelerators

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Engineering Sciences (Mechanical Engineering)

by

Alan Williams

Committee in charge:

Miroslav Krstić, Chair
Jorge Cortés
Raymond de Callafon
Jorge Poveda
Alexander Scheinker

2024

The Dissertation of Alan Williams is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

I dedicate this dissertation to my parents who fostered my interest in science and engineering.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I extend my deepest gratitude to Miroslav, whose exceptional depth of knowledge and expertise has been a guiding light throughout my research journey. His willingness to respond to my queries, even late into the night, has been immensely helpful and reflective of his commitment to my academic growth and to the research conducted during the course of my studies. Additionally, his skill in communicating complex ideas, a quality I have also greatly admired in his nonlinear system class, has significantly influenced my own ability to articulate my research. His guidance has not only shaped my work but also inspired me to strive for clarity and precision in conducting research.

I would also like to thank Alex Scheinker for his unwavering support and dedication. His investment of time and effort into my work has been a cornerstone of my progress. His constant encouragement of my ideas and his ability to foster a supportive and constructive research environment has been invaluable. The genuine concern and interest he has shown for my well-being and academic success have always reassured me that my best interests are at the heart of his mentorship. His example has been a powerful motivator, and I am grateful for the positive influence he has had on both my academic and personal development.

Thanks to my friends for your support and all of the memories: the Five Guys, Nikhil, Niall, Nick Ramey, Nick PV, Tim, Claire, Andre, Tenley, Pete, Shephard, David, Matt, Ev, Katya, and the UCSD club tennis squad, to name a few. I would not have been able to find the motivation throughout the years without you all.

Also thanks to Emma, Mom, and Dad for your continual support, praise, and love through all of it.

To my committee, thank you for your time, effort, and interest in reviewing my work.

Chapter 2 has been submitted for publication with the title “Practical Safe Extremum Seeking with Assignable Rate of Attractivity to the Safe Set” A. Williams, M. Krstic, A. Scheinker. Chapter 2 also contains a partial adaptation of the work contained in the conference paper “Practically Safe Extremum Seeking” A. Williams, M. Krstic, A. Scheinker, presented at

the IEEE Conference on Decision and Control 2022. The dissertation author was the primary investigator and author of these two papers.

Chapter 3 has been submitted for publication of the material under the title "Safety-Filtered Extremum Seeking with Unknown CBFs" A. Williams, M. Krstic, A. Scheinker. Chapter 3 also contains a partial adaptation of the work contained in the conference paper "Semi-Global Practical Extremum Seeking with Practical Safety" A. Williams, M. Krstic, A. Scheinker, presented at the IEEE Conference on Decision and Control 2023. The dissertation author was the primary investigator and author of these two papers.

Chapter 4 is a reprint of the material as it appears in IEEE Transactions on Control System Technology 2024, under the title "Experimental Safe Extremum Seeking for Accelerators" A. Williams, A. Scheinker, E. Huang, C. Taylor, M. Krstic. The dissertation author was the primary investigator and author of this paper.

Chapter 1 contains snippets of introduction sections and literature gathered from the previously mentioned publications (or submitted publications).

VITA

2017	B. S. in Mechanical Engineering, University of Florida
2020	M. S. in Engineering Sciences (Mechanical Engineering), University of California San Diego
2018-2024	Fellow in the University of California San Diego - Los Alamos National Laboratory Research and Education Collaboration
2024	Ph.D. in Engineering Sciences (Mechanical Engineering), University of California San Diego

PUBLICATIONS

A. Williams, M. Krstic, A. Scheinker, "Practical Safe Extremum Seeking with Assignable Rate of Attractivity to the Safe Set" submitted for publication 2024.

A. Williams, M. Krstic, A. Scheinker, "Safety-Filtered Extremum Seeking with Unknown CBFs" submitted for publication 2024.

A. Williams, A. Scheinker, "Estimation and Control of Particle Accelerators in Simulation using Latent Space Tuning" Accepted in the Proceedings of the 15th International Particle Accelerator Conference, IPAC-2024.

A. Williams, A. Scheinker, E. Huang, C. Taylor, M. Krstic, "Experimental Safe Extremum Seeking for Accelerators" IEEE Transactions on Control System Technology 2024.

A. Williams, M. Krstic and A. Scheinker, "Semi-Global Practical Extremum Seeking with Practical Safety" 61st IEEE Conference on Decision and Control 2023.

A. Williams, M. Krstic and A. Scheinker, "Practically Safe Extremum Seeking" 61st IEEE Conference on Decision and Control 2022.

E. Toler, J. Coleman, J. Koglin, M. McKerns, H. Morris, A. Scheinker, A. Williams, "Unwrapping Image Data Using Adaptive Machine Learning" North American Particle Accelerator Conference 2022.

A. Williams, A. Scheinker, "Safe Extremum Seeking for Real Time Optimization of Particle Accelerators" North American Particle Accelerator Conference 2022.

Compton, W., Curtin, M., Vogt, W., Scheinker, A. and Williams, A., 2022. "Deep Reinforcement Learning for Active Structure Stabilization" Data Science in Engineering, Volume 9 (pp. 5-14).

Springer, Cham.

ABSTRACT OF THE DISSERTATION

Safe Extremum Seeking with Applications to Particle Accelerators

by

Alan Williams

Doctor of Philosophy in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2024

Miroslav Krstić, Chair

This dissertation introduces a form of extremum seeking, traditionally employed for optimizing unknown objective functions, now adapted to accommodate an unknown yet measurable constraint. We consider the constraint to be a safety metric, which is maintained, practically, throughout the optimization process. We demonstrate that our approach can ensure safety violations be made arbitrarily small, parallel to how classical extremum seeking controllers achieve stability near optimal points. The power of this algorithm is particularly underscored in its application to particle accelerator systems, shown in several examples, where safety is critical to prevent substantial financial losses and operational downtime. This work broadens the scope of extremum seeking methods and establishing an approach for integrating safety considerations

into optimization processes, useful in situations where balancing optimal performance with stringent safety requirements is essential.

Chapter 1

Introduction

1.1 Particle Accelerator Control: A Motivation for Safe Extremum Seeking

Particle accelerators are sophisticated machines designed to accelerate charged particles, such as electrons or protons, to high speeds and contain them in well-defined beams. At the heart of numerous scientific endeavors, these accelerators are pivotal in pushing the boundaries of fundamental physics, allowing researchers to delve into the structure of matter at the smallest scales. Beyond their scientific applications, particle accelerators have practical uses in medical therapy, such as in cancer treatment through radiation therapy, and in industry, where they are used for materials processing and inspection.

Operating particle accelerators presents a significant challenge due to their complexity and the precision required in controlling the particle beam. Steering the beam accurately is critical; any deviation from its intended path can have severe consequences, including damage to the accelerator's components or in the worse case, burning a hole in the side of the beam pipe containing the vacuum. Additionally, if the beam is not adequately contained in a tightly formed bunch and are allowed to form a substantial halo around the center of the beam, stray particles may be gradually lost as they contact the walls of the beampipe and other components. Such incidents not only necessitate costly and time-consuming repairs but also pose safety risks. Improperly controlled beams can lead to the generation of unwanted radiation, rendering parts of

the accelerator radioactive and thus complicating maintenance efforts. The high level of precision needed in beam control highlights the importance of safety in both protecting personnel and ensuring the accelerator's integrity.

Particle accelerators are marvels of engineering, comprised of complicated assemblies that can stretch for miles and incorporate thousands of components, each playing a pivotal role in the accelerator's operation. The complexity is compounded by numerous unknown parameters that drift over time, affected by factors such as aging, thermal effects, and other environmental changes, posing significant challenges to maintaining optimal performance and beam quality. Many of these parameters, such as the initial distribution of the particles at the beginning of the accelerator, not only drift, but they are difficult to measure. This time-varying behavior necessitates continuous adjustment of various control parameters to ensure the highest beam quality, and often a critical part of the job of a beam operator is the manual hand tuning of these control parameters. Beam quality can be assessed using metrics such as beam intensity, energy uniformity, and spatial distribution. Optimizing these metrics are often crucial for maximizing the accelerator's performance, ensuring the success of its applications.

Extremum seeking control has been recently discovered as a useful tool in controlling accelerators, given the often unknown, nonlinear, and time-varying nature of the relationship between the various metrics describing the beam and the control parameters. Extremum seeking (ES) control is a model-free, adaptive control strategy designed to optimize the performance of dynamic systems in real-time by automatically adjusting control parameters to converge to the extremum (maximum or minimum) of an unknown objective function directly from measured output data. Many ES algorithms, and all ES algorithms applied to particle accelerator systems in the literature, are designed to solve the following problem

$$\min_{\theta} J(\theta),$$

where θ is a vector containing the control parameters to be adjusted and J is a measurement of

the objective function describing the performance of the system.

Historically, extremum seeking (ES) control has found diverse applications in accelerator technology, where the performance metric J directly relates to with measurements of interest. These objectives include beam loss, beam trajectory error, and the energy output of free electron lasers—a type of light-producing accelerator that relies on the acceleration of electrons. These goals can also be performed simultaneously, by summing two or more objectives to be minimized (or maximized) into a single quantity J. In these applications, the parameter θ represents various adjustable elements, including the strengths of quadrupole magnets (which confine the beam within the vacuum of the beam pipe), the strengths of kicker magnets (which are used to steer the beam), and the phase and amplitude values of radio-frequency (RF) cavities (which control the electromagnetic waves that accelerate the particles). This framework allows for the tuning of accelerator components to optimize performance, underlining the versatility and impact of ES in enhancing accelerator operations.

But why "Safe" ES? Consider the scenario where we aim to adjust control parameters but must also ensure the system's safety. Essentially, this involves automatically optimizing metrics of interest, all while maintaining the accelerator's safety over time, as gauged by specific signals we identify as indicators of safety. For instance, take the challenge of achieving a specific beam shape, where beam loss serves as a safety indicator. Thus, the primary objective becomes minimizing the error between the beam's actual shape and its desired shape. In many accelerator systems, a significant halo of particles inevitably forms due to internal repulsive forces, leading to eventual particle loss as the beam reaches higher energies. We cannot minimize this perfectly to zero and so we decide that if beam loss stays below a certain threshold, we deem the system safe. This type of problem is common in many situations, not just with beam loss as a measure of safety, but in any case where a measured signal must remain above or below a specific threshold. There is almost never a known relationship between the control parameters and beam loss, as it involves complex multi-particle dynamics, and innumerable unknowable quantities such as all particle positions and velocities, which generate the beam halo. However, merely minimizing the

combination of beam loss and the primary objective does not fully address the issue if optimizing for the primary objective does not concurrently reduce beam loss. Therefore we actually need a more general tool which can solve the following constrained optimization problem:

$$\min_{\theta} J(\theta) \text{ subject to } h(\theta) \geq 0,$$

where h is a function which is designed positive when beam loss is below the threshold, and where both h and J are analytically unknown. This tool should solve (1.1) while also guaranteeing 1) the parameters are not steered outside the safe set, if they begin in a safe operating condition 2) the parameters are steered toward the safe set, if they begin in an unsafe operating condition.

This dissertation introduces an ES tool that fundamentally tackles this issue. We develop Safe ES to address this challenge, providing theoretical guarantees and demonstrating its effectiveness across various accelerator problems.

1.2 Literature Review

The literature on extremum seeking is split into two subsections - a general overview of ES and a specific look at various constrained forms of ES. Safety literature is given its own section as the form of ES developed in this dissertation draws from the safety literature, and the contributions we make can be considered as a part of the body of work on safe control. Finally, we overview the set of methods used (which include ES methods) by the accelerator community to tune and optimize particle accelerator systems.

1.2.1 Extremum Seeking

The first ever conception of an extremum seeking design was presented in a 1922 paper by Leblanc [54]. Here a mechanism was described to maximize the power transfer from a transmission line to an electric tram car, although it is unknown if the design was ever implemented. In the 1940s the ES scheme gained popularity beginning in Russia [42, 43] and then in the

1950s in the US [15, 25, 96]. After this work, ES was modified and applied in notable early applications in the 1950s and 1960s [59, 62, 66, 68] and an early analysis was conducted in 1971 [57], although it did not provide a general guarantee of stability. The early history provided here is credited to the survey [77], which also provides a wide ranging and thorough review of extremum seeking over time, covering all significant contributions and their various applications.

The first formal stability proof for ES feedback was given in [52] combining results on averaging and singular perturbation for ordinary differential equations (ODEs). Since then various studies, extensions, and improvements of the ES algorithm have been made including the development of algorithms for semiglobal convergence [92], for maps with multiple extrema [93], extensions for multivariable maps [30], continuous and discrete stochastic generalizations [56], performance improvements for parameter convergence [33, 3, 32, 65], for stabilizing open-loop unstable time-varying systems with unknown control directions [85, 86], for application to Partial Differential Equations (PDEs) [67], for fixed-time Nash equilibrium seeking in time-varying networks [70], and for many more other problems [77]. ES methods have also enabled the use of deep neural networks for time-varying systems without retraining [81], and have improved the robustness of generative convolutional neural networks beyond the span of their training data [79].

We point the reader to the survey [77] for a wide ranging and thorough review of extremum seeking over time, covering many significant contributions and various applications of ES.

1.2.2 Constrained Extremum Seeking

Known constraints: constrained versions of various kinds of ES control laws have been studied as well, including variants where the constraint is known. In the context of population games, extremum seeking methods are useful when known simplex-like constraints must be placed on optimization parameters [71]. ES algorithms with constrained inputs have also been studied in [90] where the optimization parameters must always lie within a known domain as exceeding them may cause undesirable actuator saturation. Switching frameworks have also been

developed for constrained ES [72] where the optimization parameters must lie in a known set. A form of ES is designed such that the update to the optimization parameter is constrained to a known bound [86].

Unknown constraints: cases where the constraints are unknown have also been studied, often assuming convexity. Switching type extremum seeking algorithms are also studied [20], methods employing nonlinear programming techniques in [55] to solve convex problems by smoothly switching between gradient based scheme of the constraint and objective, and an approach which solves convex problems by designing an ES law which travels toward the saddle point of a so called "modified barrier function" [53]. Authors in [27, 28] study saddle points dynamics and a Lie bracket based analysis in their work. Authors in [13] use projections of the gradient of the objective function on the direction of gradient of the constraint function to determine a control law which satisfies the constraint.

Unknown constraint handling via data-sampled ES: a number of data sampled optimization schemes have been studied which are notable. Steady state constraint satisfaction is achieved in the framework of [34] as well as SPA stability, based on the data sampled schemes of [95]. Other data sampled frameworks are presented in [45] studying a general class of optimization schemes.

1.2.3 Control Barrier Functions and Safety Filters

The literature on safe control has boomed largely as of late due to the seminal work in [9, 99]. A large number of applications and other inspired work has arisen due to the quadratic program (QP) based control barrier function (CBF) designs, from incorporating it with learning frameworks [24, 98], applications in robotic and automotive systems [97, 73, 6], designs involving observers in the presence of disturbances in the dynamics and measurements [4], and more [61, 1, 19] to name only a small set of inspired work.

Recent work has also studied input-to-state safety CBFs, by which the violation of a safe set can bounded by the magnitude of the disturbance [49]. Inverse optimal safety filters,

in the presence of disturbances, both deterministic and stochastic, of unlimited and unknown magnitude, are introduced in [50]. Designs for CBFs of high relative degree are introduced in [51]. Independently, high-order CBF conditions were also studied subsequently in [37, 103]. Higher order CBFs have been use in optimal trajectory planning, with application for autonomous vehicles [104]. Prescribed-time safety filters are designed in [2]. As an alternative to the QP based approach, approximate optimal controller are formulated, where violation of safety is a part of value function [22].

1.2.4 Accelerator Tuning and Optimization

The Nelder-Mead simplex method [39] and ML based methods have been used to tune accelerators [5]. Extremum seeking has been studied in simulation for accelerator applications [89], and was first used experimentally in an accelerator tuning problem to minimize an objective function [87]. Since then, a bounded form of ES with guaranteed limits on parameter update rates [88], has been used extensively in several accelerator applications for free electron laser energy maximization [78], electron beam trajectory control [84], real-time multi-objective optimization [82], and for beam loss minimization [83].

Recently, various machine learning (ML)-based methods have been developed for control and optimization of particle accelerator beams [12]. Bayesian optimization has become a popular tool in tuning and in some cases has been used to design safety aware tuning algorithms [48, 47, 26]. Bayesian optimization (and methods based on it), unlike ES, constructs a probabilistic estimate of the unknown functions, in the form of a Gaussian Process (GP), and determines a new point to sample based the fitted function. Neural networks (NN) have been used as surrogate models for magnet control [40] and for simulation-based optimization studies [29]. Neural networks are also being used for uncertainty aware anomaly detection to predict errant beam pulses [17], as virtual diagnostics for 4D tomographic phase space reconstructions [102], for predicting the transverse emittance of space charge dominated beams [58], and for high resolution longitudinal phase space virtual diagnostics [105]. Neural network-based deep reinforcement learning (RL)

methods have been used for accelerator control [35], and in a sample efficient manner, which trains a policy based on data at two beam lines at CERN [41].

Although many ML-based tools have been developed they all suffer major limitations when it comes to time-varying systems. If a system changes then NN, GP, and RL methods all require new data for re-training of their models in order to be applicable for accelerator control. This major limitation is overcome by adaptive model-independent methods such as ES. Adaptive machine learning frameworks combining ES with neural networks have been demonstrated to extend the use of ML for time-varying systems. ES-based adaptive ML has been demonstrated for automatically shaping the longitudinal phase space of short intense electron beams in the LCLS FEL [81] and for creating virtual 6D diagnostics of time-varying charged particle beams [76, 79].

In all the methods described above, including those using ES, the approach to safety has been some expert-based combination of setting hard bounds on allowed parameter values, adding additional terms to the cost function, and algorithm hyper-parameter tuning. The safe ES methods we demonstrate in this dissertation reduces the amount of required hyper-parameter tuning and removes the manual design of a tradeoff between safety and optimization which depends on the weights given to safety-related cost function terms relative to objective-related cost function terms. For example, in the real-time multi-objective ES optimization application in [82], while the objective was beam spot size minimization, the safety-related term in the cost function was the beam's distance from a desired reference trajectory, and there was a tradeoff between the two depending on the weights. In our approach, practical safety is always enforced.

1.3 Contributions and Organization

Chapter 2 unveils a safe ES algorithm, termed Assignably Safe Extremum Seeking (ASfES), illustrating local stability and offering a guarantee of practical safety under initial conditions close to the constrained optimizer. The term "practical safety" refers to a violation of

safety which can be made arbitrarily small by selecting appropriate design constants. The safety feature of ASfES offers a user specified "attractivity rate" to the set set, denoted as c, which dictates the temporal behavior of h. Essentially, this allows the user to control how h—the metric defining system safety—either asymptotically approaches the unsafe threshold or moves towards safety if beginning from an unsafe state, with the behavior over time defined by the user assignable time constant c. We introduce dynamic filters that, in an average and singularly perturbed sense, mimic the feedback law typical of a quadratic programming (QP) and control barrier function (CBF) based safety filter applied to a nominal extremum seeking (ES) controller. We also extend ASfES to a Newton-based version (NB-ASfES) which has both assignable attractivity to and from the set set, and an assignable nominal convergence rate of the parameter — if the attractivity rate assignment is not violated.

Chapter 3 advances this work by making small adjustments to the algorithm – which will be referred to simply as "Safe ES", differentiating it from "ASfES" – offering proofs of semiglobal practical asymptotic stability and semiglobal practical safety. The algorithms presented in Chapter 2 and 3 are different in their design in two ways. First, the algorithm in Chapter 2 contains an additional filter which provides an estimate of the inverse of the gradient of h through the use of a nonlinear Riccati filter. This elegant solution is used to handle instances in time where the estimate of the gradient of h is zero, and was also a key part of the design in Newton based ES [30]. In contrast, the Chapter 3 algorithm adopts a more straightforward, albeit practical, approach to this issue. Second, while Chapter 2's algorithm employs a smooth control law, Chapter 3's version becomes nonsmooth. The analytical techniques employed in each chapter are distinctly different, reflecting the variations in the algorithms, despite their overall similarity. We note one final distinction in the safety result we present in Chapter 3, compared to the safety result in Chapter 2. The attractivity rate (c), which we describe in the prior paragraph, is not user assignable in the algorithm presented in Chapter 3. This is because in order to achieve semiglobality, and guarantee convergence of initial conditions farther and farther away from the constrained optimum, the rate of the parameter must be slowed more and more. This therefore

does not allow the luxury of a user assignable c.

Chapter 4 directly implements Safe ES in accelerator systems, showcasing its practical applicability. We present several scenarios in which Safe ES is shown to be effective in solving challenging tuning tasks relating to safely tuning accelerator parameters in order to achieve performance while mitigating excessive risk. The first two examples are demonstrated in simulation and the third example showcases the algorithm on a real tuning experiment conducted at the linear accelerator at Los Alamos Neutron Science Center.

Acknowledgements

This chapter contains partial adaptations and snippets of the following works: A. Williams, M. Krstic, A. Scheinker, "Safety-Filtered Extremum Seeking with Unknown CBFs" Under Review 2024. A. Williams, A. Scheinker, E. Huang, C. Taylor, M. Krstic, "Experimental Safe Extremum Seeking for Accelerators" IEEE Transactions on Control System Technology 2024. A. Williams, M. Krstic and A. Scheinker, "Semi-Global Practical Extremum Seeking with Practical Safety" 61st IEEE Conference on Decision and Control 2023. A. Williams, M. Krstic and A. Scheinker, "Practically Safe Extremum Seeking" 61st IEEE Conference on Decision and Control 2022. The dissertation author was the primary author in these publications.

Chapter 2

Practically Safe ES with Assignable Attractivity to the Safe Set

We present Assignably Safe Extremum Seeking (ASfES), an algorithm designed to minimize a measured objective function while maintaining a measured metric of safety (a control barrier function or CBF) be positive in a practical sense. We ensure that for trajectories with safe initial conditions, the violation of safety can be made arbitrarily small with appropriately chosen design constants. We also guarantee an assignable "attractivity" rate: from unsafe initial conditions, the trajectories approach the safe set, in the sense of the measured CBF, at a rate no slower than a user-assigned rate. Similarly, from safe initial conditions, the trajectories approach the unsafe set, in the sense of the CBF, no faster than the assigned attractivity rate. The feature of assignable attractivity is not present in the semiglobal version of safe extremum seeking (studied in Chapter 3), where the semiglobality of convergence is achieved by slowing the adaptation. We also demonstrate local convergence of the parameter to a neighborhood of the minimum of the objective function constrained to the safe set. The ASfES algorithm and analysis are multivariable, but we also extend the algorithm to a Newton-Based ASfES scheme (NB-ASfES) which we show is only useful in the scalar case. The proven properties of the designs are illustrated through simulation examples.

2.1 Introduction

2.1.1 Motivation

We introduce an algorithm for Assignably Safe Extremum Seeking (ASfES), where both a lower bound on the rate of convergence *to* the safe set and an upper bound on the rate of convergence *towards* the unsafe set are user-assignable.

The algorithm combines traditional Extremum Seeking (ES) methods [52] [11] with the quadratic program (QP)-based control barrier function (CBF) safety approach of [10]. This work is essentially concerned with finding an algorithm which solves the following constrained optimization problem:

$$\min_{\theta} J(\theta) \text{ subject to } h(\theta) \geq 0. \tag{2.1}$$

The objective function J should be minimized but safety, represented by an unknown, yet measured, function $h(\theta)$, should be maintained by keeping h positive over the entire course of time. The idea of "practical safety" has been described through the use of an inequality,

$$h(\theta(t)) \geq h(\theta(0))e^{-ct} + O(\epsilon) \text{ for all } t \in [0, \infty), \tag{2.2}$$

where $c > 0$ is value we term the "attractivity" rate and $O(\epsilon)$ is a violation of safety (in the worst case $O(\epsilon)$ is negative) which may be made small through appropriate choices of design constants. This inequality and the value c describes how quickly a trajectory, which starts safe — $\theta(0) \in \{h(\theta) \geq 0\}$ — is allowed to approach the practical, marginally safe boundary $h = O(\epsilon)$. But, for trajectories which start unsafe — $\theta(0) \in \{h(\theta) < 0\}$ — it describes how fast (at least) a trajectory converges to the safe set. It is therefore desirable to be able to assign the value c. In our previous work [101], the value c is fundamentally unknown to the user, and not assignable. Because of the semiglobal design in [101], the parameter $\theta(t)$ must travel *slowly*, in order to achieve semiglobality, with a small adaptation gain. With a small adaptation gain, the estimates of the unknown gradients of the optimization problem can be computed accurately such that

convergence can be achieved from any initial condition — but this sacrifices the assignability of the parameter c.

The design presented in this work allows the user to assign the value of c, achieving practical safety with an assignable attractivity rate, giving more control by the user to specify the dynamics of h. In addition, we also demonstrate local convergence of the parameter θ to a neighborhood of the constrained minimizer of the objective function on the safe set.

2.1.2 Safety-Filtered ES Applied on Average

The QP-CBF safety filter introduced in [9] is an obvious possibility for enforcing constraints in ES. However, three problems arise.

First, the QP-CBF safety filter is nonsmooth. This makes it, in its original form, incompatible with the classical averaging theory, in which the system's vector field needs to be at least continuously differentiable for the averaging theory to be applicable in the analysis. To overcome this challenge, we introduce some important modifications in the QP-CBF safety filter. In place of the nonsmooth maximum (max) operation, we introduce a smoothed approximation of that operation. More importantly, in place of the division by a term like $|L_g h|^2$ in the conventional QP-CBF safety filter, we introduce a scalar Riccati differential equation, which generates the reciprocal of $|L_g h|^2$ dynamically, rather than by an algebraic division. To eliminate the need for an analytically intractable computation of an average of a vector field with a maximum operation, we introduce a filter of an estimate of the gradient of the map being optimized. All these innovations in the QP-CBF approach are aimed not at improvements in the safeguarding capabilities of the safety filter but are made for the sake of making the averaging analysis possible.

Second, classical ES teaches us that the gradient of the unknown map can be estimated using perturbations. If the map's gradient, used in the optimization, can be allowed to be unknown and only estimated using measurements and perturbations, why would the barrier function, namely, the constraint or safety characterization, not be allowed to also be functionally unknown, and estimated using perturbations and measurements? This is indeed what we do. We

only measure the value of the barrier function, we demodulate it using the same perturbation signal as the one used to estimate the gradient of the map being optimized, and we generate an estimate of the barrier function gradient. We then filter that estimate so that the hard-to-integrate maximum operation does not act on the terms that need to be averaged.

Third, our design is motivated by a safety analysis of the time-averaged and filterless version of the system. Our engineering interest is in the safety of the system actually implemented, rather than of the system's averaged and singularly perturbed reduced version. So, our analysis needs to account for the minor violation in safety which arises due to the design being motivated by the approximated model. We spend considerable effort capturing the effect of those approximation errors on the lower bound of the CBF. Since the CBF can experience minor dips into the negative range, i.e., minor violations of the safety boundary, what we achieve is, strictly speaking, not safety but *practical safety*, a term we choose in the analogy with "practical asymptotic stability," where the attractivity is imperfect. The minor violations of the safety boundary are not fatal and can be anticipated and prevented by adding a small negative constant bias in the measured value of the barrier function. However, the amount of such a negative bias needed is not known a priori because the functional form of the CBF is considered unknown.

2.1.3 Results and Contribution

The analysis we provide is based upon classical averaging and singular perturbation techniques. The QP-CBF based safety filters presented in [10] are inherently nonsmooth, due to the term $\max\{x,0\}$ appearing in the safety filter term. Therefore our design contains a smooth approximation of this quantity, $\max_\delta\{x\}$ — a term which is inherently more conservative with respect to safety, and makes our algorithm more conservative than the design presented in [101], yielding convergence to a neighborhood of the constrained equilibrium with a bias favoring safety (with conservativeness increasing with large δ). Furthermore, our results are fundamentally local in nature relying on more restrictive assumption of J and h than those used in [101], due to the analysis techniques used. The design also relies on a dynamic estimate of the quantity $||G_h||^{-2}$

by use of a Ricatti filter, reminiscent of [30].

What is gained by the ASfES algorithm presented here over its semiglobal cousin [101] is that the user has enhanced control over the temporal behavior of h. Namely, we show that the value of c, the attractivity rate, in (2.2) can be assigned by the user. We achieve: 1) local, practical asymptotic convergence to the constrained minimum of the objective function on the safe set 2) practical convergence to the safe set at least as fast as the assignable attractivity, for initial conditions which start unsafe 3) an approach towards the unsafe set, practically, no faster than the assigned attractivity for initial conditions which start safe.

Furthermore, we also present an extension to a Newton-Based ASfES (NB-ASfES) scheme which we show should only be used when considering a scalar parameter. NB-ASfES shares all of the features of ASfES with the additional feature that the nominal convergence rate of the parameter, k, is also assigned. Therefore we achieve an assignable rate of convergence of the parameter only if the assigned safety attractivity is not violated — otherwise the parameter rate is slowed from its assigned rate in favor of maintaining the assigned safe attractivity rate.

This chapter elucidates a tradeoff: one can either use slow adaptation and achieve semiglobality [101] or, as shown here, assign the attractivity rate locally.

The conference version of this work [100] studied the ASfES algorithm (see [100] for a intuitive derivation of the dynamics), but with analysis only given in 1 dimension. This work presents an extensive n dimensional analysis. Other work in [101] studies a similar algorithm, which is semiglobal in nature, but has nonassignable attractivity. Practical safety is a feature shared by ASfES and the algorithm in [101]. Namely, that for an arbitrarily small violation of safety (the $O(\epsilon)$ term in (2.2)), there exist design constants which guarantee it *for all time*.

Organization: first we introduce the ASfES design and the main results in Sections 2.2 and 2.3, then we perform the averaging and convergence analysis in n dimensions in Section 2.4. In Section 2.5, the safety of the system is studied with the use of singular perturbation techniques, showing that the attractivity rate, c, can be assigned a prior by the user. Finally, Section 2.7 presents simulations of the algorithms for 1 and 2 dimensions.

2.2 Algorithm

2.2.1 Assignably Safe Extremum Seeking

We will first introduce the set of differential equations describing the ASfES algorithm in n dimensions:

$$\dot{\hat{\theta}} = -kG_J + \gamma \max_\delta\{kG_J^T G_h - c\eta_h\}G_h \tag{2.3}$$

$$\dot{G}_J = -\omega_{\mathrm{f}} G_J + \omega_{\mathrm{f}}(J(\hat{\theta}(t) + S(t)) - \eta_J)M(t) \tag{2.4}$$

$$\dot{\eta}_J = -\omega_{\mathrm{f}}\eta_J + \omega_{\mathrm{f}} J(\hat{\theta}(t) + S(t)) \tag{2.5}$$

$$\dot{G}_h = -\omega_{\mathrm{f}} G_h + \omega_{\mathrm{f}}(h(\hat{\theta}(t) + S(t)) - \eta_h)M(t) \tag{2.6}$$

$$\dot{\eta}_h = -\omega_{\mathrm{f}}\eta_h + \omega_{\mathrm{f}} h(\hat{\theta}(t) + S(t)) \tag{2.7}$$

$$\dot{\gamma} = \omega_{\mathrm{f}}\gamma(1 - \gamma||G_h||^2) \tag{2.8}$$

where the state variables $\hat{\theta}, G_J, G_h \in \mathbb{R}^n$ and $\eta_J, \eta_h, \gamma \in \mathbb{R}$ so overall the dimension of the system is $3n+3$. The map is evaluated at θ, defined by

$$\theta(t) = \hat{\theta}(t) + S(t)\ . \tag{2.9}$$

The integer n denotes the number of parameters one wishes to optimize over. The design coefficients are $k, c, \delta, \omega_{\mathrm{f}} \in \mathbb{R}_{>0}$. The perturbation signal S and demodulation signal M are given by $S_i(t) = a\sin(\omega_i t)$ and $M_i(t) = \frac{2}{a}\sin(\omega_i t)$ the standard signals found in classical ES [11]. The functions $J, h : \mathbb{R}^n \mapsto \mathbb{R}$ are scalar valued functions defined over the optimization variables θ. To make the classical averaging theorem applicable, the algorithm must have a continuously differential right-hand side. Because the standard maximum operation is nonsmooth in the QP-CBF formulation we approximate it by the function $\max_\delta : \mathbb{R} \mapsto \mathbb{R}_{>0}$ given by

$$\max_\delta\{x\} := \frac{1}{2}\left(x + \sqrt{x^2 + \delta}\right) \approx \max\{x, 0\}, \tag{2.10}$$

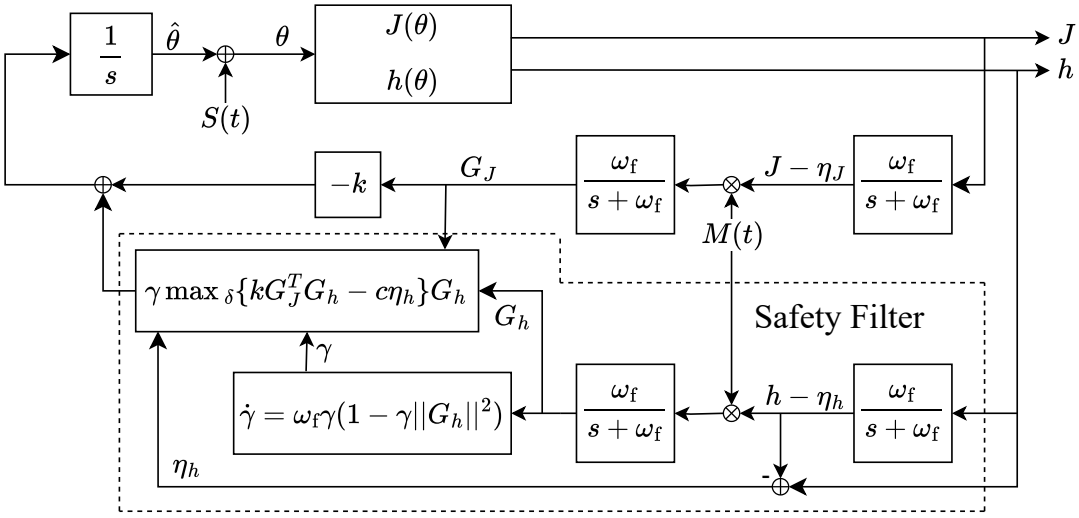

Figure 2.1. A block diagram of the ASfES algorithm. Removing the "Safety Filter" block recovers the classical extremum seeking algorithm.

for a small parameter $\delta > 0$. Note that $\max_\delta\{x\} > \max\{x, 0\}$ for all $x \in \mathbb{R}$.

2.2.2 Notation and Coordinate Transformations

The variable θ, is transformed several times before averaging is performed, and then again to the equilibrium of the constrained system. We summarize the various coordinates in Table 2.1 for clarity, using "param." to mean "parameter", "est." to mean "estimate" or "estimation", "err." to mean "error", and "avg." to mean "average".

Table 2.1. Transformations of the parameter θ.

Var.	Definition	Description
$\theta(t)$	n/a	optimization param. (input to the map)
$\hat{\theta}(t)$	$= \theta(t) - S(t)$	unconstrained param. est.
$\tilde{\theta}(t)$	$= \hat{\theta}(t) - \theta^*$	unconstrained param. est. err.
$\tilde{\theta}^{\mathrm{a}}(t)$	avg. of $\tilde{\theta}(t)$	unconstrained avg. param est. err.
$\tilde{\theta}^{\mathrm{a,e}}$	eqm. of $\tilde{\theta}^{\mathrm{a}}(t)$	eqm. of unconstrained avg. param. est. err.
$\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}(t)$	$= \tilde{\theta}^{\mathrm{a}}(t) - \tilde{\theta}^{\mathrm{a,e}}$	constrained avg. param. est. err.

In unconstrained extremum seeking, the equilibrium $\tilde{\theta}^{\mathrm{a,e}} = 0$. This is not true for the dynamics we present. The final transformation in Table 2.1 defines a new state with equilibrium of zero $\tilde{\theta}^{\mathrm{a,e}}_{\mathrm{c}} = 0$. The rest of the states in the algorithm, G_J, η_J, G_h, η_h, and γ, undergo fewer

transformations than θ.

We use the variable 'x' to describe the full state of the solutions or equilibrium:

$$x = [\theta; G_J; \eta_J; G_h; \eta_h; \gamma], \tag{2.11}$$

$$\hat{x} = [\hat{\theta}; G_J; \eta_J; G_h; \eta_h; \gamma], \tag{2.12}$$

$$\tilde{x} = [\tilde{\theta}; G_J; \eta_J; G_h; \eta_h; \gamma], \tag{2.13}$$

$$x^{\mathrm{a}} = [\tilde{\theta}^{\mathrm{a}}; G_J^{\mathrm{a}}; \eta_J^{\mathrm{a}}; G_h^{\mathrm{a}}; \eta_h^{\mathrm{a}}; \gamma^{\mathrm{a}}], \tag{2.14}$$

$$x^{\mathrm{a,e}} = [\tilde{\theta}^{\mathrm{a,e}}; G_J^{\mathrm{a,e}}; \eta_J^{\mathrm{a,e}}; G_h^{\mathrm{a,e}}; \eta_h^{\mathrm{a,e}}; \gamma^{\mathrm{a,e}}], \tag{2.15}$$

$$x_c^{\mathrm{a}} = [\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}}; G_{J,\mathrm{c}}^{\mathrm{a}}; \eta_{h,\mathrm{c}}^{\mathrm{a}}; G_{h,\mathrm{c}}^{\mathrm{a}}; \gamma_{\mathrm{c}}^{\mathrm{a}}; \eta_{J,\mathrm{c}}^{\mathrm{a}}]. \tag{2.16}$$

Notation: the superscript "a" always denotes a variable defined from an averaged system, and the superscript "e" always denotes an equilibrium value. For example, $x^{\mathrm{a,e}}$ is the equilibrium of the state x^{a}, which is the averaged version of $\tilde{x}$ (shown in the next section). Here, "averaged" is in the sense in the standard averaging theory [44]. The subscript "c" relate to the "constrained" dynamics of ASfES, shown later to be defined as $x_c^{\mathrm{a}} := x^{\mathrm{a}} - x^{\mathrm{a,e}}$. Note that the ordering of the state variables have also been changed in x_c^{a}, which will help identify block structures in a linearization procedure later on. Also, when stacking vectors $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ we denote $z = [x; y] := [x_1, ..., x_n, y_1, ..., y_m]^T \in \mathbb{R}^{n+m}$. We define the signum function as: $\mathrm{sgn}(x) := -1$ for $x < 0, 0$ for $x = 0, 1$ for $x > 0$. We define the unit step function as: $u(x) := 0$ for $x < 0, 0.5$ for $x = 0, 1$ for $x > 0$.

2.3 Main Results

We consider the analysis of ASfES under the following assumptions.

Assumption 2.1. *With the symmetric Hessian matrix $H \succ 0$ and unconstrained minimum $\theta^* \in \mathbb{R}^n$,*

the objective function is unknown and has the quadratic form

$$J(\theta) = J^* + \frac{1}{2}(\theta - \theta^*)^T H(\theta - \theta^*). \tag{2.17}$$

Assumption 2.2. *With $h_0 \in \mathbb{R}$ and $h_1 \in \mathbb{R}^n_{\neq 0}$, the barrier function is unknown and has the linear form*

$$h(\theta) = h_0 + h_1^T(\theta - \theta^*). \tag{2.18}$$

Since the results are local, a linear form of the barrier function is assumed essentially with no loss of generality. In (2.18), $h_0 \geq 0$ implies the unconstrained minimum of J is safe, and $h_0 < 0$ implies the unconstrained minimum is unsafe.

Assumption 2.3. *The design constants are chosen as $\omega_{\mathrm{f}}, \omega_i, \delta, a, k, c > 0$, where ω_i/ω_j are rational with frequencies ω_i chosen such that $\omega_i \neq \omega_j$ and $\omega_i + \omega_j \neq \omega_k$ for distinct i, j, and k.*

The first main result of this chapter concerns convergence of the parameter to the (constrained) minimizer of the objective function J on the safe set.

Theorem 2.1 (Local Convergence to the Safe Optimum)**.** *Under Assumptions 2.1–2.3, there exists positive constants ρ, ω^* such that if $||\tilde{x}(0) - x^{\mathrm{a,e}}|| < \rho$ then for all $\omega \in [\omega^*, \infty)$,*

$$\limsup_{t \to \infty} J(\theta(t)) = J_s^* + O(1/\omega + a + \delta), \tag{2.19}$$

where J_s^* is the minimum of $J(\theta)$ on the safe set $C = \{\theta : h(\theta) \geq 0\}$.

The local region $||\tilde{x}(0) - x^{\mathrm{a,e}}|| < \rho$ describes initial conditions which are close to the constrained minimizer. This implies $\hat{\theta}(0)$ must be close to $\theta^* + \tilde{\theta}^{\mathrm{a,e}}$ — a quantity which is the global minimizer (θ^*) and a large offset, describing the shift of the parameter into the safe set. Additionally the filtered estimates $G_J, \eta_J, G_h, \eta_h, \gamma$ must also lie close to the true values.

In the second main result, we show that the violation of safety of the original system is of order $O(1/\omega + 1/\omega_{\mathrm{f}} + a)$ and the decay of the time dependent term has a time constant c, which

is specified by the user and contained in the algorithm dynamics (2.3) - (2.8).

Theorem 2.2 (Practical Safety with Assignable Attractivity Rate)**.** *Under Assumptions 2.1–2.3, there exists positive constants* $\rho, \omega_{\mathrm{f}}^*, \omega^*$ *such that if* $||\tilde{x}(0) - x^{\mathrm{a,e}}|| < \rho$*, then for all* $\omega_{\mathrm{f}} \in [\omega_{\mathrm{f}}^*, \infty)$ *and* $\omega \in [\omega^*, \infty)$*,*

$$h(\theta(t)) \geq h(\theta(0))\mathrm{e}^{-ct} + O(1/\omega_{\mathrm{f}} + 1/\omega + a). \tag{2.20}$$

2.4 Convergence

The next subsections present the main steps in deriving Theorem 2.1. First an average system is derived, with a unique equilibrium, and then linearization is performed to conclude the convergence result.

2.4.1 Deriving the Average System

With the change of variables, introducing $\hat{\theta}$ and $\tilde{\theta}$, we have the relations

$$\tilde{\theta}(t) = \hat{\theta}(t) - \theta^* \tag{2.21}$$

$$\theta(t) = \hat{\theta}(t) + S(t) = \tilde{\theta}(t) + S(t) + \theta^* \tag{2.22}$$

$$\theta(t) - \theta^* = \tilde{\theta}(t) + S(t) \tag{2.23}$$

where

$$S_i(t) = a\sin(\omega_i t), \tag{2.24}$$

and with the time transformation $\tau = \omega t$ we have

$$\omega \frac{d\tilde{x}}{d\tau} = \begin{bmatrix} -kG_J + \gamma \max_\delta \{kG_J^T G_h - c\eta_h\} G_h \\ -\omega_f G_J + \\ \quad \omega_f (J(\tilde{\theta} + S(\frac{\tau}{\omega}) + \theta^*) - \eta_J) M(\frac{\tau}{\omega}) \\ -\omega_f \eta_J + \omega_f J(\tilde{\theta} + S(\frac{\tau}{\omega}) + \theta^*) \\ -\omega_f G_h + \\ \quad \omega_f (h(\tilde{\theta} + S(\frac{\tau}{\omega}) + \theta^*) - \eta_h) M(\frac{\tau}{\omega}) \\ -\omega_f \eta_h + \omega_f h(\tilde{\theta} + S(\frac{\tau}{\omega}) + \theta^*) \\ \omega_f \gamma (1 - \gamma ||G_h||^2) \end{bmatrix} \tag{2.25}$$

$$= f(\tilde{x}, \tau). \tag{2.26}$$

Now the system is in the correct form to perform averaging. We express the the signals S and M with frequencies scaled by ω as $S_i(\frac{\tau}{\omega}) = a \sin(\omega_i' \tau)$ and $M_i(\tau) = \frac{2}{a} \sin(\omega_i' \tau)$, with $\omega_i := \omega \omega_i'$. We compute the average of the right-hand side given by the formula (2.26), taking the period to be Π given as

$$\Pi = 2\pi \times \text{LCM} \left\{ \frac{1}{\omega_i'} \right\}, \quad i \in \{1, 2, ..., n\}, \tag{2.27}$$

where LCM denotes the least common multiple. We compute the average system as

$$f^{\text{a}}(x^{\text{a}}) := \frac{1}{\Pi} \int_0^\Pi f(x^{\text{a}}, \tau) d\tau, \tag{2.28}$$

$$x^{\text{a}} := [\tilde{\theta}^{\text{a}}; G_J^{\text{a}}; \eta_J^{\text{a}}; G_h^{\text{a}}; \eta_h^{\text{a}}; \gamma^{\text{a}}]. \tag{2.29}$$

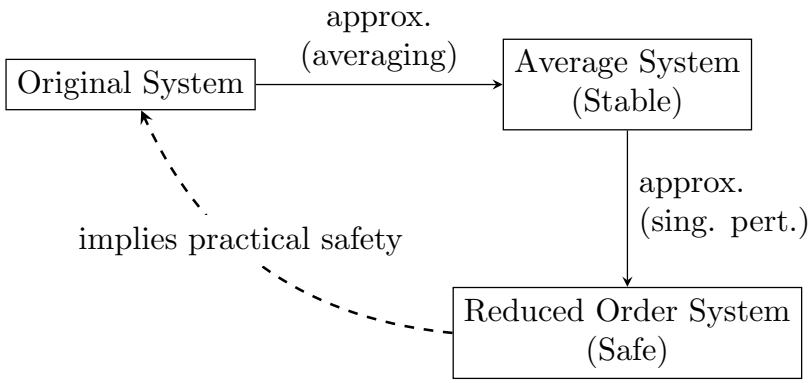

Figure 2.2. Road map of the analysis steps for ASfES. The approximation between the original system and average is performed via linearization of the average system (Section 2.4) and the approximation between the average system and the reduced order system is performed via singular perturbation (Section 2.5) by taking ω_{f} large.

Using classical averaging [44], we arrive at the average system

$$\omega \frac{dx^{\mathrm{a}}}{d\tau} = \begin{bmatrix} -kG_J^{\mathrm{a}} + \gamma^{\mathrm{a}} \max_{\delta}\{kG_J^{\mathrm{a}T} G_h^{\mathrm{a}} - c\eta_h^{\mathrm{a}}\} G_h^{\mathrm{a}} \\ -\omega_{\mathrm{f}} G_J^{\mathrm{a}} + \omega_{\mathrm{f}} H\tilde{\theta}^{\mathrm{a}} \\ -\omega_{\mathrm{f}}\eta_J^{\mathrm{a}} + \omega_{\mathrm{f}}(J(\tilde{\theta}^{\mathrm{a}} + \theta^*) + \frac{a^2}{4}\mathrm{Tr}(H)) \\ -\omega_{\mathrm{f}} G_h^{\mathrm{a}} + \omega_{\mathrm{f}} h_1 \\ -\omega_{\mathrm{f}}\eta_h^{\mathrm{a}} + \omega_{\mathrm{f}} h(\tilde{\theta}^{\mathrm{a}} + \theta^*) \\ \omega_{\mathrm{f}}\gamma^{\mathrm{a}}(1 - \gamma^{\mathrm{a}}||G_h^{\mathrm{a}}||^2) \end{bmatrix} \tag{2.30}$$

$$= f^{\mathrm{a}}(x^{\mathrm{a}}). \tag{2.31}$$

2.4.2 Equilibrium of the Average System

We find the equilibrium of the average system (2.31), taking the positive equilibrium point of γ^{a}. Immediately we calculate the following equilibria

$$\eta_J^{\mathrm{a,e}} = J^* + \frac{1}{2}\tilde{\theta}^{\mathrm{a,e}T} H \tilde{\theta}^{\mathrm{a,e}} + \frac{a^2}{4}\mathrm{Tr}(H) \tag{2.32}$$

$$G_h^{\mathrm{a,e}} = h_1 \tag{2.33}$$

$$\gamma^{\mathrm{a,e}} = \frac{1}{||h_1||^2}. \tag{2.34}$$

for three components of the system state. And we have the following relations:

$$0 = -kG_J^{\mathrm{a,e}} + \gamma^{\mathrm{a,e}} \max_\delta \{kG_J^{\mathrm{a,e}T} G_h^{\mathrm{a,e}} - c\eta_h^{\mathrm{a,e}}\} G_h^{\mathrm{a,e}} \tag{2.35}$$

$$G_J^{\mathrm{a,e}} = H\tilde{\theta}^{\mathrm{a,e}} \tag{2.36}$$

$$\tilde{\theta}^{\mathrm{a,e}} = H^{-1} G_J^{\mathrm{a,e}} \tag{2.37}$$

$$\eta_h^{\mathrm{a,e}} = h(\theta^* + \tilde{\theta}^{\mathrm{a,e}}) = h_0 + h_1^T \tilde{\theta}^{\mathrm{a,e}}. \tag{2.38}$$

From (2.35) we have the following quadratic vector equation in $G_J^{\mathrm{a,e}}$

$$k||h_1||^2 G_J^{\mathrm{a,e}} = h_1 \max_\delta (kG_J^{\mathrm{a,e}T} h_1 - ch_0 - ch_1^T H^{-1} G_J^{\mathrm{a,e}}). \tag{2.39}$$

To solve (2.39), we use the following fact: for some $\nu > 0$ we have

$$k||h_1||^2 G_J^{\mathrm{a,e}} = h_1 \nu \tag{2.40}$$

because $\max_\delta(x) > 0$. Writing $G_J^{\mathrm{a,e}} = h_1 \frac{\nu}{k||h_1||^2}$, and using the definition of $\max_\delta$, we achieve the expansion of (2.39):

$$h_1\nu = h_1 \frac{1}{2}(\nu - ch_0 - d\nu) + h_1 \frac{1}{2}\sqrt{(\nu - ch_0 - d\nu)^2 + \delta} \tag{2.41}$$

letting the quantity $d > 0$ (because $H > 0$) be

$$d = \frac{c}{k||h_1||^2} h_1^T H^{-1} h_1. \tag{2.42}$$

After solving for the positive solution of the quadratic equation in ν, the (unique) equilibrium of the average system can be written as below

$$\tilde{\theta}^{\mathrm{a,e}} = H^{-1} G_J^{\mathrm{a,e}}, \tag{2.43}$$

$$G_J^{\mathrm{a,e}} = \frac{|h_0|}{2h_1^T H^{-1} h_1} \left(-\mathrm{sgn}(h_0) + \sqrt{1 + \delta \frac{d}{c^2 h_0^2}} \right) h_1, \tag{2.44}$$

$$\begin{aligned} \eta_J^{\mathrm{a,e}} = {} & J^* + \frac{a^2}{4} \mathrm{Tr}(H) + \\ & \frac{h_0^2}{8h_1^T H^{-1} h_1} \left(-\mathrm{sgn}(h_0) + \sqrt{1 + \delta \frac{d}{c^2 h_0^2}} \right)^2, \end{aligned} \tag{2.45}$$

$$G_h^{\mathrm{a,e}} = h_1, \tag{2.46}$$

$$\eta_h^{\mathrm{a,e}} = \frac{|h_0|}{2} \left(\mathrm{sgn}(h_0) + \sqrt{1 + \delta \frac{d}{c^2 h_0^2}} \right), \tag{2.47}$$

$$\gamma^{\mathrm{a,e}} = \frac{1}{||h_1||^2}, . \tag{2.48}$$

We see that no matter the sign of h_0, this equilibrium always lies in the interior of the safe region because $\eta_h^{\mathrm{a,e}} = h(\theta^* + \tilde{\theta}^{\mathrm{a,e}}) > 0$. Also, we can think of $\eta_J^{\mathrm{a,e}}$ as the equilibrium value of the filtered measurement of the objective function. It settles to J^* with a small error due to the dithering amplitude, and a possibly large offset if $h_0 < 0$, resulting from the safety dynamics constraining the equilibrium to the safe set.

2.4.3 Linearization

Using the in error variables of the average system (to the constrained equilibrium) we define the following coordinates with new ordering as

$$\begin{aligned} x_c^{\mathrm{a}}(t) := [&\tilde{\theta}^{\mathrm{a}}(t) - \tilde{\theta}^{\mathrm{a,e}}; G_J^{\mathrm{a}}(t) - G_J^{\mathrm{a,e}}; \eta_h^{\mathrm{a}}(t) - \eta_h^{\mathrm{a,e}}; \\ &G_h^{\mathrm{a}}(t) - G_h^{\mathrm{a,e}}; \gamma^{\mathrm{a}}(t) - \gamma^{\mathrm{a,e}}; \eta_J^{\mathrm{a}}(t) - \eta_J^{\mathrm{a,e}}] \\ = [&\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}}(t); G_{J,\mathrm{c}}^{\mathrm{a}}(t); \eta_{h,\mathrm{c}}^{\mathrm{a}}(t); G_{h,\mathrm{c}}^{\mathrm{a}}(t); \gamma_{\mathrm{c}}^{\mathrm{a}}(t); \eta_{J,\mathrm{c}}^{\mathrm{a}}(t)], \end{aligned} \tag{2.49}$$

and the system

$$\omega \frac{dx_c^{\mathrm{a}}}{d\tau} = g(x_c^{\mathrm{a}}), \tag{2.50}$$

from (2.31). We have changed the ordering of the dynamic equations to help identify block triangular structures and have therefore have defined the system under with a new function g, instead of f^{a}. For example, $\omega \dot{G}_{h,\mathrm{c}}^{\mathrm{a}} = g_3(x_c^{\mathrm{a}}) = -\omega_{\mathrm{f}}(G_{h,\mathrm{c}}^{\mathrm{a}} + G_h^{\mathrm{a,e}}) + \omega_{\mathrm{f}} h_1$, with a slight abuse of notation as g_i denotes the the vector valued function associated with the dynamics of the ith vectored valued components of $x_c^{\mathrm{a}}(t)$ listed in (2.49).

The Jacobian matrix of (2.50) which is $J = \frac{\partial g}{\partial x_c^{\mathrm{a}}}$, is a matrix of size $(3n+3) \times (3n+3)$

with structure:

$$J = \left[\begin{array}{ccc:cc:c} 0 & \frac{\partial g_1}{\partial G^{\mathrm{a}}_{J,\mathrm{c}}} & \frac{\partial g_1}{\partial \eta^{\mathrm{a}}_{h,\mathrm{c}}} & \frac{\partial g_1}{\partial G^{\mathrm{a}}_{h,\mathrm{c}}} & \frac{\partial g_1}{\partial \gamma^{\mathrm{a}}_{\mathrm{c}}} & 0 \\ \frac{\partial g_2}{\partial \tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}} & \frac{\partial g_2}{\partial G^{\mathrm{a}}_{J,\mathrm{c}}} & 0 & 0 & 0 & 0 \\ \frac{\partial g_3}{\partial \tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}} & 0 & \frac{\partial g_3}{\partial \eta^{\mathrm{a}}_{h,\mathrm{c}}} & 0 & 0 & 0 \\ \hdashline 0 & 0 & 0 & \frac{\partial g_4}{\partial G^{\mathrm{a}}_{h,\mathrm{c}}} & 0 & 0 \\ 0 & 0 & 0 & \frac{\partial g_5}{\partial G^{\mathrm{a}}_{h,\mathrm{c}}} & \frac{\partial g_5}{\partial \gamma^{\mathrm{a}}_{\mathrm{c}}} & 0 \\ \hdashline \frac{\partial g_6}{\partial \tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}} & 0 & 0 & 0 & 0 & \frac{\partial g_6}{\partial \eta^{\mathrm{a}}_{J,\mathrm{c}}} \end{array}\right] \tag{2.51}$$

Evaluating at the equilibrium yields

$$J|_{x^{\mathrm{a}}_c=0} = \begin{bmatrix} J_{11}|_{x^{\mathrm{a}}_c=0} & J_{12}|_{x^{\mathrm{a}}_c=0} & 0 \\ 0 & J_{22}|_{x^{\mathrm{a}}_c=0} & 0 \\ J_{31}|_{x^{\mathrm{a}}_c=0} & 0 & J_{33}|_{x^{\mathrm{a}}_c=0} \end{bmatrix}. \tag{2.52}$$

The eigenvalues of this block triangular matrix are the eigenvalues of the diagonal blocks. Therefore we analyze J_{11} of size $(2n+1)\times(2n+1)$, J_{22} of size $(n+1)\times(n+1)$, and the scalar J_{33} separately. We can easily compute

$$J_{33}|_{x^{\mathrm{a}}_c=0} = \left.\frac{\partial g_6}{\partial \eta^{\mathrm{a}}_{J,\mathrm{c}}}\right|_{x^{\mathrm{a}}_c=0} = -\omega_{\mathrm{f}} \tag{2.53}$$

and the block J_{22} evaluated at the origin can be written as

$$J_{22}|_{x_c^{\mathrm{a}}=0} = \begin{bmatrix} -\omega_{\mathrm{f}} I & 0 \\ l & -\omega_{\mathrm{f}} \end{bmatrix} \tag{2.54}$$

where l need not be calculated. Because this matrix is itself block triangular, it's eigenvalues lie on the diagonal and are $\lambda_i = -\omega_{\mathrm{f}}$ for $i \in \{1, 2, ..., n+1\}$, therefore (2.54) is Hurwtiz. We have discovered $2n+1$ eigenvalues equal to $-\omega_{\mathrm{f}}$.

After careful algebra and differentiation we express

$$J_{11}|_{x_c^{\mathrm{a}}=0} = \begin{bmatrix} 0 & -M & -c\alpha||h_1||^{-2}h_1 \\ \omega_{\mathrm{f}} H & -\omega_{\mathrm{f}} I & 0 \\ \omega_{\mathrm{f}} h_1^T & 0 & -\omega_{\mathrm{f}} \end{bmatrix}, \tag{2.55}$$

with the symmetric, positive definite matrix $M \succ 0$ and the scalar $0 < \alpha < 1$ defined as

$$M = k\left(I - \alpha \frac{h_1 h_1^T}{||h_1||^2}\right), \tag{2.56}$$

$$\alpha = \frac{1}{2}\left(\frac{kc_1||h_1||^2 - c\eta_h^{\mathrm{a,e}}}{\sqrt{(kc_1||h_1||^2 - c\eta_h^{\mathrm{a,e}})^2 + \delta}} + 1\right), \tag{2.57}$$

$$c_1 = \frac{|h_0|}{2h_1^T H^{-1} h_1}\left(-\operatorname{sgn}(h_0) + \sqrt{1 + \delta \frac{d}{c^2 h_0^2}}\right). \tag{2.58}$$

We provide an extensive analysis of matrices in the form of (2.55) in Lemma 2.2 the appendix, but present the most essential result below.

Corollary 2.1. *Under Assumptions 2.1–2.3 and for $\alpha \in [0, 1]$, the matrix $J_{11}|_{x_c^{\mathrm{a}}=0}$ in* (2.55) *is always Hurwtiz, and therefore the Jacobian of the constrained average error system $J_{11}|_{x_c^{\mathrm{a}}=0}$ in* (2.52) *is always Hurwtiz.*

This result is slightly more general than what absolutely necessary for the convergence

analysis because it gives results for α potentially being equal to zero or one, even though the assumption of $\delta > 0$ precludes this possibility. This is still nonetheless useful knowledge because it tells us that as $\delta \to 0$, the eigenvalues of the linearization do not tend to undesirable values and provides hope for the future in developing a nonsmooth version of this algorithm with $\delta = 0$.

From Theorem 10.4 in [44], Corollary 2.1, and Proposition 2.5 given in the appendix, we conclude the following results.

Proposition 2.1 (Approximation by the Average System)**.** *Consider the exponentially stable equilibrium point* $x^{\mathrm{a,e}}$ (2.43)–(2.48) *of the average system in* (2.31)*. Under Assumptions 2.1–2.3, there exists positive constants* ρ, ω^* *such that if* $||\tilde{x}(0) - x^{\mathrm{a,e}}|| < \rho$ *then for all* $\omega \in [\omega^*, \infty)$*,*

$$||x^{\mathrm{a}}(t) - \tilde{x}(t)|| = O(1/\omega) \text{ for all } t \in [0, \infty), \tag{2.59}$$

$$||\theta(t) - (\tilde{\theta}^{\mathrm{a}}(t) + \theta^*)|| = O(1/\omega + a) \text{ for all } t \in [0, \infty). \tag{2.60}$$

The equilibrium of the average system (2.31), given in (2.43)–(2.48), resides in the interior of the safe set regardless of whether the unconstrained minimizer $\theta_{\min}$ of J lies in the safe set or not. One can show this by evaluating h in the original coorindates at the value $\tilde{\theta}^{\mathrm{a,e}}$. Additionally, as evident from (2.47), the parameter δ in the softened max operation, (2.10), creates a bias in the equilibrium that favors safety.

The proof of Theorem 2.1 utilizes the results of Section 2.4, and is shown in the appendix.

2.5 Safety

This section provides a description of safety, through the use of singular perturbation results. The following results also motivates the addition the $\max_\delta$ safety term which differentiates ASfES from standard extremum seeking. In the previous section we showed that through a linearization of the average system, the dynamics of ASfES are locally stable in n dimensions. We also showed that the average system and the original system are close. In the this section,

we begin by showing that the parameter $\tilde{\theta}^{\mathrm{a}}(t)$ of the average system is close to a reduced order model, denoted by its parameter $\theta_{\mathrm{r}}(t)$ - this reduced model is created by taking the the filter gains $\omega_{\mathrm{f}} \to \infty$. These results are used in the proof of Theorem 2.2 shown in the appendix. See Fig. 2.2 for a visual depiction of the high level steps in the analysis.

Consider the singular perturbation as ω_{f} approaches ∞. We can express (2.50) in the form of the standard singular perturbation model with $\epsilon = \frac{1}{\omega_{\mathrm{f}}}$ in the original time $t = \omega/\tau$ with the original ordering:

$$\dot{\tilde{\theta}}^{\mathrm{a}}_{\mathrm{c}} = w(z), \tag{2.61}$$

$$\epsilon \dot{z} = p(\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}, z), \tag{2.62}$$

where $z = [G^{\mathrm{a}}_{J,\mathrm{c}}; \eta^{\mathrm{a}}_{J,\mathrm{c}}; G^{\mathrm{a}}_{h,\mathrm{c}}; \eta^{\mathrm{a}}_{h,\mathrm{c}}; \gamma^{\mathrm{a}}_{\mathrm{c}}]$ and

$$w(z) = -k(G^{\mathrm{a}}_{J,\mathrm{c}} + G^{\mathrm{a,e}}_{J}) + (\gamma^{\mathrm{a}}_{\mathrm{c}} + \gamma^{\mathrm{a,e}})(G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h})\mathrm{max}_{\delta}\{\beta\}, \tag{2.63}$$

$$p(\tilde{\theta}^{\mathrm{a}}, z) = \begin{bmatrix} -\omega_{\mathrm{f}}(G^{\mathrm{a}}_{J,\mathrm{c}} + G^{\mathrm{a,e}}_{J}) + \omega_{\mathrm{f}} H(\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}} + \tilde{\theta}^{\mathrm{a,e}}) \\ -\omega_{\mathrm{f}}(\eta^{\mathrm{a}}_{J,\mathrm{c}} + \eta^{\mathrm{a,e}}_{J}) + \omega_{\mathrm{f}}(J(\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}} + \tilde{\theta}^{\mathrm{a,e}} + \theta^{*}) + \\ \qquad \frac{a^2}{4}\mathrm{Tr}(H)) \\ -\omega_{\mathrm{f}}(G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h}) + \omega_{\mathrm{f}} h_1 \\ -\omega_{\mathrm{f}}(\eta^{\mathrm{a}}_{h,\mathrm{c}} + \eta^{\mathrm{a,e}}_{h}) + \omega_{\mathrm{f}} h_0 + \omega_f h_1(\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}} + \tilde{\theta}^{\mathrm{a,e}} + \theta^{*}) \\ \omega_{\mathrm{f}}(\gamma^{\mathrm{a}}_{\mathrm{c}} + \gamma^{\mathrm{a,e}})(1 - (\gamma^{\mathrm{a}}_{\mathrm{c}} + \gamma^{\mathrm{a,e}})||G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h}||^2) \end{bmatrix}, \tag{2.64}$$

with the quantity

$$\beta = k(G^{\mathrm{a}}_{J,\mathrm{c}} + G^{\mathrm{a,e}}_{J})^{T}(G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h}) - c(\eta^{\mathrm{a}}_{h,\mathrm{c}} + \eta^{\mathrm{a,e}}_{h}) \tag{2.65}$$

In an effort to meet the conditions of Theorem 11.2 'Singular Perturbation on the Infinite Interval' in [44], we must show that the reduced model and the boundary layer model are exponentially

stable. Let us denote the root of $0 = p(\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}}, z)$ as the quasi-steady state

$$z = \begin{bmatrix} G_{J,\mathrm{c}}^{\mathrm{a}} \\ \eta_{J,\mathrm{c}}^{\mathrm{a}} \\ G_{h,\mathrm{c}}^{\mathrm{a}} \\ \eta_{h,\mathrm{c}}^{\mathrm{a}} \\ \gamma_{\mathrm{c}}^{\mathrm{a}} \end{bmatrix} = r(\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}}) = \begin{bmatrix} H\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}} \\ \frac{1}{2}\tilde{\theta}_{\mathrm{c}}^{\mathrm{a}T} H \tilde{\theta}_{\mathrm{c}}^{\mathrm{a}} \\ 0 \\ h_1^T \tilde{\theta}_{\mathrm{c}}^{\mathrm{a}} \\ 0 \end{bmatrix}. \tag{2.66}$$

The reduced model is n-dimensional differential equation

$$\dot{\theta}_{\mathrm{r}} = w(r(\theta_{\mathrm{r}})). \tag{2.67}$$

To show stability of the reduced order model we must determine if the Jacobian

$$J_r = \left. \frac{\partial w(r(\theta_{\mathrm{r}}))}{\partial \theta_{\mathrm{r}}} \right|_{\theta_{\mathrm{r}}=0} \tag{2.68}$$

is Hurwtiz. We are able to use the calculations in Section 2.4.3 to help us differentiate, using the terms

$$\left. \frac{\partial g_1(x_c^{\mathrm{a}})}{\partial G_{J,\mathrm{c}}^{\mathrm{a}}} \right|_{x_c^{\mathrm{a}}=0} = -M \tag{2.69}$$

$$\left. \frac{\partial g_1(x_c^{\mathrm{a}})}{\partial \eta_{h,\mathrm{c}}^{\mathrm{a}}} \right|_{x_c^{\mathrm{a}}=0} = -c\alpha ||h_1||^{-2} h_1 \tag{2.70}$$

from (2.55). Equation (2.69) is simply $\frac{\partial w(z)}{\partial G_{J,\mathrm{c}}^{\mathrm{a}}}$ and (2.70) is $\frac{\partial w(z)}{\partial \eta_{h,\mathrm{c}}^{\mathrm{a}}}$, both evaluated at $z = 0$. Now consider the following boundary layer variable as functions of θ_{r} as follows: $G_{J,\mathrm{c}}^{\mathrm{a}} = G_{J,\mathrm{c}}^{\mathrm{a}}(\theta_{\mathrm{r}}) = H\theta_{\mathrm{r}}$

and $\eta^{\mathrm{a}}_{h,\mathrm{c}} = \eta^{\mathrm{a}}_{h,\mathrm{c}}(\theta_{\mathrm{r}}) = h_1^T\theta_{\mathrm{r}}$. Therefore we can write the RHS of (2.67) as

$$w(r(\theta_{\mathrm{r}})) = -k(G^{\mathrm{a}}_{J,\mathrm{c}}(\theta_{\mathrm{r}}) + G^{\mathrm{a,e}}_{J}) + (\gamma^{\mathrm{a}}_{\mathrm{c}} + \gamma^{\mathrm{a,e}})(G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h})\max_{\delta}\{\beta\}, \tag{2.71}$$

$$\beta = k(G^{\mathrm{a}}_{J,\mathrm{c}}(\theta_{\mathrm{r}}) + G^{\mathrm{a,e}}_{J})^T(G^{\mathrm{a}}_{h,\mathrm{c}} + G^{\mathrm{a,e}}_{h}) - c(\eta^{\mathrm{a}}_{h,\mathrm{c}}(\theta_{\mathrm{r}}) + \eta^{\mathrm{a,e}}_{h}) \tag{2.72}$$

where we have yet to use the relations in (2.66). Computing the Jacobian, and linearizing (2.67), we can write (2.68) as

$$J_r = \left(\frac{\partial w(r(\theta_{\mathrm{r}}))}{\partial G^{\mathrm{a}}_{J,\mathrm{c}}}\frac{\partial G^{\mathrm{a}}_{J,\mathrm{c}}}{\partial \theta_{\mathrm{r}}} + \frac{\partial w(r(\theta_{\mathrm{r}}))}{\eta^{\mathrm{a}}_{h,\mathrm{c}}}\frac{\eta^{\mathrm{a}}_{h,\mathrm{c}}}{\partial \theta_{\mathrm{r}}}\right)\Bigg|_{\theta_{\mathrm{r}}=0} \tag{2.73}$$

$$= \frac{\partial w(r(\theta_{\mathrm{r}}))}{\partial G^{\mathrm{a}}_{J,\mathrm{c}}}\Bigg|_{\theta_{\mathrm{r}}=0}\frac{\partial G^{\mathrm{a}}_{J,\mathrm{c}}}{\partial \theta_{\mathrm{r}}} + \frac{\partial w(r(\theta_{\mathrm{r}}))}{\eta^{\mathrm{a}}_{h,\mathrm{c}}}\Bigg|_{\theta_{\mathrm{r}}=0}\frac{\eta^{\mathrm{a}}_{h,\mathrm{c}}}{\partial \theta_{\mathrm{r}}} \tag{2.74}$$

$$= -MH - c\alpha\frac{h_1 h_1^T}{||h_1||^2} \tag{2.75}$$

By Lemma 2.1 we conclude the eigenvalues of J_r are real and negative and the reduced model is exponentially stable at the origin.

The boundary layer model is

$$\dot{z}_{\mathrm{b}} = p(\theta, z_{\mathrm{b}} + r(\theta)) \tag{2.76}$$

for a fixed θ. The boundary layer model can be shown to be

$$\dot{z}_{\mathrm{b}} = \begin{bmatrix} -z_{\mathrm{b},1} \\ -z_{\mathrm{b},2} \\ -z_{\mathrm{b},3} \\ -z_{\mathrm{b},4} \\ (z_{\mathrm{b},5} + \frac{1}{||h_1||^2})(1 - (z_{\mathrm{b},5} + \frac{1}{||h_1||^2})||z_{\mathrm{b},3} + h_1||^2) \end{bmatrix}. \tag{2.77}$$

with a slight notation abuse denoting $z_{\mathrm{b,i}}$ as either a vector or scalar variable corresponding with

the vectors and scalars $G^{\mathrm{a}}_{J,\mathrm{c}}, \eta^{\mathrm{a}}_{J,\mathrm{c}}, \ldots, \gamma^{\mathrm{a}}_{\mathrm{c}}$. It is evident by inspection that the linearization of (2.77) at the origin has $2n+3$ eigenvalues at -1. (Note that the linearization of the last row will yield a cross term due to $z_{\mathrm{b},3}$, but since the linearization is triangular we need only consider the diagonal entrees.) Hence, the boundary layer model is exponentially stable.

The reduced order model in (2.78)

$$\dot{\theta}_{\mathrm{r}} = -k(H\theta_{\mathrm{r}} + G^{\mathrm{a,e}}_J) + \frac{1}{||h_1||^2}\max_\delta\{k(H\theta_{\mathrm{r}} + G^{\mathrm{a,e}}_J)^T h_1 - c(h_1^T\theta_{\mathrm{r}} + \eta^{\mathrm{a,e}}_h)\}, \tag{2.78}$$

can be expressed in the coordinates $\theta_{\mathrm{r}} = \tilde{\theta}_{\mathrm{r}} - \tilde{\theta}^{\mathrm{a,e}}$, as

$$\dot{\tilde{\theta}}_{\mathrm{r}} = -kH\tilde{\theta}_{\mathrm{r}} + \frac{h_1}{||h_1||^2}\max_\delta\{k\tilde{\theta}_{\mathrm{r}}^T H h_1 - c(h_0 + h_1^T\tilde{\theta}_{\mathrm{r}})\}, \tag{2.79}$$

using the relations $\eta^{\mathrm{a,e}}_h = h_0 + h_1^T H^{-1} G^{\mathrm{a,e}T}_J$ and $\tilde{\theta}^{\mathrm{a,e}} = H^{-1}G^{\mathrm{a,e}}_J$. We can now state the following result having satisfied the conditions in [44].

Proposition 2.2 (Singular Perturbation of Average System)**.** *Let the solution of* (2.31) *be given by* $x^{\mathrm{a}}(t)$*, with its first component* $\tilde{\theta}^{\mathrm{a}}(t)$ *in the time scale* t*, the solution of reduced model* (2.79) *be given by* $\tilde{\theta}_{\mathrm{r}}(t)$*, and suppose Assumptions 2.1–2.3 hold. Then there exist positive constants* μ_1, μ_2 *and* ω^*_{f} *such that for all*

$$||\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}(0)|| < \mu_1, \;\; ||z(0) - r(\tilde{\theta}^{\mathrm{a}}_{\mathrm{c}}(0))|| < \mu_2, \;\; \omega_{\mathrm{f}} > \omega^*_{\mathrm{f}} \tag{2.80}$$

the singular perturbation problem (2.61)–(2.62) *has a unique solution for all* $t > 0$ *and*

$$||\tilde{\theta}^{\mathrm{a}}(t) - \tilde{\theta}_{\mathrm{r}}(t)|| = O(1/\omega_{\mathrm{f}}) \quad \textit{for all } t \in [0, \infty). \tag{2.81}$$

We omit other bounds given by [44] (those based on the quasi-steady state (2.66)) as they do not relate to safety. Proposition 2.1 describes closeness of the original system with that of the average system. Proposition 2.2 describes a closeness in the average system and the reduced

model *of the average system* (2.78) (with ω_f large). Now we show the reduced order model (2.79) is indeed safe.

Proposition 2.3 (Safety of the Reduced System)**.** *Under the dynamics of the system* (2.79)*, the set* $\{\tilde{\theta}_r \in \mathbb{R} : h(\tilde{\theta}_r + \theta^*) \geq 0\}$ *is forward invariant and* $h(\tilde{\theta}_r(t) + \theta^*) \geq h(\tilde{\theta}_r(0) + \theta^*)e^{-ct}$ *for all* $t \geq 0$.

Proof. We can verify that (2.79) renders the set forward invariant by showing $\dot{h}(\tilde{\theta}_r + \theta^*) + ch(\tilde{\theta}_r + \theta^*) \geq 0$ [10]:

$$\dot{h} + ch = -kh_1^T H\tilde{\theta}_r + ch(\tilde{\theta}_r + \theta^*) + \max_\delta\{kh_1^T H\tilde{\theta}_r - ch(\tilde{\theta}_r + \theta^*)\} > 0 \tag{2.82}$$

where we use the relation $h_0 + h_1^T\tilde{\theta}_r = h(\tilde{\theta}_r + \theta^*)$. The inequality holds because for any x, $-x + \max_\delta\{x\} > 0$. □

The safety of the reduced system, and its closeness to that of the average system provide the key intermediate results in achieving Theorem 2.2 — the proof is shown in the appendix.

2.6 Newton-Based Assignably Safe Extremum Seeking

One might imagine that if we can assign a bound on the rate of h, while conducting gradient-based optimization, then we should be assign safety when performing Newton-based optimization [30], which has an assignable rate of convergence of the parameter itself. A Newton-Based Assignably Safe ES (NB-ASfES) scheme would conceivably achieve 1) assignable safety and 2) assignable convergence of the parameter. In NB-ASfES we hope for a convergence rate of the parameter to be assigned a rate k when the assigned safety condition (with rate c) is not violated – and if the safety condition is violated, the convergence of the parameter is such that the safety rate assignment is maintained. We show in this section that this is *not possible* with the NB-ASfES approach in multiple dimensions, because the NB-ASfES scheme will not solve the optimization problem in (2.1) in general. Furthermore, we also show that a NB-ASfES scheme is only useful for $n = 1$.

To understand the basic dynamics of the ASfES scheme, we can think of the algorithm approximating the following dynamics:

$$\dot{\theta} = u_0 + \frac{\nabla h(\theta)}{||\nabla h(\theta)||^2} \max\{-u_0^T \nabla h(\theta) - ch(\theta), 0\}, \tag{2.83}$$

with $u_0 = -k\nabla J(\theta)$. The form of (2.83) guarantees safety and can be derived for a nominal control law u_0 using the standard QP formulation given in [10]. Now consider a NB-ASfES controller, with the nominal control $u_0 = -kH^{-1}\nabla J(\theta) = -k(\theta - \theta^*)$ where $k > 0$ is the assigned rate of convergence. Note that in the extremum seeking form of (2.83), the NB-ASfES algorithm contains a state $\Gamma(t)$ which provides the estimate of H^{-1} [30].

Consider the case of $h_0 < 0$ where the minimizer of J on $\{h(\theta) \geq 0\}$ is unsafe and given by

$$\theta_{\text{smin}} = \frac{|h_0| H^{-1} h_1}{h_1^T H^{-1} h_1} + \theta^*. \tag{2.84}$$

See the proof of Proposition 2.5 in the Appendix for this fact. Now let us ask the question: is the equilibrium of (2.83) equal to θ_{smin} when $u_0 = -k(\theta - \theta^*)$ and $h_0 < 0$ under Assumptions 2.1 - 2.3? Solving for

$$-k(\theta_{\text{smin}} - \theta^*) + \frac{\nabla h(\theta_{\text{smin}})}{||\nabla h(\theta_{\text{smin}}||^2} \max\{k(\theta_{\text{smin}} - \theta^*)^T \nabla h(\theta_{\text{smin}}) - ch(\theta_{\text{smin}}), 0\} = 0 \tag{2.85}$$

yields the condition

$$H^{-1} h_1 = \frac{h_1 H^{-1} h_1}{||h_1||^2} h_1. \tag{2.86}$$

This is nothing more than an eigenvector equation stating that h_1 must be an eigenvector of H^{-1}. Because H and H^{-1} share eigenvectors, h_1 must also be an eigenvector of H. This is a restrictive condition, stating that the gradient of h, must point along one of the principle axes of the ellipsoids formed by the levels of J.

One can also come to the conclusion that NB-ASfES will not solve the constrained

optimization problem with some geometric intuition for the case $n = 2$. Intuition says that ASfES, with the nominal control law of $u_0 = -k\nabla J$, causes trajectories to descend down the level curves of J, finding the smallest level curve of J which intersects the safe set at some point $p = \theta_{\text{smin}}$. The level curve $J(\theta) = J(\theta_{\text{smin}})$ will in general be an ellipse in 2 dimensions. If the nominal control law $u_0 = -k(\theta - \theta^*)$ is used instead, trajectories will now travel down level curves of some other function $\tilde{J}(\theta) = \frac{1}{2}(\theta - \theta^*)^T(\theta - \theta^*)$ because $(\theta - \theta^*)$ is the gradient of $\tilde{J}$. The smallest level set of $\tilde{J}$ touching the safe set is, in general, a circle and will intersect at some other point $\tilde{p}$, different from the point p given by the ellipse generated from J.

We cannot assume that the restrictive condition (2.86) holds when h_1 and H are unknown to the user, and therefore we do not present a NB-ASfES scheme for the generic n dimensional case. The only situation where (2.86) holds is in the 1 dimensional case.

Based on this discussion, we propose a NB-ASfES scheme only for the case of $n = 1$, as

$$\dot{\hat{\theta}} = -k\Gamma G_J + \gamma \max_\delta\{k\Gamma G_J G_h - c\eta_h\}G_h \tag{2.87}$$

$$\dot{G}_J = -\omega_{\text{f}} G_J + \omega_{\text{f}}(J(\hat{\theta}(t) + S(t)) - \eta_J)M(t) \tag{2.88}$$

$$\dot{\eta}_J = -\omega_{\text{f}}\eta_J + \omega_{\text{f}} J(\hat{\theta}(t) + S(t)) \tag{2.89}$$

$$\dot{G}_h = -\omega_{\text{f}} G_h + \omega_{\text{f}}(h(\hat{\theta}(t) + S(t)) - \eta_h)M(t) \tag{2.90}$$

$$\dot{\eta}_h = -\omega_{\text{f}}\eta_h + \omega_{\text{f}} h(\hat{\theta}(t) + S(t)) \tag{2.91}$$

$$\dot{\gamma} = \omega_{\text{f}}\gamma(1 - \gamma|G_h|^2) \tag{2.92}$$

$$\dot{\Gamma} = \omega_{\text{f}}\Gamma(1 - \Gamma J(\hat{\theta}(t) + S(t))N(t)) \tag{2.93}$$

where $N(t) = \frac{16}{a^2}(\sin^2(\omega t) - \frac{1}{2})$.

The resulting theory is identical to that of Theorem 2.1 and 2.2 for the scalar algorithm given in (2.87) - (2.93), and follow similar analysis steps (but for $n = 1$) shown in Section 2.4 and Section 2.5. The scalar NB-ASfES algorithm shares the assignable safety property with that of the gradient-based algorithm ASfES, with the additional feature that the nominal convergence

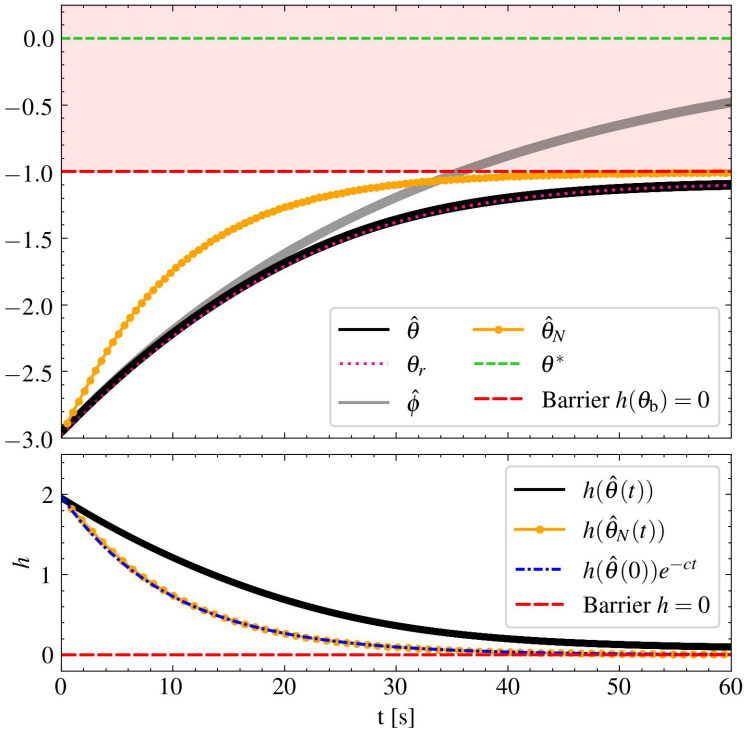

Figure 2.3. A demonstration of the algorithm in 1D with the optimizer lying in the unsafe region — trajectories of the original and reduced system of ASfES are given by $\hat{\theta}$ and θ_r and the NB-ASfES trajectory is given by $\hat{\theta}_N$. The classical ES solution is given by $\hat{\phi}$, and the red shaded region marks the unsafe set where $h < 0$.

rate of the parameter can also be assigned.

In Section 2.7.1 we illustrate the behavior of the NB-ASfES algorithm.

2.7 Simulations

2.7.1 Scalar System

To demonstrate the closeness of trajectories provided by our analysis, we show the solutions of the original ASfES algorithm (2.3)–(2.8) and reduced system (2.78), under Assumptions 2.1–2.3 in Figure 2.3. The unknown function parameters are $\theta^* = 0$ (so $\theta = \tilde{\theta}$), $J^* = 0$, $H = 0.1$, $h_0 = -1$ and $h_1 = -1$. The ASfES and NB-ASfES controller parameters are $a = 0.25$, $k = 0.3$, $c = 0.1$, $\delta = 10^{-3}$, $\omega = 200$, and $\omega_f = 3$. The initial conditions are $\hat{\theta}(0) = \hat{\theta}_N(0) = \theta_r(0) = -3$, while all other estimator states are initialized to their exact quantities. The trajectories of the

ASfES and NB-ASfES algorithm parameters are denoted by $\hat{\theta}(t)$ and $\hat{\theta}_N(t)$ respectively. Note that although we do not know the gradients of J and h to initialize G_J and G_h to exactly $\nabla J(\theta(0))$ and $\nabla h(\theta(0))$, we can generate an accurate estimate by "warming up" the algorithm in (2.3)–(2.8) by setting $\dot{\hat{\theta}} = 0$, and integrating over time until the states $G_J, \eta_J, G_h, \eta_h \gamma$ (and Γ for NB-ASfES) converge adequately.

We compare the ASfES and NB-ASfES algorithms with the classical ES algorithm denoted by solution $\phi(t)$ which obey the dynamics

$$\begin{aligned}
\dot{\hat{\phi}} &= -kG_J \\
\dot{G}_J &= -\omega_{\mathrm{f}} G_J + \omega_{\mathrm{f}}(J(\hat{\phi}(t) + S(t)) - \eta_J) M(t) \\
\dot{\eta}_J &= -\omega_{\mathrm{f}} \eta_J + \omega_{\mathrm{f}} J(\hat{\phi}(t) + S(t))
\end{aligned}$$

with $\phi = \hat{\phi} + S(t)$. The initial condition is $\phi(0) = -3$ and other filtered estimates are initialized to their exact quantities. The gains $a, k, \omega_{\mathrm{f}}, \omega$ are set the same as in the ASfES and NB-ASfES schemes.

The reduced system in the upper part of Fig. 2.3 maintains safety while the ASfES and NB-ASfES schemes maintain practical safety. In the lower part of Fig. 2.3 we observe the behavior of $h(\hat{\theta}(t))$ and $h(\hat{\theta}_N(t))$ have the assigned time constant c bound associated with its approach to the barrier. The NB-ASfES exhibits faster convergence due to the cancellation of the small Hessian, $H = 0.1$, yet maintains the assigned practical safety bound, which can be observed by the closeness between the plots of $h(\hat{\theta}_N(t))$ and $h(\hat{\theta}(0))e^{-ct}$.

2.7.2 Quadratic Objective and Linear CBF in 2 Dimensions

In this final example we illustrate the effect of c on ASfES in 2 dimensions. The unknown functions are $h = \theta_1 + \theta_2 - 1$ and $J = \theta_1^2 + \theta_2^2$. The controller parameters are $a = 0.25$, $k = 0.1$, $\delta = 10^{-3}$, $\omega_1 = 75$, $\omega_2 = 100$, and $\omega_{\mathrm{f}} = 3$. Trajectories are plotted in Fig. 2.4 for various initial conditions for $c = 1$ and $c = 0.1$ with all other controller states are initialized to their exact

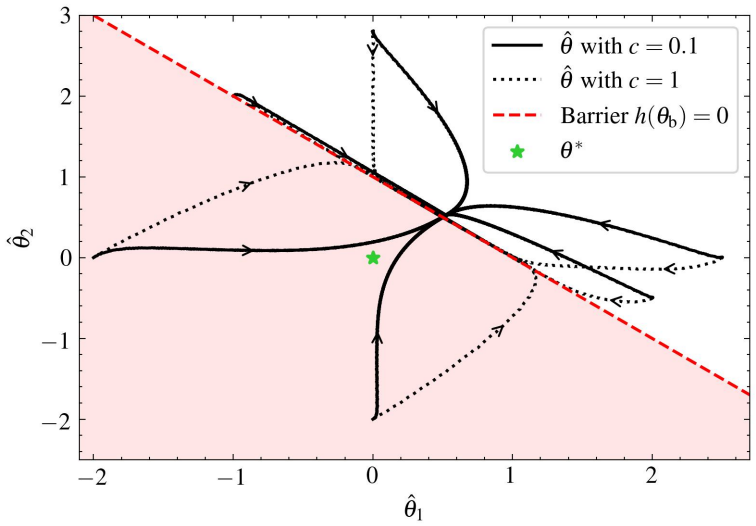

Figure 2.4. Trajectories of $\hat{\theta}$ in 2D with a linear h and quadratic J. Six initial conditions, starting both inside and outside the safe set, and their trajectories are plotted for both $c = 1$ and $c = 0.1$. The global optimizer θ^* is outside of the safe set.

quantities.

Notice that a larger value of c allows a direct approach (and a truly faster approach in time, see Theorem 2.2) to the unsafe region when starting from the safe region but a more direct approach to the safe region when starting from the unsafe region.

2.7.3 Barrier Islands and the Effect of the Attractivity Rate

In this setting we utilize the 2D algorithm given in (2.3)–(2.8). The unknown functions are $h = \cos(0.2\pi\theta_1) + \theta_1 \sin(0.3\pi\theta_2)$ and $J = (\theta_1 - 4)^2 + (\theta_2 - 4)^2$. The controller parameters are $a = 0.25$, $k = 0.1$, $c = \{1.00, 0.10, 0.01\}$, $\delta = 10^{-3}$, $\omega_1 = 75$, $\omega_2 = 100$, and $\omega_f = 3$. The initial conditions are $\theta(0) = [0, -4]^T$, while all other controller states are initialized to their exact quantities.

The behavior demonstrated in Fig. 2.5 shows behavior that the classical safety filter design often yields, with the trajectory 'smoothly sliding' around the unsafe barrier island. In this case the trajectory manages to stabilize near the optimizer with a small conservative bias due to δ. This simulation also illustrates, qualitatively, that ASfES may have stronger convergence

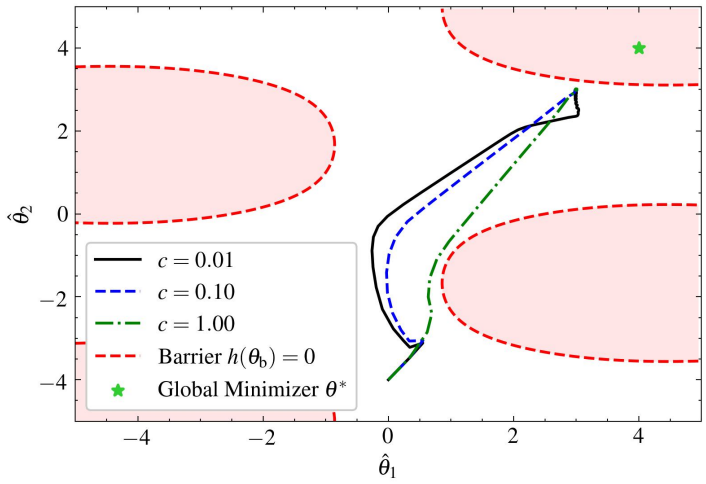

Figure 2.5. Trajectory of $\hat{\theta}_1, \hat{\theta}_2$ in 2D with island-shaped barriers and the optimizer lying in an unsafe region. The global optimizer θ^* is selected outside of the safe set.

properties than the local properties proved here, and may be used for the constrained optimization of J from a larger set of initial conditions and not merely within a ball of initial conditions near the constrained optimum - this is an open problem. Additionally we see the effect that c has qualitatively. Namely, that a lower c value takes a more conservative route around the parameter space, being more cautious to avoid unsafe regions.

2.8 Conclusion

We present ASfES and NB-ASfES which achieve an assignable safety attractivity rate, specifying a bound on the rate at which h, the measurement of safety, is allowed to decay — with trajectories moving towards the safe set from the unsafe region, or moving towards the unsafe set from a safe region. We demonstrate that the NB-ASfES scheme is only useful in the case of a scalar parameter, and may lead to non-optimal convergence if used in multiple dimensions, yet features an assignable rate of parameter convergence which is favored only if it agrees with the safety attractivity rate.

The algorithms presented here are local in nature, but provide benefits over their semiglobal versions (studied in Chapter 3), with assignable safety attractivity and assignable (nominal)

parameter convergence.

2.9 Chapter Appendix

This section includes the analysis of the matrix (2.55) and the proof of Lemma 2.1. It also includes additional details on the spectrum of matrices of this form. The first result describes a general algebraic relationship between the spectra of two matrices, X and Z. The second result show the eigenvalues of a particular Z are real and positive. These two results and Proposition 2.5 are used in Theorem 2.1. We then present the proof of Theorem 2.2. We use the notation of a unit vector $\hat{h}_1 = h_1||h_1||^{-1}$.

2.9.1 Eigenvalues of X and Z

Proposition 2.4. *Consider the* $\mathbb{R}^{2n+1\times 2n+1}$ *matrix*

$$X = \begin{bmatrix} 0 & -M & -\tilde{c}h_1 \\ \omega_{\mathrm{f}}H & -\omega_{\mathrm{f}}I & 0 \\ \omega_{\mathrm{f}}h_1^T & 0 & -\omega_{\mathrm{f}} \end{bmatrix}$$

with $\omega_{\mathrm{f}}, \tilde{c} > 0$, $h_1 \in \mathbb{R}^n$ *and matrices* $M, H \in \mathbb{R}^{n\times n}$.

Then, for all $\lambda \in \sigma(X)$ *such that* $\lambda \neq -\omega_{\mathrm{f}}$, *there exists a* $\bar{\lambda} \in \sigma(\omega_{\mathrm{f}}MH + \omega_{\mathrm{f}}\tilde{c}h_1h_1^T)$ *satisfying* $\lambda^2 + \omega_{\mathrm{f}}\lambda + \bar{\lambda} = 0$.

Proof. Any eigenvalue $\lambda \neq -\omega_{\mathrm{f}}$, of the matrix X, must satisfy the eigenvalue expression for a nonzero eigenvector,

$$\begin{bmatrix} 0 & -M & -\tilde{c}h_1 \\ \omega_{\mathrm{f}}H & -\omega_{\mathrm{f}}I & 0 \\ \omega_{\mathrm{f}}h_1^T & 0 & -\omega_{\mathrm{f}} \end{bmatrix} \begin{pmatrix} u \\ v \\ w \end{pmatrix} = \lambda \begin{pmatrix} u \\ v \\ w \end{pmatrix} \tag{2.94}$$

with $u, v \in \mathbb{C}^n$, $w \in \mathbb{C}$. From the second and third rows of (2.94) we have

$$\omega_{\mathrm{f}} H u = (\lambda + \omega_{\mathrm{f}}) v \tag{2.95}$$

$$\omega_{\mathrm{f}} h_1^T u = (\lambda + \omega_{\mathrm{f}}) w \tag{2.96}$$

and from the first row of (2.94) we have

$$-(\lambda + \omega_{\mathrm{f}}) M v - (\lambda + \omega_{\mathrm{f}}) \tilde{c} h_1 w = (\lambda + \omega_{\mathrm{f}}) \lambda u, \tag{2.97}$$

where we have multiplied on both sides $\lambda + \omega_{\mathrm{f}} \neq 0$. Now we substitute expressions in (2.97) for $(\lambda + \omega_{\mathrm{f}}) w$ and $(\lambda + \omega_{\mathrm{f}}) v$,

$$(\omega_{\mathrm{f}} M H + \omega_{\mathrm{f}} \tilde{c} h_1 h_1^T) u = -(\lambda^2 + \omega_{\mathrm{f}} \lambda) u. \tag{2.98}$$

What we arrive at, is another eigenvalue equation describing the scaling of the vector u by the value $-(\lambda^2 + \omega_{\mathrm{f}} \lambda)$.

It is also guaranteed that $u \neq 0$ for the eigenvector equation (2.94) under the following argument: suppose not and $u = 0$, which implies in that $v = 0$ and $w = 0$ from (2.95) and (2.96) because $\lambda + \omega_{\mathrm{f}} \neq 0$. Because the equation (2.94) must have nonzero eigenvector then we have a contraction which implies $u \neq 0$.

Therefore we can say that $\bar{\lambda}$ is an eigenvalue of Z defined as

$$Z := \omega_{\mathrm{f}} M H + \omega_{\mathrm{f}} \tilde{c} h_1 h_1^T, \tag{2.99}$$

where

$$\bar{\lambda} = -(\lambda^2 + \omega_{\mathrm{f}} \lambda). \tag{2.100}$$

□

2.9.2 Eigenvalues of Z are real and positive

Lemma 2.1. *Consider the matrix* $Z := \omega_{\mathrm{f}} MH + \omega_{\mathrm{f}} \tilde{c} h_1 h_1^T$ *with the following parameters:* $H \succ 0$, $M = k\left(I - \alpha \frac{h_1 h_1^T}{||h_1||^2}\right)$, $k, \omega_{\mathrm{f}} > 0$, $\tilde{c} \geq 0$, $0 \leq \alpha \leq 1$, *and* $h_1 \in \mathbb{R}^n_{\neq 0}$ *such that if the scalar* $\alpha = 1$ *then* $\tilde{c} > 0$ *strictly. Then, each* $\bar{\lambda} \in \sigma(Z)$ *is real and strictly positive.*

Proof. We consider three cases of α in this proof. Case 1 is when $\alpha = 0$, Case 2 is $\alpha \in (0, 1)$, and Case 3 is when $\alpha = 1$.

Case 1: If $\alpha = 0$, the matrix Z can readily be shown to have real and strictly positive eigenvalues because it is the sum of a positive definite matrix and a positive semi definite matrix.

Case 2: Consider $\alpha \in (0, 1)$. The positive definite matrix $M \succ 0$ is recalled below and its' inverse is also shown:

$$M = k(I - \alpha \hat{h}_1 \hat{h}_1^T) \succ 0, \tag{2.101}$$

$$M^{-1} = k^{-1}\left(I + \frac{\alpha}{1-\alpha} \hat{h}_1 \hat{h}_1^T\right) \succ 0, \tag{2.102}$$

where $\hat{h}_1 = h_1 ||h_1||^{-1}$. The reader can check by hand that $MM^{-1} = M^{-1}M = I$ holds, see [60] for background. We use (2.102) to derive the expression

$$\hat{h}_1 \hat{h}_1^T = \frac{(1-\alpha)}{\alpha} k M^{-1} - \frac{(1-\alpha)}{\alpha} I \quad \text{for } \alpha \neq 0, \tag{2.103}$$

and rewrite Z as

$$Z = \omega_{\mathrm{f}} MH + \omega_{\mathrm{f}} \tilde{c} h_1 h_1^T, \tag{2.104}$$

$$Z = \omega_{\mathrm{f}} MH + \beta\left((1-\alpha) k M^{-1} - (1-\alpha) I\right). \tag{2.105}$$

for some positive scalar $\beta = \omega_{\mathrm{f}} \tilde{c} / \alpha \geq 0$. Now we perform a similarity transformation $Z' = T^{-1} Z T$.

Let us take $T = M^{1/2} > 0$ and compute

$$Z' = M^{-1/2} Z M^{1/2} = \omega_f M^{1/2} H M^{1/2} + \beta\alpha \hat{h}_1 \hat{h}_1^T. \tag{2.106}$$

The matrix Z' is the sum of a positive definite matrix $\omega_f M^{1/2} H M^{1/2}$ (which can be shown by definition) and a positive semi definite matrix $\beta\alpha \hat{h}_1 \hat{h}_1^T$. Because two similar matrices have the same eigenvalues, it follows that Z shares its' eigenvalues with that of a symmetric, positive definite matrix Z' which is known to have real and positive eigenvalues.

Case 3: Let us consider the case of $\alpha = 1$, $\tilde{c} > 0$ and rewrite Z as

$$Z = \omega_f k \left(I - \frac{h_1 h_1^T}{||h_1||^2} \right) H + \omega_f \tilde{c} h_1 h_1^T \tag{2.107}$$

$$= \omega_f k \left(I - \hat{h}_1 \hat{h}_1 \right) H + \mu \hat{h}_1 \hat{h}_1^T \tag{2.108}$$

with $\mu = \omega_f \tilde{c} ||h_1||^2 > 0$ .The symmetric, rank 1 matrix $\hat{h}_1 \hat{h}_1$, with spectrum $\sigma(\hat{h}_1 \hat{h}_1) = \{1, 0, 0, ..., 0\}$, is diagonalizable by an orthogonal matrix U such that $U^T = U^{-1}$. We can write

$$\hat{h}_1 \hat{h}_1^T = U D U^T = U \begin{bmatrix} 1 & 0_{1\times n-1} \\ 0_{n-1\times 1} & 0_{n-1\times n-1} \end{bmatrix} U^T \tag{2.109}$$

with D having a single nonzero element $d_{11} = 1$. Then it follows that

$$\begin{aligned} Z &= \omega_f k (I - U D U^T) H + \mu U D U^T \\ &= U \left(\omega_f k (I - D) U^T H U + \mu D \right) U^T \\ &= U \left(\omega_f k (I - D) \tilde{H} + \mu D \right) U^T \end{aligned} \tag{2.110}$$

where $\tilde{H} \equiv U^T H U > 0$ is a matrix similar to H. Therefore Z is similar to $Z' \equiv \omega_f k (I - D) \tilde{H} + \mu D$

which can be written as block triangular:

$$\begin{aligned} Z' &= \omega_{\mathrm{f}} k (I - D)\tilde{H} + \mu D \\ &= \omega_{\mathrm{f}} k \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \tilde{h}_{11} & \tilde{H}_{12}^T \\ \tilde{H}_{12} & \tilde{H}_{22} \end{bmatrix} + \mu \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} \mu & 0 \\ \omega_{\mathrm{f}} k \tilde{H}_{12} & \omega_{\mathrm{f}} k \tilde{H}_{22} \end{bmatrix}. \end{aligned} \tag{2.111}$$

The matrix $\tilde{H}_{22}$ is a diagonal block of a positive definite matrix, and is positive definite itself [36]. Because Z' is block triangular with positive definite diagonal blocks, Z has real and positive eigenvalues. □

2.9.3 Properties of X

In the next lemma, we extend the conclusions from the prior two results in this section and make more specific claims on the eigenvalues of the linearizing matrix. Statement 2 assigns each eigenvalue of Z to a pair of eigenvalues of X. This is an extension of the results in Proposition 2.4, which only concerns the existence of an eigenvalue of Z.

Lemma 2.2. *Under Assumptions 2.1 - 2.3, consider*

$$X = \begin{bmatrix} 0 & -M & -\tilde{c}h_1 \\ \omega_{\mathrm{f}} H & -\omega_{\mathrm{f}} I & 0 \\ \omega_{\mathrm{f}} h_1^T & 0 & -\omega_{\mathrm{f}} \end{bmatrix}, \tag{2.112}$$

with $0 \leq \alpha \leq 1$, $M = k\left(I - \alpha \hat{h}_1 \hat{h}_1^T\right)$, *and* $\tilde{c} = c\alpha ||h_1||^{-2}$. *Then*

1. *X is Hurwtiz.*

2. *The spectrum* $\sigma(X) = \{-\omega_{\mathrm{f}}, \lambda_1^+, \lambda_1^-, ..., \lambda_n^+, \lambda_n^-\}$ *where* $\lambda_i^+ / \lambda_i^- \neq -\omega_{\mathrm{f}}$ *are the positive and*

negative solutions to $\lambda_i^2 + \omega_f \lambda_i + \bar{\lambda}_i = 0$ *for each* $\bar{\lambda}_i \in \sigma(Z)$ *and for all* $\alpha \in [0,1]$*, where* $Z := \omega_f MH + \omega_f \tilde{c} h_1 h_1^T$.

(a) If $\alpha = 0$ *then* $\sigma(Z) = \sigma(\omega_f k H)$.

(b) If $\alpha = 1$ *then* $\sigma(Z) = \{\omega_f c\} \cup \sigma(\omega_f k \tilde{H}_{22})$ *where* $\tilde{H}_{22} > 0$ *is the lower right* $n-1 \times n-1$ *block of* $\tilde{H} = U^T H U$ *for an orthogonal matrix* U *which diagonalizes* $\hat{h}_1 \hat{h}_1^T$.

Proof. Statement 1 follows directly from Proposition 2.4 and Lemma 2.1. From Proposition 2.4: each eigenvalue $\lambda \in \sigma(X)$, for which $\lambda \neq -\omega_f$, there exists a $\bar{\lambda} \in \sigma(\omega_f MH + \omega_f \tilde{c} h_1 h_1^T)$ satisfying $\lambda^2 + \omega_f \lambda + \bar{\lambda} = 0$. Because $\bar{\lambda} > 0$ (Lemma 2.1), by the Routh-Hurwtiz criterion X is Hurwtiz.

Statement 2 can be shown by taking an eigenvalue of Z, denoted as the positive scalar $\bar{\lambda} > 0$ (positivity is shown in Lemma 2.1) and showing that it generates 2 eigenvalues of X. Take the logic and algebra of Proposition 2.4 in reverse. Express $\bar{\lambda} = -(\lambda_i^2 + \omega_f \lambda_i)$ for either $\lambda_i = \lambda_i^+$ or $\lambda_i = \lambda_i^-$ satisfying $\lambda_i^2 + \omega_f \lambda_i + \bar{\lambda} = 0$. The eigenvalue equation is

$$(\omega_f MH + \omega_f \tilde{c} h_1 h_1^T) u = -(\lambda_i^2 + \omega_f \lambda_i) u. \tag{2.113}$$

for some distinct eigenvector $u \in \mathbb{C}^n_{\neq 0}$. We introduce $v \in \mathbb{C}^n, w \in \mathbb{C}$ satisfying

$$\omega_f H u = (\lambda_i + \omega_f) v, \tag{2.114}$$

$$\omega_f h_1^T u = (\lambda_i + \omega_f) w, \tag{2.115}$$

and make substitutions to get

$$-(\lambda_i + \omega_f) M v - (\lambda_i + \omega_f) \tilde{c} h_1 w = (\lambda_i + \omega_f) \lambda_i u. \tag{2.116}$$

where $\lambda_i \neq -\omega_f$ because $\lambda_i^2 + \omega_f \lambda_i + \bar{\lambda} = 0$ and $\bar{\lambda} > 0$. Then we recover the eigenvector equation in (2.94) with the eigenvalue λ_i satisfying $X\nu = \lambda_i \nu$ with $\nu = [u^T, v^T, w]^T$. Because this holds for $2n$ eigenvalues ($\lambda_i = \lambda_i^+$ and $\lambda_i = \lambda_i^-$ for each $\bar{\lambda} \neq -\omega_f$), we have accounted for all but one of the

$2n+1$ eigenvalues in $\sigma(X)$.

Finally, the spectrum of X always contains an eigenvalue $-\omega_{\mathrm{f}}$, the last unaccounted eigenvalue. When $\alpha = 0$, the matrix Y has an eigenvalue of $-\omega_{\mathrm{f}}$, due to the triangular structure. For $\alpha \in (0,1]$, the determinant $\det(X+\omega_{\mathrm{f}} I)$ is

$$\begin{aligned}
|X+\omega_{\mathrm{f}} I| &= \begin{vmatrix} I\omega_{\mathrm{f}} & -M & -\tilde{c}h_1 \\ \omega_{\mathrm{f}} H & 0 & 0 \\ \omega_{\mathrm{f}} h_1^T & 0 & 0 \end{vmatrix} = \omega_{\mathrm{f}}^n \begin{vmatrix} HM & \tilde{c}Hh_1 \\ h_1^T M & \tilde{c}||h_1||^2 \end{vmatrix} \\
&= \tilde{c}||h_1||^2 \omega_{\mathrm{f}}^n \det(HM - \tilde{c}Hh_1(\tilde{c}||h_1||^2)^{-1} h_1^T M) \\
&= \tilde{c}||h_1||^2 \omega_{\mathrm{f}}^n \det(H(I - \hat{h}_1 \hat{h}_1^T)M) = 0
\end{aligned} \tag{2.117}$$

using rules for determinants of block matrices. Because $(I - \hat{h_1}\hat{h_1}^T)$ is singular and the determinant of a product is the product of the determinants, then $\det(X+\omega_{\mathrm{f}} I) = 0$ and $-\omega_{\mathrm{f}}$ is an eigenvalue.

Statement 2a can be shown by noting that Z for $\alpha = 0$, simply reduces to $k\omega_{\mathrm{f}} H$.

Statement 2b can be shown by looking at the analysis of this case in the proof in Lemma 2.1, Case 3. The matrix Z is similar to the block diagonal matrix Z' which has spectra of its diagonal blocks: $\sigma(\mu) \cup \sigma(\omega_{\mathrm{f}} k \tilde{H}_{22})$. The matrix $\tilde{H}_{22} \succ 0$ is the lower right $n-1 \times n-1$ block of $\tilde{H} = U^T H U$ for a orthogonal matrix U which diagonalizes $\hat{h}_1 \hat{h}_1^T$. Finally, making the substitutions for μ and $\tilde{c}$, Statement 2b is shown. □

2.9.4 Constrained Minimum

Proposition 2.5. *The minimum of J on the set $C = \{\theta : h(\theta) \geq 0\}$, for J and h given in Assumptions 2.1-2.2, is*

$$J_s^* = J^* + \frac{h_0^2}{2h_1^T H^{-1} h_1} u(-h_0)$$

where u is the unit step function.

Proof. We use the following convex analysis result: "A point θ_{smin} is the minimum of J relative to S if and only if $-\nabla J(x)$ is normal to S at θ_{smin}". This result is restated from Theorem 27.4

[74]. Additionally, one can find this result in [18] Section 4.2.3.

When $h_0 \geq 0$ the minimum of J on S is simply J^*.

When $h_0 < 0$, we use the normality condition to yield

$$ph_1 = H(\theta_{\text{smin}} - \theta^*) \tag{2.118}$$

for some p using the fact that $\nabla J(\theta) = H(\theta - \theta^*)$ and $\nabla h(\theta) = h_1$. Also,

$$h(\theta_{\text{smin}}) = h_0 + h_1^T(\theta_{\text{smin}} - \theta^*) = 0. \tag{2.119}$$

Solving, we get

$$p = -\frac{h_0}{h_1^T H^{-1} h_1}. \tag{2.120}$$

The value of J at θ_{smin} is

$$\begin{aligned} J(\theta_{\text{smin}}) &= J^* + \frac{1}{2}(\theta_{\text{smin}} - \theta^*)^T H(\theta_{\text{smin}} - \theta^*) \\ &= J^* + \frac{p^2}{2} h_1^T H^{-1} h_1 = J^* + \frac{h_0^2}{2h_1^T H^{-1} h_1}. \end{aligned} \tag{2.121}$$

Combining cases $h \geq 0$ and $h < 0$ we reach the result. □

2.9.5 Proof Theorem 2.1

Proof. From $\theta(t) - \theta^* = \tilde{\theta}(t) + S(t) = \tilde{\theta}(t) + O(a)$ we have

$$J(\theta(t)) = J^* + \frac{1}{2}\tilde{\theta}(t)^T H \tilde{\theta}(t) + O(a). \tag{2.122}$$

From Proposition 2.1, because $\|\tilde{\theta}(t) - \tilde{\theta}^{\text{a}}(t)\| = O(1/\omega)$ and $\tilde{\theta}^{\text{a}}(t) \to \tilde{\theta}^{\text{a,e}}$, then

$$\limsup_{t\to\infty} J(\theta(t)) = J^* + \frac{1}{2}\tilde{\theta}^{\text{a,e}T} H \tilde{\theta}^{\text{a,e}} + O(a) + O(1/\omega). \tag{2.123}$$

Before expanding $\theta^{\mathrm{a,e}}$ above, we use the Taylor series approximation for small δ to derive $\sqrt{1+r\delta} = 1 + O(\delta)$ (for constant r) and derive the expression

$$\begin{aligned}(-\operatorname{sgn}(h_0) + \sqrt{1+r\delta}) &= 1 - \operatorname{sgn}(h_0) + O(\delta) \\ &= 2(1 - u(h_0)) + O(\delta) \\ &= 2u(-h_0) + O(\delta) \end{aligned} \tag{2.124}$$

where u is the unit step function. We can write

$$\theta^{\mathrm{a,e}} = \frac{\|h_0\| u(-h_0)}{h_1^T H^{-1} h_1} H^{-1} h_1 + O(\delta) \tag{2.125}$$

from the approximation above and (2.43). Expanding part of (2.123) we have

$$\frac{1}{2}\tilde{\theta}^{\mathrm{a,e}T} H \tilde{\theta}^{\mathrm{a,e}} = \frac{h_0^2}{2h_1^T H^{-1} h_1} u(-h_0) + O(\delta). \tag{2.126}$$

Using Proposition 2.5, we achieve the result. □

2.9.6 Proof of Theorem 2.2

Proof. From Proposition 2.1 we have $||\theta(t) - (\tilde{\theta}^{\mathrm{a}}(t) + \theta^*)|| = O(1/\omega + a)$ for all $t \geq 0$. From Proposition 2.2 we have $||\tilde{\theta}^{\mathrm{a}}(t) - \tilde{\theta}_{\mathrm{r}}(t)|| = O(1/\omega_{\mathrm{f}})$ for all $t \geq 0$. Therefore $||\theta(t) - (\tilde{\theta}_{\mathrm{r}}(t) + \theta^*)|| = O(1/\omega_{\mathrm{f}} + 1/\omega + a)$ for all $t \geq 0$. Because h is Lipschitz, $||h(\theta(t)) - h(\tilde{\theta}_{\mathrm{r}}(t) + \theta^*)|| = O(1/\omega_{\mathrm{f}} + 1/\omega + a)$ for all $t \geq 0$. Using Proposition 2.3 we have the result. □

Acknowledgements

Chapter 2 has been submitted for publication with the title "Practical Safe Extremum Seeking with Assignable Rate of Attractivity to the Safe Set" A. Williams, M. Krstic, A. Scheinker. Chapter 2 also contains a partial adaptation of the work contained in the conference paper "Practically Safe Extremum Seeking" A. Williams, M. Krstic, A. Scheinker, presented at

the IEEE Conference on Decision and Control 2022. The dissertation author was the primary investigator and author of these two papers.

Chapter 3

Semiglobal Practically Safe ES

We introduce a safe extremum-seeking (Safe ES) algorithm which achieves the minimization of an unknown objective function while ensuring that an unknown, yet measured, control barrier function (CBF) remains above an arbitrarily small negative value for all time. In other words, "practical safety" is maintained during the entire period of convergence to the constrained extremum. Our design is based on quadratic program (QP) CBF style filters for safety, which is applied in an average and estimated sense. Using nonsmooth analysis tools, we guarantee semiglobal practical asymptotic (SPA) stability of the global constrained optimum, practical convergence to the safe set if starting in a condition violating the CBF, and practical safety for all time—semiglobally—if starting in safe set. The safety result of the paper is analogous with modern notions of SPA stability, guaranteeing that, for any small violation of safety, there exist design coefficients which guarantee that such a small violation is not exceeded. The paper outlines a set of sufficient conditions on the barrier and objective functions, and by way of a Lyapunov argument, we demonstrate that nonconvex constrained optimization problems can be solved. We present these results in the setting of a static map and a dynamical system. A simulation example illustrates the results.

3.1 Introduction

3.1.1 The Goal and How it is Achieved

This chapter presents an ES algorithm designed to minimize an unknown objective function $J(\theta)$ over parameters θ, while ensuring system safety. Safety, represented by an unknown, yet measured, function $h(\theta)$, is maintained by keeping h positive. Therefore our algorithm tackles the following problem:

$$\min_{\theta} J(\theta) \text{ subject to } h(\theta) \geq 0.$$

In essence, the classical unconstrained extremum seeking algorithm approximates the following dynamics, in an average sense, on the parameter: $\dot{\theta} = -\nabla J(\theta)$ i.e. it imparts a gradient descent term on the parameter dynamics in order to minimize J. In the framework of [10] using the QP-CBF safety filter designs, the term $-\nabla J(\theta)$ can be considered a "nominal control" and therefore a safety term u_s, can be algebraically computed in terms of $h(\theta), \nabla h(\theta), \nabla h(\theta)$ [100]. The safety filtered dynamics, $\dot{\theta} = -\nabla J(\theta) + u_s$, now satisfies a safety inequality $\dot{h}(\theta) + ch(\theta) \geq 0$ thereby guaranteeing that the safe set remain forward invariant. A prime achievement of ES comes from the fact that it is capable of estimating $\nabla J(\theta)$, and therefore it is also capable of estimating gradients of other unknown functions, given measurements, such as $\nabla h(\theta)$. This allows the safety term to be computed and applied (in an approximated and average sense) to the standard ES dynamics. In the end, the Safe ES algorithm approximates $\dot{\theta} = -\nabla J(\theta) + u_s$, where u_s is an approximation of the QP-derived safety term which would result if $\nabla h(\theta)$ was known.

3.1.2 Main Results of Semiglobal Safe ES

The framework from [64] and significant prior works [91, 92, 94] establishes semiglobal practical asymptotic (SPA) stability properties of extremum seeking. Using these results, we demonstrate that the Safe ES algorithm achieves SPA stability to the constrained optimum, and

furthermore we demonstrate that "practical safety" may be achieved in addition to practical stability. Practical safety guarantees that for any violation of safety, however small, there exist design constants which can achieve it for all time. We make use of the two time scales built into the design, in which the estimated gradients converge fast, relative to the convergence of the parameter θ. Due to the approximation that ES provides us, we therefore do not expect to exactly solve the constrained optimization problem. Overall the results are encompassed by the following three statements: 1) practical asymptotic stabilization—semiglobally 2) practical convergence to the safe set—semiglobally 3) practical safety for all time—semiglobally—if starting in the safe set.

Safe ES allows the designer two degrees of freedom concerning safety: 1) the choice of the c, and 2) the design of CBF h. Choosing c small slows the overall movement of the parameter towards the unsafe region when initialized in the safe set, and specifies the tradeoff between the favoring of safety and the favoring the objective. The design of h may be derived from any signal such that a sufficient safety margin is already built in. For example, consider some safe signal $s(t) \in [0, 1]$ which must never go below 0.5. Then one may conservatively design $h(t) = s(t) - 0.6$ if one is sure that the chosen perturbation signal and transient of the algorithm will not cause a large safety violation. Furthermore, the conditions on J and h we provide, allow for various kinds of nonconvex constrained optimization problems to be solved. We provide one such example at the end of this chapter.

3.1.3 Differences from ASfES

In the Chapter 2 we presented ASfES which was shown to have an assignable attractivity which was previously denoted with the parameter c. In "Safe ES", which is the name we give to the algorithm presented here, the parameter c has a slightly different characterization due to small differences between the algorithms. Due to the design of Safe ES we achieve a safety result which is essentially given by the inequality $h(\theta(t)) \geq h(\theta(0))e^{-ck\omega_{\mathrm{f}}(t)} + O(\delta)$ for all $t > 0$ – in contrast with the ASfES result of $h(\theta(t)) \geq h(\theta(0))e^{-ct} + O(\delta)$ for all $t > 0$. The adaptation gain, $k\omega_{\mathrm{f}}$,

appears in the time constant in the safe inequality of Safe ES, whereas it does not in ASfES. This means that the user of the Safe ES algorithm does not have control over the attractivity – which has been sacrificed in order to achieve theoretical semiglobality. It is an open problem to design a safe ES algorithm which achieves both assignable attractivity *and* (semi)global convergence.

3.1.4 Notation

We denote $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ as the set of positive and non-negative real numbers, respectively. For a differentiable function $Q : \mathbb{R}^n \to \mathbb{R}$ we denote the gradient as $\nabla Q : \mathbb{R}^n \to \mathbb{R}^n$. For a vector $v \in \mathbb{R}^n$, the notation $\|v\|$ denotes the Euclidean norm. The continuous function $\beta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class $\mathcal{K}$ if $\beta(0) = 0$ and it is strictly increasing. The continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class $\mathcal{KL}$ if it is strictly increasing in its first argument and strictly decreasing to zero in its second argument. The image of a function h is denoted by $\mathrm{Im}(h)$. The open ball around a point p is $B_r(p) = \{\theta \in \mathbb{R}^n : \|\theta - p\| < r\}$ and the closed ball is $\bar{B}_r(p)$. For some nonzero vector v we may denote the normalized vector as $\hat{v} = v/\|v\|$. The boundary and interior of a set C are denoted as $\mathrm{bd}(C)$ and $\mathrm{Int}(C)$ respectively. We use the term "SPA stability" to refer to the notion of semiglobal practical asymptotic stability [91]. A function $f(x, \epsilon)$ is $O(\epsilon)$ if for any compact set Ω there exists a positive pair (ϵ^*, k) such that $\|f(x, \epsilon)\| \leq k\epsilon$ for all $\epsilon \in (0, \epsilon^*]$ for all $x \in \Omega$. The convex hull of a set $\mathcal{S}$ is $\mathrm{co}(\mathcal{S})$. Given $a \in \mathbb{R}^n$ and $b \in \mathbb{R}^m$ we often use the notation $u = (a, b) = [a^T, b^T]^T \in \mathbb{R}^{n+m}$ when stacking vectors, and a slight abuse of notation when indexing components of a stacked vector so that $u_1 = a \in \mathbb{R}^n$ and $u_2 = b \in \mathbb{R}^m$.

3.2 Assumptions

We define

$$C = \{\theta \in \mathbb{R}^n : h(\theta) \geq 0\} \tag{3.1}$$

called the 'safe set'. We also define the notation of a superlevel set of h, parameterized by the parameter $\rho \leq 0$ as

$$C_\rho = \{\theta \in \mathbb{R}^n : h(\theta) \geq \rho, \rho \in \mathrm{Im}(h) \cap \mathbb{R}_{\leq 0}\}. \tag{3.2}$$

The sets of the form C_ρ always contain C along with some unsafe region given by ρ, a non-positive value in the image of h. We also use the following assumptions throughout.

Assumption 3.1 (Objective Function Conditions)**.** *The objective function $J : \mathbb{R}^n \to \mathbb{R}$ is differentiable with locally Lipschitz Jacobian and satisfies:*

1. *$\theta_c^* \in C$ is the unique constrained minimizer of J on C denoted as $J^* = J(\theta_c^*)$,*
2. *if there exists a θ such that $\nabla J(\theta) = 0$ for $\theta \in C$, then $\theta = \theta_c^*$.*

Assumption 3.2 (Barrier Function Conditions)**.** *The barrier (or safety) function $h : \mathbb{R}^n \to \mathbb{R}$ is differentiable with locally Lipschitz Jacobian and satisfies:*

1. *the safe set C is non-empty,*
2. *for any C_ρ, there exists a $L \in (0, \infty)$ such that $\|\nabla h(\theta)\| > L$ for $\theta \in \{h(\theta) \leq 0\} \cap C_\rho$.*

Assumption 3.3 (Optimizer Condition)**.** *If $\nabla h(\theta)^T \nabla J(\theta) = \|\nabla h(\theta)\| \, \|\nabla J(\theta)\|$ ($\nabla h(\theta)$ and $\nabla J(\theta)$ are collinear) for $\theta \in \mathrm{bd}(C)$, then $\theta = \theta_c^*$.*

Assumption 3.4 (Angle Condition)**.** *For any $\rho \in \mathrm{Im}(h) \cap \mathbb{R}_{<0}$, there exists a $r^* > 0$ and $f^* \in [0, 1)$ such that*

$$\frac{\nabla J(\theta)^T \nabla h(\theta)}{\|\nabla J(\theta)\| \, \|\nabla h(\theta)\|} \leq f^*, \tag{3.3}$$

for $\theta \in \{\rho \leq h(\theta) \leq 0\}$, $\theta \notin B_{r^*}(\theta_c^*)$.

Assumption 3.5 (Radial Unboundedness)**.** *The function $V(\theta) = \max\{-h(\theta), 0\} + \max\{J(\theta) - J^*, 0\}$ is positive definite and $\|\theta - \theta_c^*\| \to \infty \implies V \to \infty$.*

Assumption 3.6 (Bounded Levels of h)**.** *All sets C_ρ are compact.*

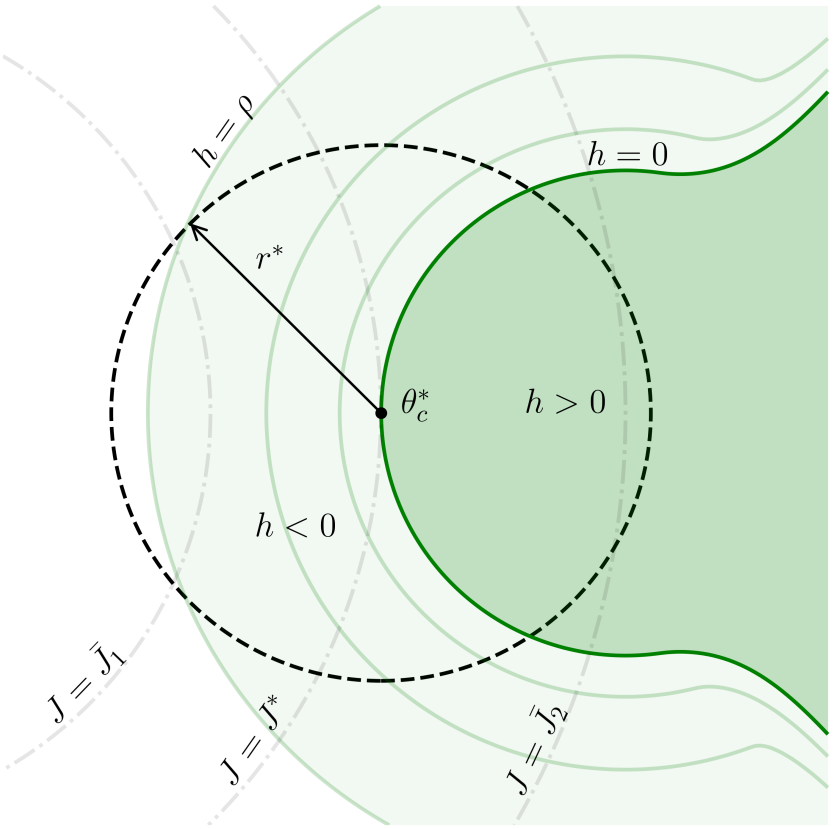

Figure 3.1. Depiction of sets and levels of h and J with $\bar{J}_1 < J^* < \bar{J}_2$. The safe set C is shaded in dark green, which is contained by the set C_ρ defined by both dark and light green shading. The radius r^* relates to the "angle condition" given in Assumption 3.4.

Assumption 3.7 (ES Constants)**.** *The design constants are chosen as* $\omega_{\mathrm{f}}, \omega_i, \delta, a, k, c > 0$*, where* ω_i/ω_j *are rational with frequencies* ω_i *chosen such that* $\omega_i \neq \omega_j$ *and* $\omega_i + \omega_j \neq \omega_k$ *for distinct* $i, j,$ *and* k*.*

In our analysis we use Assumptions 3.1, 3.2, and 3.3 to establish a unique equilibrium in the dynamics of a so called 'exact system', which our ES scheme approximates. This equilibrium is identically the constrained optimizer θ_c^*. Assumptions 3.4, 3.5, and 3.6 facilitate the subsequent Lyapunov analysis.

Assumption 3.3 ensures a single equilibrium by imposing constraints on the gradients of J and h at their boundary. This condition aligns with principles from optimization solutions, such as the 'method of Lagrange multipliers'. The gradient condition is also apparent in Fig. 3.1

in the tangency of the levels of J with $h = 0$.

Assumption 3.4 focuses on the cosine of the angle between $\nabla J(\theta)$ and $\nabla h(\theta)$ far away from the constrained equilibrium, but only in the unsafe set, $\theta \in \{\rho \leq h(\theta) \leq 0\}$. It aids in deriving a Lyapunov function with a universally negative time derivative, applicable, for instance, in scenarios with semi-infinite safe sets, as depicted in Fig. 3.1. Proposition 3.1 illustrates that for a large class of problems, quadratic and positive definite J with linear h, Assumption 3.4 holds. The proof is given in the appendix.

Proposition 3.1. *If J is quadratic, positive definite function and h is a linear function satisfying Assumption 3.2, then Assumption 3.4 is satisfied.*

Assumption 3.5 ensures that the objective function does not asymptotically plateau as θ extends away from the equilibrium in safe directions. Our main results hinge upon the satisfaction of either Assumptions 3.4-3.5 or Assumption 3.6. The latter, asserting all sets C_ρ are compact, exhibits sufficient strength to establish Semiglobal Practical Asymptotic (SPA) stability in the reduced system. However, for scenarios with semi-infinite safe sets, the Assumptions 3.4-3.5 become necessary for our main results to hold.

The assumptions listed here a minimal set of a non restrictive conditions we have gathered which guarantee the main results of the chapter. There are indeed slightly less restrictive assumptions that may be taken instead, relaxing some of the conditions. For example, Assumptions 3.4 is truly only necessary if some C_ρ is an unbounded set and $\|\nabla J(\theta)\|$ grows unbounded on C_ρ. If $\|\nabla J(\theta)\|$ does not grow unbounded on this C_ρ (a set which can be bounded or not), Assumption 3.4 is unnecessary.

3.3 Preliminaries on Nonsmooth Lyapunov Analysis

In this section we introduce concepts of regularity, the set-valued Lie derivative, and stability theorems. Nonsmooth analysis tools are necessary as we use the Lyapunov function $V_\alpha(\theta) = \max\{-\alpha h(\theta), 0\} + \max\{J(\theta) - J^*, 0\}$ for the main stability results of the chapter.

3.3.1 The Generalized Gradient and Set-Valued Lie Derivative

Definition 3.1. *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz and let Ω_V be a set of measure zero containing the points where V is not differentiable. The generalized gradient is defined as*

$$\partial V(x) = \operatorname{co}\left\{\lim_{i\to\infty} \nabla V(x_i) : x_i \to x, x_i \notin \Omega_V\right\}. \tag{3.4}$$

Regularity of a Lyapunov function proves to be very useful in calculating generalized derivatives. The classical definition of regularity can be seen in [21], but the following known class of regular functions is sufficient background for this chapter, rewritten from [21, Proposition 2.3.12].

Proposition 3.2. *For $k \in \{1,\dots,m\}$, let $g_k : \mathbb{R}^n \to \mathbb{R}$ be differentiable functions. Let $V : \mathbb{R}^n \to \mathbb{R}$ be defined as*

$$V(x) = \max\{g_k(x) : k \in \{1,\dots,m\}\}. \tag{3.5}$$

Then:

1. *V is regular and locally Lipschitz.*
2. *Let $I_V(x)$ denote the set of indices k for which $V(x) = g_k(x)$, then*

$$\partial V(x) = \operatorname{co}\{\nabla g_i(x) : i \in I_V(x)\}. \tag{3.6}$$

The generalized gradient has been defined and a rule given to calculate it for functions of the form (3.5).

We now define a set-valued form of the Lie derivative, utilizing the generalized gradient. It can be more generally defined for a differential inclusion, but we have defined it below for a differential equation,

$$\dot{x} = f(x) \tag{3.7}$$

with $f : \mathbb{R}^n \to \mathbb{R}^n$ locally Lipschitz.

Definition 3.2. *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz. The set-valued Lie derivative of V with respect to f in* (3.7) *is defined as*

$$\mathcal{L}_f V(x) = \{a \in \mathbb{R} : f(x)^T p = a \text{ for all } p \in \partial V(x)\}. \tag{3.8}$$

If $V(x)$ is differentiable at x, then the usual time derivative coincides with the set-valued Lie derivative $\dot{V} = \mathcal{L}_f V(x) = \{f(x)^T \nabla V(x)\}$. However, if $V(x)$ is not differentiable at x then $\mathcal{L}_f V(x)$ will represent the convex hull of the set of all possible time derivatives obtained at x when approaching x from all possible directions as defined by the generalized gradient.

3.3.2 Stability and Invariance Results

The next two stability results are nonsmooth versions of Lasalle's invariance principle and Nagumo's theorem. The following result is based on [14, 23].

Theorem 3.1. *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz and regular. Let $\mathcal{S} \subset \mathbb{R}^n$ be compact and positively invariant for* (3.7). *Suppose* $\max L_f V(x) \leq 0$ *for all $x \in \mathcal{S}$. Then all solutions of* (3.7), *with $x(t_0) \in \mathcal{S}$, converge to the largest invariant set $\mathcal{M}$ contained in*

$$\mathcal{S} \cap \overline{\{x \in \mathbb{R}^n : 0 \in L_f V(x)\}}. \tag{3.9}$$

The next result is a statement of Nagumo's theorem [63], more simply reformulated from [16, Section 4.2].

Theorem 3.2. *Suppose $\mathcal{S} \subseteq \mathbb{R}^n$ is a closed set. Then $\mathcal{S}$ is invariant with respect to* (3.7) *if and only if*

$$f(x) \in \mathcal{T}_{\mathcal{S}}(x) \text{ for all } x \in \mathrm{bd}(\mathcal{S}), \tag{3.10}$$

where $\mathcal{T}_{\mathcal{S}}(x)$ is Bouligand's tangent cone defined as

$$\mathcal{T}_{\mathcal{S}}(x) = \left\{ z : \liminf_{\tau \to 0} \frac{\text{dist}(x + \tau z, \mathcal{S})}{\tau} = 0 \right\} \tag{3.11}$$

with the distance from a set defined as $\text{dist}(y, \mathcal{S}) = \inf_{w \in \mathcal{S}} \|y - w\|$ *for some* $y \in \mathbb{R}^n$.

The tangent cone $\mathcal{T}_{\mathcal{S}}(x)$ essentially describes the set of vectors pointing tangent along the boundary of $\mathcal{S}$ or "inward". For example, consider $C = \{h(\theta) \geq 0\}$, for h with properties given in Assumption 3.2. Then for $\theta \in \text{bd}(C)$, $\mathcal{T}_C(\theta) = \{z : \nabla h(\theta)^T z \geq 0\}$.

The following lemma makes clear many of the assumptions and restrictions on J and h we have considered in this chapter.

Lemma 3.1. *Consider the locally Lipschitz function $V_\alpha : \mathbb{R}^n \to \mathbb{R}$ defined as*

$$V_\alpha(\theta) = \max\{-\alpha h(\theta), 0\} + \max\{J(\theta) - J^*, 0\}, \tag{3.12}$$

for some parameter $\alpha > 0$ under the Assumptions 3.1 - 3.3. Then

1. *V_α is regular and can be written as*

$$V_\alpha(\theta) = \max\left\{-\alpha h(\theta), J(\theta) - J^*, -\alpha h(\theta) + J(\theta) - J^*\right\}. \tag{3.13}$$

2. *Let either Assumption 3.4 - 3.5 hold, or 3.6 hold. Then for any $\rho < 0$, there is a sufficiently large $\alpha > 0$ such that the following holds for the set $\mathcal{R} = \{V_\alpha(\theta) \leq \bar{V}\} \cap C_\rho$ for any $\bar{V} > 0$:*

   (a) the tangent cone $\mathcal{T}_{\mathcal{R}}(\theta)$ is:

$$\mathcal{T}_{\mathcal{R}}(\theta) = \{v : \max\{\xi^T v : \xi \in \partial W(\theta)\} \leq 0\}, \tag{3.14}$$

for $\theta \neq \theta_c^*$, *with*

$$W(\theta) = \max\{-h(\theta) + \max\{\rho, -\bar{V}/\alpha\}, \\ J(\theta) - J^* - \bar{V}, \\ J(\theta) - J^* - \bar{V} - \alpha h(\theta)\}. \tag{3.15}$$

(b) the set $\mathcal{R} = \{\theta : W(\theta) \leq 0\}$ *is compact and* $\{W(\theta) = 0\} \equiv \text{bd}(\mathcal{R})$.

(c) for $W(\theta) = 0$, $0 \notin \partial W(\theta)$.

The above lemma shows the regularity of V_α, under the assumed properties of J and h, when α is chosen sufficiently large. It also provides simple expressions for tangent cones on sets involving the levels of V_α using mainly the results of [21], and useful properties of these sets.

3.4 Global Convergence of the Exact Algorithm

Before conducting analysis of the ES scheme for unknown functions, we study the optimization algorithm in its exact form in which we assume that all gradients are known, in order to find the appropriate Lyapunov function which will be used in the next section. Consider the following dynamics:

$$\dot{\theta} = F(\theta) = -\nabla J(\theta) + \frac{\nabla h(\theta)}{\|\nabla h(\theta)\|^2} \max\{\nabla J(\theta)^T \nabla h(\theta) - ch(\theta), 0\}. \tag{3.16}$$

Note that there is no risk of dividing by zero in the expression of (3.16) as $\nabla h(\theta) = 0 \implies h(\theta) > 0$ by Assumption 3.2. Therefore, $\lim_{\|\nabla h(\theta)\| \to 0} F(\theta) = -\nabla J(\theta)$ on compact sets.

The differential inequality $\dot{h} + ch \geq 0$ is commonly used to show the forward invariance of the safe set [10], where it is assumed that the initial condition of any given trajectory is safe. But it also shows attractivity to the safe set in the case where $h(\theta(t)) < 0$, for some t, implying that $\dot{h}(\theta(t)) > 0$ and therefore h is increasing. So unsafe trajectories become 'more safe' in an

exponential fashion. Therefore, the family of sets C_ρ is also positively invariant, and not just the case of $\rho = 0$. We state this formally below.

Proposition 3.3. *Under Assumption 3.2, the dynamics* (3.16) *satisfy* $\frac{dh(\theta(t))}{dt} + ch(\theta(t)) \geq 0$ *for all* $\theta \in \mathbb{R}^n$ *and* C *is forward invariant. Moreover, all sets* C_ρ *are forward invariant and* $h(\theta(t)) \geq h(\theta(t_0))e^{-ct}$ *for all* $\theta(t_0) \in \mathbb{R}^n$.

Proof. We can express the exact dynamics as (3.16) under Assumption 3.2 and simply compute

$$\frac{dh(\theta(t))}{dt} + ch(\theta) = -\nabla J(\theta)^T \nabla h(\theta) + ch(\theta) + \max\{\nabla J(\theta)^T \nabla h(\theta) - ch(\theta), 0\} \geq 0 \tag{3.17}$$

where we use the fact that $-x + \max\{x, 0\} \geq 0$ for any $x \in \mathbb{R}$. By the comparison principle, $h(\theta(t)) \geq h(\theta(t_0))e^{-ct}$ for all $\theta(t_0) \in \mathbb{R}^n$ and all sets C_ρ are forward invariant. □

This proposition will help us later establish global convergence of the algorithm for trajectories starting outside of C.

We now discuss the equilibrium of (3.16). The previous proposition has guaranteed that the equilibrium of the algorithm cannot lie in an unsafe region, as $\dot{h}(\theta) > 0$ for $h(\theta) < 0$. Therefore, consider solving for $F(\theta^e) = 0$ for $h(\theta^e) \geq 0$ under Assumptions 3.1, 3.2, and 3.3. Clearly if $\nabla J(\theta^e) = 0$ then $\theta^e = \theta_c^*$ by Assumption 3.1. If $\nabla J(\theta^e) \neq 0$ for $h(\theta^e) \geq 0$, computing the inner product $\nabla J(\theta^e)^T F(\theta^e)$ shows that $F(\theta^e) = 0$ only when the following two things are true at the same time: 1) $h(\theta^e) = 0$ and 2) $\nabla h(\theta^e)$ and $\nabla J(\theta^e)$ are collinear. Then, by Assumption 3.3 this implies $\theta^e = \theta_c^*$. All in all, we have that

$$F(\theta^e) = 0 \implies \theta^e = \theta_c^*. \tag{3.18}$$

It turns out that the inner product $\nabla J(\theta)^T F(\theta)$ is also always negative for $h(\theta) \geq 0$ and $\theta \neq \theta_c^*$. Therefore we can formulate the Lyapunov function $V_1(\theta) = J(\theta) - J(\theta^e)$ in order to prove

stability for $\theta \in C$. Computing $\dot{V}_1 = \nabla J(\theta)^T F(\theta)$ we have,

$$\dot{V}_1 = -\|\nabla J(\theta)\|^2 + \frac{\nabla J(\theta)^T \nabla h(\theta)}{\|\nabla h(\theta)\|^2} \max\{\nabla J(\theta)^T \nabla h(\theta) - ch(\theta), 0\} \tag{3.19}$$

which is clearly negative definite if $\nabla J(\theta)^T \nabla h(\theta) - ch(\theta) \leq 0$. If not,

$$\dot{V}_1 = -\nabla J(\theta)^T \left(I - \frac{\nabla h(\theta) \nabla h(\theta)^T}{\|\nabla h(\theta)\|^2} \right) \nabla J(\theta) - c \frac{\nabla J(\theta)^T \nabla h(\theta)}{\|\nabla h(\theta)\|^2} h(\theta). \tag{3.20}$$

The quadratic first term is negative (unless $\theta = \theta_c^*$) and the second term is negative semidefinite on $\theta \in \{h(\theta) \geq 0\} \cup \{\nabla J(\theta)^T \nabla h(\theta) - ch(\theta) > 0\}$, $\theta \neq \theta_c^*$. Because the set $C \coloneqq \{h(\theta) \geq 0\}$ is forward invariant, and V_1 is a smooth Lyapunov function on C, we have the following Lemma.

Lemma 3.2. *Let Assumptions 3.1-3.3 hold. The function $V_1(\theta) = J(\theta) - J(\theta^e)$ is a Lyapunov function for the equilibrium θ_c^* on C, yielding strictly $\dot{V}_1 < 0$ for all $\theta \in C \setminus \{\theta_c^*\}$. The dynamics* (3.16) *are asymptotically stable to the constrained minimizer θ_c^* for $\theta(t_0) \in C$.*

Proof. **Step 1: construct a forward invariant set $\mathcal{R} = \{V_1(\theta) \leq \bar{V}\} \cap C$ for any $\bar{V} > 0$.** We represent this set as $\mathcal{R} = \{W(\theta) \leq 0\}$ with $W(\theta) = \max\{V_1(\theta) - \bar{V}, -h(\theta)\}$. The set is compact, and can be shown to be forward invariant with Theorem 3.2 by showing $F(\theta) \in \mathcal{T}_R(\theta)$. The tangent cone can be expressed using generalized gradients of W:

$$\mathcal{T}_{\mathcal{R}}(\theta) = \{v : \max\{\xi^T v : \xi \in \partial W(\theta)\} \leq 0\}, \tag{3.21}$$

For details showing that $\mathcal{R}$ has the regularity properties necessary to prove the statement above, see the arguments in the proof of Statement 2a, Lemma 3.1 - we omit demonstrating these regularity properties due to space. Now, check the negativity of $F(\theta)^T p$ where $p \in \partial W(\theta)$ for $\theta \in \mathrm{bd}(\mathcal{R})$. If $\theta \notin \{h(\theta) = 0\}$ then $F(\theta)^T p = \dot{V}_1 < 0$, as shown in (3.19) and (3.20). If $\theta \in \{h(\theta) = 0\}$ then $p = -\lambda_1 \nabla h(\theta) + \lambda_2 \nabla J(\theta)$, for $\lambda_1 + \lambda_2 = 0, \lambda_1 \geq 0, \lambda_2 \geq 0$. Then $F(\theta)^T p = -\lambda_1 \dot{h}(\theta) + \lambda_2 \dot{V}_1(\theta) < 0$ by Proposition 3.3 and using the fact that $\dot{V}_1(\theta) < 0$. By Theorem 3.2, $\mathcal{R}$ is forward invariant

Step 2: apply Theorem 3.1. The result holds because $\dot{V}_1(\theta) < 0$ on the compact invariant set $\mathcal{R} \setminus \{\theta_c^*\}$. □

Because Theorem 3.1 requires a compact invariant set, the proof above uses the boundary of the safe set (a potentially unbounded set) and a level of the Lyapunov function $V_1 = \bar{V}$ in order to construct this compact invariant set. This idea is used again in the next lemma where we use the boundary $h = \rho$ of a set C_ρ and the levels of a different Lyapunov function V_α.

In light of Proposition 3.3 (attractivity to C) and Lemma 3.2 (convergence within C), it should be intuitive to the reader that all trajectories eventually converge on the equilibrium point θ_c^*. We can demonstrate this fact with the following Lyapunov argument on any invariant set in Lemma 3.3.

We use Theorem 3.1 in the following lemma to prove global asymptotic stability. In order to use Theorem 3.1 we first construct a *compact* invariant set $\mathcal{R}$ using the intersection of C_ρ and a sublevel set of the Lyapunov function V_α. A function W is used to define the set $\mathcal{R}$ and check that the dynamics $F(\theta)$ point 'inward' to $\mathcal{R}$ at it's boundary. The boundary is not smooth, which is why generalized gradients are key tool. Based on the assumptions on the functions J and h, an essential idea in achieving global stability is the fact that we can design the set $\mathcal{R}$ to be arbitrarily large, encompassing any initial condition.

Lemma 3.3. *Let Assumptions 3.1-3.3 hold, and either Assumptions 3.4-3.5 or Assumption 3.6. Consider the Lypaunov function*

$$V_\alpha(\theta) = \max\{-\alpha h(\theta), 0\} + \max\{J(\theta) - J(\theta_c^*), 0\}. \tag{3.22}$$

For any invariant set C_ρ, there exists $\alpha \in (0, \infty)$ such that $\max L_f V_\alpha(x) < 0$ *for $\theta \in C_\rho \setminus \{\theta_c^*\}$. The dynamics* (3.16) *are globally asymptotically stable to the constrained minimizer θ_c^*.*

Proof. **Step 1: show that for any set C_ρ there is a sufficiently large $\alpha > 0$ such that $\mathcal{R} = \{V_\alpha(\theta) \leq \bar{V}\} \cap C_\rho$ is compact and invariant for any $0 < \bar{V} < \infty$.** From Lemma 3.1, V_α is

regular and $\mathcal{R}$ is compact, therefore we can determine if $F(\theta)$ lies in $\mathcal{T}_{\mathcal{R}}(\theta)$ by showing that $F(\theta)^T p < 0$ for $p \in \partial W(\theta)$ and $\theta \in \{W(\theta) = 0\} \setminus \{\theta_c^*\}$ with

$$W(\theta) = \max\{-h(\theta) + \max\{\rho, -\bar{V}/\alpha\}, J(\theta) - J^* - \bar{V}, J(\theta) - J^* - \bar{V} - \alpha h(\theta)\}. \tag{3.23}$$

Also in Lemma 3.1 we showed that for sufficiently large $\alpha > 0$, $p \neq 0$ for $W(\theta) = 0$. Therefore we assume, by default, in the following cases of h that 1) $\alpha > 0$ is sufficiently large based on any fixed $\rho < 0$ and 2) $\theta \in \{W(\theta) = 0\} \setminus \{\theta_c^*\}$.

Case 1): $h(\theta) > 0$. Then $W(\theta) = J(\theta) - J^* - \bar{V}$ and $\partial W(\theta) = \{\nabla J(\theta)\} \neq \{0\}$. With $p = \nabla J(\theta)$ it is clear from (3.19) and (3.20) that $\dot{V}_1 = F(\theta)^T p < 0$ for any $\alpha > 0$.

Case 2): $h(\theta) = 0$. Then $\partial W(\theta) = \text{co}\{\nabla J(\theta), \nabla J(\theta) - \alpha \nabla h(\theta)\}$. For some $p \in \partial W(\theta)$, $p = \nabla J(\theta) - \alpha\lambda \nabla h(\theta)$, for $0 \leq \lambda \leq 1$. If $\nabla J(\theta)^T \nabla h(\theta) \leq 0$ then

$$F(\theta)^T p = -\|\nabla J(\theta)\|^2 + \alpha\lambda \nabla J(\theta)^T \nabla h(\theta) < 0 \tag{3.24}$$

for any $\alpha > 0$.

If $\nabla J(\theta)^T \nabla h(\theta) > 0$ then by Assumption 3.3,

$$F(\theta)^T p = \nabla J(\theta)^T \left(\frac{\nabla h(\theta)^T \nabla h(\theta)}{\|\nabla h(\theta)\|^2} - I \right) \nabla J(\theta) - \alpha\lambda \nabla J(\theta)^T \nabla h(\theta) < 0 \tag{3.25}$$

for any $\alpha > 0$.

Case 3): $h(\theta) < 0$. So $\partial W(\theta) = \text{co}\{-\nabla h(\theta), \nabla J(\theta) - \alpha \nabla h(\theta)\}$. For some $p \in \partial W(\theta)$, we have $p = -\lambda_1 \nabla h(\theta) + \lambda_2(-\alpha \nabla h(\theta) + \nabla J(\theta))$, with $0 \leq \lambda_1, \lambda_2 \leq 1$ and $\lambda_1 + \lambda_2 = 1$.

We define the lumped quantities

$$m(\theta) = \nabla h(\theta)^T \nabla J(\theta) - c h(\theta), \tag{3.26}$$

$$f(\theta) = \nabla J(\theta)^T \nabla h(\theta) / (\|\nabla J(\theta)\| \|\nabla h(\theta)\|) \leq 1. \tag{3.27}$$

If θ lies in a region where $m(\theta) \leq 0$ then $\nabla J(\theta)^T h(\theta) \leq ch(\theta) < 0$ and one can easily show $F(\theta)^T p < 0$. Otherwise, if $m(\theta) > 0$ then

$$F(\theta)^T p = c|h(\theta)| \left(-\lambda_1 + \lambda_2 \left(-\alpha + f(\theta) \frac{\|\nabla J(\theta)\|}{\|\nabla h(\theta)\|} \right) \right) - \lambda_2 \|\nabla J(\theta)\|^2 \left(1 - f^2(\theta) \right). \tag{3.28}$$

We consider θ to lie in a region such that $m(\theta) > 0$ in the following two sub-cases.

Case 3a): Assumption 3.6 holds. Then set C_ρ is compact so choose

$$\alpha > \sup_{\theta \in C_\rho} \frac{\|\nabla J(\theta)\|}{L}, \tag{3.29}$$

using Assumption 3.2 (which states $\|\nabla h(\theta)\| > L > 0$ for $h(\theta) < 0$). This yields $F(\theta)^T p < 0$ for all $\theta \in C_\rho \setminus \{\theta_c^*\}$, therefore $\mathcal{R}$ is forward invariant by Theorem 3.2. Note that this argument also holds if $\|\nabla J(\theta)\|$ is bounded on a potentially unbounded set C_ρ.

Case 3b): Assumptions 3.4 - 3.5 hold. Assume the non trivial scenario where $\|\nabla J(\theta)\|$ is unbounded on C_ρ. Then from Assumption 3.4, there exists a scalar f^* and set $B_{r^*}(\theta_c^*)$ such that $0 < f(\theta) \leq f^* < 1$ for all $\theta \notin B_{r^*}(\theta_c^*)$. Let $\tilde{f} = 1 - f^{*2}$ with $\tilde{f} \in (0, 1)$. Then for any $\theta \notin B_{r^*}(\theta_c^*)$,

$$F(\theta)^T p = -c|h(\theta)|(\lambda_1 + \alpha\lambda_2) + \lambda_2 \|\nabla J(\theta)\| \left(\frac{c|\rho|}{L} - \tilde{f} \|\nabla J(\theta)\| \right), \tag{3.30}$$

using Assumption 3.2. From the expression in the equation above, $\|\nabla J(\theta)\| > |\rho|c/L\tilde{f}$ implies $F(\theta)^T p < 0$ for $\theta \notin B_{r^*}(\theta_c^*)$ regardless of α. Therefore, consider $\|\nabla J(\theta)\| \leq |\rho|c/L\tilde{f}$. From (3.28) we have

$$F(\theta)^T p \leq c|h(\theta)| \left(-\lambda_1 + \lambda_2 \left(-\alpha + \frac{c|\rho|}{L^2 \tilde{f}} \right) \right) - \lambda_2 \|\nabla J(\theta)\|^2 \left(1 - f^2(\theta) \right) \tag{3.31}$$

for any θ such that $\|\nabla J(\theta)\| \leq |\rho|c/L\tilde{f}$. Therefore choosing

$$\alpha > \max \left\{ \sup_{\theta \in B_{r^*}(\theta_c^*)} \frac{\|\nabla J(\theta)\|}{L}, \frac{c|\rho|}{L^2 \tilde{f}} \right\} \tag{3.32}$$

yields $F(\theta)^T p < 0$ for all $\theta \in C_\rho \setminus \{\theta_c^*\}$.

So $F(\theta)^T p < 0$ for all $\theta \in C_\rho \setminus \{\theta_c^*\}$, and $\mathcal{R}$ is forward invariant by Theorem 3.2. Note that α does not depend on $\bar{V}$, see (3.29) and (3.32), although α in general depends on ρ. This is very important, because otherwise the set $\mathcal{R}$ cannot be guaranteed to be made arbitrarily large (by means of choosing a large $\bar{V}$) in order to encompass any initial condition of (3.16).

For any $\Delta > 0$ and initial condition, $\theta(t_0) \in \bar{B}_\Delta(\theta_c^*)$, we are always able to find a C_ρ such that $\bar{B}_\Delta(\theta_c^*) \subseteq C_\rho$. Then we choose $\alpha > 0$ as in the cases 1-3, as a function of ρ. Then we find a $\bar{V} > 0$ such that $\bar{B}_\Delta(\theta_c^*) \subseteq \mathcal{R}$, recalling that $\mathcal{R} = \{V_\alpha(\theta) \leq \bar{V}\} \cap C_\rho$. In the case 3b we use Assumption 3.5 to guarantee the existence of this $\bar{V}$. In Case 3a, any chosen set C_ρ is compact by assumption and $\bar{V}$ can be taken as the maximum of V_α on C_ρ. So any trajectory of (3.16) is contained on $\mathcal{R}$. See Fig. 3.2 for the depiction of the sets.

Step 2: show $\max L_f V_\alpha(\theta) < 0$ **for all** $\theta \in \mathcal{R} \setminus \{\theta_c^*\}$**.** Now that we have demonstrated, using the generalized gradient of W, that $\mathcal{R}$ is a compact forward invariant set, we must show that $\max L_f V_\alpha(\theta) < 0$ for all $\theta \in \mathcal{R} \setminus \{\theta_c^*\}$ in order to claim asymptotic stability to θ_c^* using Theorem 3.1. Using Proposition 3.2, for any $p \in \partial V_\alpha(\theta)$, p takes the same form as that of $p \in \partial W(\theta)$ using the the three cases in Step 1. This is due to the similar form of V_α in (3.13) and W in (3.15) and the rule for computing the generalized gradients in Proposition 3.2. We showed in Step 1 that $F(\theta)^T p < 0$ in each case, therefore for the same α chosen in Step 1, $\max L_f V_\alpha(\theta) < 0$ for all $x \in \mathcal{R} \setminus \{\theta_c^*\}$. Therefore the dynamics (3.16) are asymptotically stable to the constrained minimizer θ_c^*. □

The proof above relies on the invariance of C_ρ where we are guaranteed that for any unsafe initial condition $h(\theta(0)) = \rho$ we will never leave the set C_ρ and additionally, the boundaries of the sublevel set of Lyapunov function V_α. The choice of α is simple if sets C_ρ are compact as we know trajectories will never leave the bounded super-level sets of h, and the supremum in (3.29) exists. If C_ρ are not all compact, then we must have a radially unbounded V_α to argue that trajectories will never leave the bounded sublevel sets of V and some other set, C_ρ.

The Lyapunov function parameter α depends on the size of the set of the initial conditions, because it is dependant on the invariant set C_ρ. The parameter α used to prove stability of some initial condition $h(\theta(0)) = -1$ may be insufficient to prove the stability of the initial condition $h(\theta(0)) = -2$. But, a Lyapunov function with α chosen to prove the stability of the initial condition $h(\theta(0)) = -2$, will indeed be sufficient in showing the stability of $h(\theta(0)) = -1$.

This Lyapunov function has an interesting connection to recent literature and we note that a similar Lyapunov function for 'gradient flow' systems in a more general (nonlinear programming) setting was also discovered and can be used to show local stability [7].

3.5 Safe Extremum Seeking

We introduce the Safe ES scheme for analytically unknown functions h and J in this section. We will then derive a 'reduced model' following the framework of [64]. The reduced model is essentially derived from (3.33)-(3.37), making the filtered estimate equations (3.34)-(3.37) algebraic under the appropriate time transformation and singular perturbation. The authors in [64] showed that the reduced model must be SPA stable in a for the algorithm itself in (3.33)-(3.37) to be SPA stable. We demonstrate that indeed SPA stability can be shown with the help of an invariant set from the boundaries of a sub level set of the Lyapunov function used in Lemma 3.3, $\{V_\alpha(\theta) \leq \bar{V}\}$, and a sufficiently large set C_ρ. Then we present our notion of *practical safety* and extend the results to dynamical systems for a complete picture of the set of applications of this ES scheme.

3.5.1 Algorithm Design

We introduce the algorithm:

$$\dot{\hat{\theta}} = k\omega_{\mathrm{f}}(-G_J + \min\{\|G_h\|^{-2}, M^+\}\max\{G_J^T G_h - c\eta_h, 0\}G_h) \tag{3.33}$$

$$\dot{G}_J = -\omega_{\mathrm{f}}(G_J - (J(\hat{\theta}(t) + S(t)) - \eta_J)M(t)) \tag{3.34}$$

$$\dot{\eta}_J = -\omega_{\mathrm{f}}(\eta_J - J(\hat{\theta}(t) + S(t))) \tag{3.35}$$

$$\dot{G}_h = -\omega_{\mathrm{f}}(G_h - (h(\hat{\theta}(t) + S(t)) - \eta_h)M(t)) \tag{3.36}$$

$$\dot{\eta}_h = -\omega_{\mathrm{f}}(\eta_h - h(\hat{\theta}(t) + S(t))) \tag{3.37}$$

where the state variables $\hat{\theta}, G_J, G_h \in \mathbb{R}^n$, $\eta_J, \eta_h \in \mathbb{R}$. The overall the dimension of the system is $3n+2$. The map is evaluated at θ, defined by

$$\theta(t) \coloneqq \hat{\theta}(t) + S(t)\ . \tag{3.38}$$

The integer n denotes the number of parameters one wishes to optimize over. The design coefficients are $k, c, \delta, \omega_{\mathrm{f}}, M^+ \in \mathbb{R}_{>0}$. The perturbation signal S and demodulation signal M are given by

$$S(t) = a\left[\sin(\omega_1 t),\ \ldots\ , \sin(\omega_n t)\right]^T, \tag{3.39}$$

$$M(t) = \frac{2}{a}\left[\sin(\omega_1 t),\ \ldots\ , \sin(\omega_n t)\right]^T, \tag{3.40}$$

and contain additional design parameters $\omega_i, a \in \mathbb{R}_{>0}$.

3.5.2 Deriving the Reduced Model

First, we derive the reduce model. Defining

$$F_0(\xi) \coloneqq -\xi_1 + \min\{\|\xi_2\|^{-2}, M^+\}\max\{\xi_1^T\xi_2 - c\xi_4, 0\}\xi_3, \tag{3.41}$$

with

$$\xi^T \coloneqq [G_J^T, \eta_J, G_h^T, \eta_h], \tag{3.42}$$

$$\zeta^T(t,\theta,\xi,a) \coloneqq [(J(\theta)-\xi_2)M(t)^T, J(\theta), \tag{3.43}$$
$$(h(\theta)-\xi_4)M(t)^T, h(\theta)],$$

we can rewrite (3.33) - (3.37) as

$$\dot{\hat{\theta}} = k\omega_{\mathrm{f}} F_0(\xi), \tag{3.44}$$

$$\dot{\xi} = -\omega_{\mathrm{f}}(\xi - \zeta(t,\theta,\xi,a)), \tag{3.45}$$

recalling $\theta = \hat{\theta} + S(t)$. Letting $\tau = \omega_{\mathrm{f}} t$, the system in the new time scale is

$$\frac{d\hat{\theta}}{d\tau} = kF_0(\xi), \tag{3.46}$$

$$\frac{d\xi}{d\tau} = -\left(\xi - \zeta\left(\frac{\tau}{\omega_{\mathrm{f}}}, \hat{\theta} + S\left(\frac{\tau}{\omega_{\mathrm{f}}}\right), \xi, a\right)\right). \tag{3.47}$$

We can take the average of the system (see Proposition 3.4 in Appendix) to compute

$$\frac{d\tilde{\theta}_{av}}{d\tau} = kF_0(\xi_{av}), \tag{3.48}$$

$$\frac{d\xi_{av}}{d\tau} = -\left(\xi_{av} - \mu(\hat{\theta}_{av}, a)\right), \tag{3.49}$$

$$D(\hat{\theta}_{av}) \coloneqq [\nabla J(\hat{\theta}_{av})^T, J(\hat{\theta}_{av}), \nabla h(\hat{\theta}_{av})^T, h(\hat{\theta}_{av})]^T, \tag{3.50}$$

$$\mu(\hat{\theta}_{av}, a) \coloneqq D(\hat{\theta}_{av}) + O(a). \tag{3.51}$$

Making another time transformation $s = k\tau$ we have

$$\frac{d\hat{\theta}_{av}}{ds} = F_0(\xi_{av}), \tag{3.52}$$

$$k\frac{d\xi_{av}}{ds} = -\left(\xi_{av} - \mu(\hat{\theta}_{av}, a)\right). \tag{3.53}$$

Taking $k = 0$, we now derive the singularly perturbed (or "reduced") system with a quasi steady state

$$z_s \coloneqq \xi_{av} = \mu(\hat{\theta}_{av}, a). \tag{3.54}$$

Defining

$$y = \xi_{av} - \mu(\hat{\theta}_{av}, a), \tag{3.55}$$

the boundary layer system (with $\tau = s/k$) is

$$\frac{dy}{d\tau} = -\left(\xi_{av} - \mu(\hat{\theta}_{av}, a)\right) = -y. \tag{3.56}$$

The boundary layer system is UGAS uniformly in ξ_{av} and t_0. The reduced system is

$$\frac{d\theta_r}{ds} = F_0(\mu(\theta_r, a)) = F_0\left(D(\theta_r) + O(a)\right). \tag{3.57}$$

The reduced order model is an estimate of the original system (3.33)-(3.37) (in a different time scale) because we will later choose k to be small. A small k implies the *slowness* of the dynamics of (3.33) relative to (3.34)-(3.37). Therefore, the singular perturbation makes the filters in (3.34)-(3.37) *fast*, capturing the effect of a small k. Note that the singular perturbation is actually performed on an averaged version of the original system.

3.5.3 SPA Stability of the Reduced Model

The reduced system in (3.57) and (3.16) are similar with the distinct difference that the gradients and measurements of J and h are perturbed by a small $O(a)$ disturbance. We make the claim below that under a sufficiently small a, the reduced model is SPA stable. Using results from [64], this demonstrates that the algorithm itself in (3.33)-(3.37) to be SPA stable.

Lemma 3.4. *Let Assumptions 3.1-3.3, 3.7 hold. Also, let either 3.4 and 3.5 hold or 3.6 hold. Then there exists a $\beta_\theta \in \mathcal{KL}$ such that: for any positive pair Δ, ν, there exists a $a^* > 0$ such that*

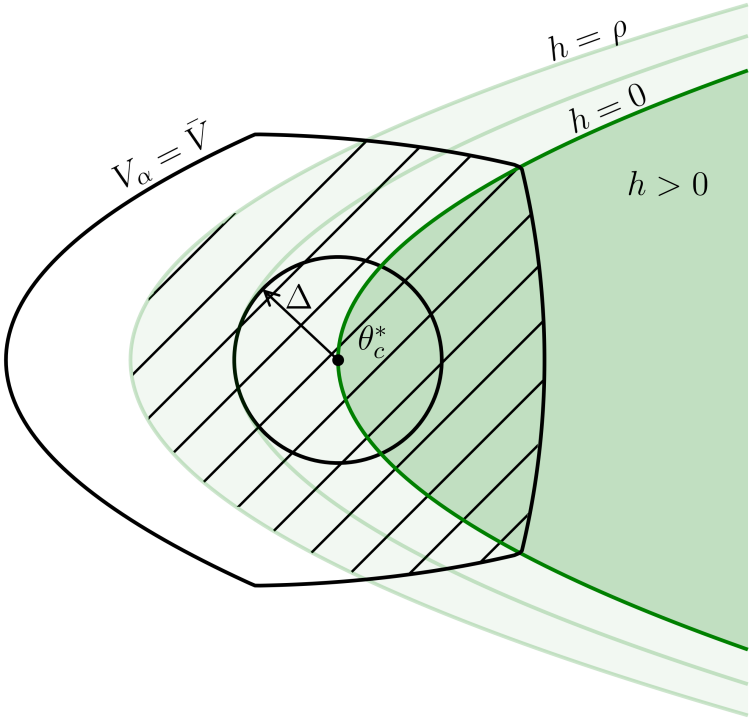

Figure 3.2. Depiction of $V_\alpha = \bar{V}$ subsuming the set of initial conditions in $\|\theta - \theta_c^*\| \leq \Delta$, where the shaded region marks $\mathcal{R}$.

for all $a \in (0, a^*)$ *the reduced system* (3.57) *satisfies*

$$\|\theta(t) - \theta_c^*\| \leq \beta_\theta(\|\theta(0) - \theta_c^*\|, s) + \nu \tag{3.58}$$

for all $\|\theta(0) - \theta_c^*\| \leq \Delta$.

Proof. To show the SPA stability of (3.57) we make use of analogous arguments used in the proof of Lemma 3.3, assuming either Assumptions 3.4-3.5 or Assumption 3.6. Consider a set of initial conditions around the constrained minimizer $\theta_r(t_0) \in \bar{B}_\Delta(\theta_c^*)$, then there exist a larger ball $\bar{B}_\Delta(\theta_c^*) \subset \bar{B}_{\Delta^*}(\theta_c^*)$ with $0 < \Delta < \Delta^*$. There must exist a C_ρ such that $\bar{B}_\Delta(\theta_c^*) \subset \bar{B}_{\Delta^*}(\theta_c^*) \subset C_\rho$. Then, there exists an $\alpha > 0$, corresponding to $V_\alpha(\theta)$, and a $\bar{V}$ such that such that $\mathcal{R} = \{V_\alpha(\theta) \leq \bar{V}\} \cap C_\rho$ is compact and invariant for (3.16) any $0 < \bar{V} < \infty$ by Lemma 3.3. Therefore we choose a $\bar{V}$ such that $\bar{B}_\Delta(\theta_c^*) \subset \bar{B}_{\Delta^*}(\theta_c^*) \subset \mathcal{R}$. See the depiction in Fig. 3.2 - the set $\mathcal{R}$ subsumes the ball of initial conditions along with a slightly larger ball around the Δ ball.

Under Assumptions 3.1 - 3.2, there exists a $M^+, a_{max} > 0$ such that for all $a \in (0, a_{max})$ the disturbed dynamics in (3.57) can be written as

$$\frac{d\theta_r}{ds} = F_0(D(\theta_r) + O(a)) = F(\theta_r) + O(a) \tag{3.59}$$

for all $\theta_r \in \mathcal{R}$, Lemma 3.5 in the appendix.

In other words, if one chooses a small enough and M^+ large enough, then the reduced dynamics (3.57) are simply the exact dynamics in (3.16) with an additive $O(a)$ disturbance. Furthermore, one can make the set $\mathcal{R}$ forward invariant with sufficiently small a under the same arguments used in Step 1 in the proof of Lemma 3.3. The vector $F(\theta_r) + O(a)$ lies in $\mathcal{T}_\mathcal{R}(\theta_r)$ if $(F(\theta_r) + O(a))^T p < 0$ for each θ_r on $\text{bd}(\mathcal{R})$. Since $F(\theta_r)^T p < 0$ on this compact region (Lemma 3.3) there exists a sufficiently small a that that $(F(\theta_r) + O(a))^T p < 0$ for $\theta \in \text{bd}(\mathcal{R})$. This proves semiglobal stability of (3.57). To achieve the 'practical' part of the SPA stability of (3.57), a further reduction of a is needed and an analogous argument to that used in Step 2 in the proof

of Lemma 3.3. Recall for $p \in \partial V_\alpha(\theta_r)$ we have that $F(\theta_r)^T p < 0$ and we are interested in the negativity of $F(\theta_r)^T p + O(a)^T p < 0$. For any $\delta > 0$ ball, $\bar{B}_\delta(\theta_c^*)$, there exists a $\bar{V}_\delta > 0$ such that $\{V_\alpha(\theta_r) \leq \bar{V}_\delta\} \subseteq \bar{B}_\delta(\theta_c^*)$ and a sufficiently small $a > 0$ such that trajectories of (3.57) converge to the set $\{V_\alpha(\theta_r) \leq \bar{V}_\delta\}$ by Theorem 3.2. This can be accomplished, with $\mathcal{R}' = \mathcal{R} \setminus \{V_\alpha(\theta_r) \leq \bar{V}_\delta\}$ by choosing a such that $\|O(a)^T p\| < |C|$ for $\theta_r \in \text{bd}(\mathcal{R})$ where C is the maximum of $F(\theta_r)^T p$ on the set $\mathcal{R}'$. Therefore (3.57) is SPA stable for small a. □

3.5.4 Main Results in Static Plants

We make the conclusion below with $z = \xi - \mu(\tilde{\theta}, a)$ and $\tilde{\theta} = \hat{\theta} - \theta_c^*$.

Theorem 3.3 (Semiglobal Practical Asymptotic Stability)**.** *Let Assumptions 3.1-3.3, 3.7 hold. Also, let either 3.4 and 3.5 hold or 3.6 hold. Then there exists $\beta_\theta, \beta_\xi \in \mathcal{KL}$ such that: for any positive pair (Δ, ν) there exist $M^+, \omega_{\text{f}}^*, a^* > 0$, such that for any $\omega_{\text{f}} \in (0, \omega_{\text{f}}^*)$, $a \in (0, a^*)$, there exists $k^*(a) > 0$ such that for any $k \in (0, k^*(a))$ the solutions to* (3.44)-(3.45) *satisfy*

$$\|\tilde{\theta}(t)\| \leq \beta_\theta \left(\|\tilde{\theta}(t_0)\|, k \cdot \omega_{\text{f}} \cdot (t - t_0) \right) + \nu, \tag{3.60}$$

$$\|z(t)\| \leq \beta_\xi \left(\|z(t_0)\|, \omega_{\text{f}} \cdot (t - t_0) \right) + \nu, \tag{3.61}$$

for all $\|(\tilde{\theta}(t_0), z(t_0))\| \leq \Delta$, and all $t \geq t_0 \geq 0$.

Sketch of proof: the proof of Theorem 3.3 follows the proof of Theorem 1 in [64] - we will summarize it this paragraph using terminology from [91]. The reduced system (3.57) is SPA stable in a uniformly in small a, demonstrated in Lemma 3.4. Using Lemma 1 in [92], this implies the average system (3.48) - (3.49) is SPA stable in k, a. This further implies the original system (3.46) - (3.47) is SPA stable in k, a, ω_{f} uniformly in k, a, ω_{f} in time scale s. Performing the appropriate transformations in from s to t we achieve the result of Theorem 3.3.

Clarifying some key points: this argument uses the two results, Lemma 2 in [91] and Lemma 1 in [92], which are fundamentally based on Theorem 1 in [94]. Theorem 1 in [94]

states several important key facts: 1) There is a time scale separation (by a factor k) between the convergence of parameter $\tilde{\theta}$ and the convergence of the boundary layer variable z. 2) The class $\mathcal{KL}$ functions governing the convergence of the reduced model and boundary layer model are the same functions which govern the slow and fast states of the actual system (see Remarks 8 and 15 in [94] for the choice of measuring functions). 3) We use may use the same $\mathcal{KL}$ function to describe the convergence of an average system as that of the actual system, see [94, Section V, "Classical Averaging"].

Equation (3.61) tell us about convergence of estimated quantities $\xi(t)$ to their exact values $D(\tilde{\theta}(t)+\theta_c^*)$. Using (3.51) we can write

$$e \coloneqq \xi - D(\tilde{\theta}+\theta_c^*), \tag{3.62}$$

$$\|e(t)\| \leq \beta_\xi(\|z(t_0)\|, \omega_{\mathrm{f}}(t-t_0)) + v + |O(a)|. \tag{3.63}$$

The variable e can be thought of as the error of the estimated quantities - various measurements and gradients of the static maps. Note, the functions β_θ, β_ξ are independent of $a, k, \omega_{\mathrm{f}}$ [91]. Because $\dot{\tilde{\theta}}$ is proportional to k (3.44), we can choose a k such that the change in $\tilde{\theta}(t)$ over some time interval can be small relative to the change in $e(t)$ over the same interval. The steady state offset in the error $e(t)$ can also be made small after some amount of time, with a small a. From Lemma 3.5 we know that there must be a small enough disturbance in the estimated quantities such that the dynamics can be written linearly (in the disturbance), but only after the transient of β_ξ has been sufficiently diminished. Therefore we make the following claim.

Theorem 3.4. (Semiglobal Practical Convergence to Safe Set and Semiglobal Practical Safety) *Let Assumptions 3.1–3.3, 3.7 hold. Also, let either 3.4 and 3.5 hold or 3.6 hold. For any $\Delta > 0$ and any $\delta > 0$: there exists $M^+, a^{**}, \omega_{\mathrm{f}}^{**} > 0$ such that for any $a \in (0, a^{**})$, $\omega_{\mathrm{f}} \in (0, \omega_{\mathrm{f}}^{**})$, there exists $k^{**}(a) > 0$ such that for any $k \in (0, k^{**}(a))$,*

$$h(\theta(t)) \geq h(\theta(t_0))e^{-ck\omega_{\mathrm{f}}(t-t_0)} + O(\delta) \tag{3.64}$$

for all $t \in [t_0, \infty]$ *and there exists a* $T \in [t_0, \infty]$ *such that*

$$h(\hat{\theta}(t)) \geq h(\theta(t_0)) - \delta \qquad \text{for } t_0 \leq t \leq T, \tag{3.65}$$

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_f c(t-T)} - \delta \qquad \text{for } T \leq t < \infty \tag{3.66}$$

for all $\|[\tilde{\theta}(t_0)^T, z(t_0)^T]\| \leq \Delta$

Proof. **Step 1: show for appropriate ranges of design coefficients, there exists a** $T > 0$ **such that for** $t \geq T$**, the dynamics of** $\hat{\theta}$ **simplify.**

Consider (3.44)-(3.45) written as,

$$\dot{\hat{\theta}} = k\omega_f F_0 \left(D(\hat{\theta}) + e(t)\right), \tag{3.67}$$

$$\|e(t)\| \leq \beta_\xi(\|z(t_0)\|, \omega_f(t - t_0)) + \nu + |O(a)|. \tag{3.68}$$

with the bounds (3.63) restated, for some $\Delta, \nu > 0$ using Theorem 3.3, which follows from the assumptions. Theorem 3.3 states that there exist ranges of coefficients of M^+, a, ω_f, k achieving the bound above, as well as the existence of an invariant set $\mathcal{R}$ which contains the set of initial conditions. Furthermore, there exists an ϵ^*, such that we can choose a $\nu > 0$ and restrict a to a smaller range such that $\nu + |O(a)| < \epsilon$ in (3.68), for $0 < \epsilon < \epsilon^*$ by Lemma 3.5.

In summary, for an appropriate ν and a (and appropriate ranges for the other design coefficients), there exist a time $T > 0$, after which (3.67) can be written as

$$\dot{\hat{\theta}} = k\omega_f \left(F(\hat{\theta}) + r(t)\right) \tag{3.69}$$

where $r(t) = O(\nu + a) = O(\epsilon)$.

We will summarize the contents of Step 1 in this paragraph. After a time $T > 0$, the dynamics of the parameter can be written as the sum of the 'exact' gradient based scheme and a time dependent disturbance. The time $T > 0$ depends on the first choosing v (by Theorem

3.3, this choice defines the ranges of coefficients) and then restricting a further. The time $T > 0$ represents the time during which the estimated quantities' error below a threshold $\epsilon^* > 0$, which exists by Lemma 3.5.

Step 2: construct an approximate safe inequality for $T \leq t \leq \infty$.

For any $\delta > 0$, take ν and a, satisfying the condition $\nu + |O(a)| < \epsilon$ in Step 1, and sufficiently small to satisfy $\|r(t)\| < \delta c/b$ because $r(t) = O(\nu + a)$. The quantity $b > 0$ is the supremum of $\|\nabla h(\hat{\theta})\|$ on $\mathcal{R}$. From the result of Step 1, there exists a $T > 0$ such that (3.69) is true.

Next, compute $\dot{h} + ck\omega_{\mathrm{f}} h$:

$$\dot{h}(\hat{\theta}) + ck\omega_{\mathrm{f}} h(\hat{\theta}) = k\omega_{\mathrm{f}}(F(\hat{\theta})^T \nabla h(\hat{\theta}) + ch(\hat{\theta}) + r(t)^T \nabla h(\hat{\theta})), \tag{3.70}$$

$$\dot{h}(\hat{\theta}) + ck\omega_{\mathrm{f}} h(\hat{\theta}) \geq -\delta ck\omega_{\mathrm{f}}, \tag{3.71}$$

because $F(\hat{\theta})^T \nabla h(\hat{\theta}) + ch(\hat{\theta}) \geq 0$ by Proposition 3.3 and because $r(t)^T \nabla h(\hat{\theta}) \geq -\delta c$.

The approximate safety inequality $\dot{h} + ck\omega_{\mathrm{f}} h + \delta ck\omega_{\mathrm{f}} \geq 0$ implies, by the comparison principle, for $T \leq t < \infty$,

$$h(\hat{\theta}(t)) \geq (h(\hat{\theta}(T)) + \delta) e^{-k\omega_{\mathrm{f}} c(t-T)} - \delta. \tag{3.72}$$

Step 3: summarize and refine bounds.

For the period of time $t_0 \leq t \leq T$, one can find a range of k, sufficiently small, to achieve an arbitrarily small change the parameter (without affecting $T > 0$). Using the range of k given by Theorem 3.3, take k small such that

$$h(\hat{\theta}(t)) \geq h(\theta(t_0)) - \delta, \tag{3.73}$$

for $t_0 \leq t \leq T$. Using $h(\hat{\theta}(T)) + \delta \geq h(\theta(t_0))$ which follows from above, we summarize the

bounds for all $t > 0$:

$$h(\hat{\theta}(t)) \geq h(\theta(t_0)) - \delta \qquad \text{for } t_0 \leq t \leq T, \tag{3.74}$$

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-T)} - \delta \qquad \text{for } T \leq t < \infty \tag{3.75}$$

Already, one can simply stop here and present the result as is. In summary, for any $\delta > 0$, one achieves the bounds above by means of further restriction of the design coefficients that are given in Theorem 3.3.

Step 5: when $h(\theta(t_0)) \geq 0$, bound $h(\hat{\theta}(t))$ with a single inequality.

Consider first (3.74). Because $e^{-k\omega_{\mathrm{f}}c(t-t_0)} \leq 1$ on $t_0 \leq t \leq T$, we multiply the term $h(\theta(t_0))$ by $e^{-k\omega_{\mathrm{f}}c(t-t_0)}$ in (3.74) to yield

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-t_0)} - \delta \text{ for } t_0 \leq t \leq T. \tag{3.76}$$

Next, consider (3.75), and using the fact that $e^{-k\omega_{\mathrm{f}}c(t_0-T)} \leq 1$, we have the following, similarly:

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-t_0)} - \delta \text{ for } T \leq t < \infty. \tag{3.77}$$

Step 6: construct a single inequality which holds for $t_0 \leq t < \infty$ and all $h(\theta(t_0))$.

Consider the Taylor expansion of $e^{-k\omega_{\mathrm{f}}c(t-t_0)}$ about $k = 0$ for $t_0 \leq t \leq T$, which is $e^{-k\omega_{\mathrm{f}}c(t-t_0)} = 1 + O(k)$. Then, $h(\theta(t_0)) = h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-t_0)} + O(k)$ and

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-t_0)} - \delta + O(k) \text{ for } t_0 \leq t \leq T, \tag{3.78}$$

following from (3.76). Using a similar argument with $e^{-k\omega_{\mathrm{f}}c(t_0-T)} = 1 + O(k)$, we achieve

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_{\mathrm{f}}c(t-t_0)} - \delta + O(k) \text{ for } T \leq t < \infty. \tag{3.79}$$

from (3.77). With an appropriately small k, we have

$$h(\hat{\theta}(t)) \geq h(\theta(t_0))e^{-k\omega_f c(t-t_0)} + O(\delta) \text{ for } t_0 \leq t < \infty. \tag{3.80}$$

□

Note that this result makes use of the fact that the time T is independent of k. Because the the estimator error convergence rate given by $\beta_\xi(\Delta, \omega_f(t-t_0))$ is independent of k, there will be some finite time at which estimators are very close to their true values. During this finite time, we are able to shrink the change in h by restricting k. This result also relies on Theorem 3.3 and a specific choice of ν to select particular intervals on a, ω_f, k which yield us the exponential safety result we seek. Therefore the intervals of a, ω_f, k given in Theorem 3.4 are perhaps a more strict set of intervals than the ones given in Theorem 3.3 - if the user desires the type of safety given in (3.64).

Furthermore, this safety result is analogous to the statement on stability. The stability result says that for any set of initial conditions, one should be able to adjust a, ω_f, k such that trajectories are ν-practically stable. The safety result says that for any set of initial conditions, one should be able to adjust a, ω_f, k such that trajectories are δ-practically stable.

3.5.5 Main Results for Dynamical Systems

We follow a similar outline of the transformations and problem formulation given in [64]. Consider the dynamic system

$$\dot{x} = f(x,u), \qquad y_h = g_h(x), \qquad y_J = g_J(x), \tag{3.81}$$

where $f : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ and $g_h, g_J : \mathbb{R}^n \to \mathbb{R}$. We will take y_J as our signal of performance and y_h as a measure of safety. Assume there exist the control law parameterized by θ as

$$u = \beta(x, \theta), \tag{3.82}$$

where $\theta \in \mathbb{R}^n$. The closed loop system is

$$\dot{x} = f(x, \beta(x, \theta)). \tag{3.83}$$

Assumption 3.8. *There exists a function $l : \mathbb{R}^n \to \mathbb{R}^n$ such that $f(x, \beta(x, \theta)) = 0$ if and only if $x = l(\theta)$. And, for each $\theta \in \mathbb{R}^n$, the equilibrium $x = l(\theta)$ of* (3.83) *is global asymptotically stable, uniformly in θ.*

Under a singular perturbation and time scale transformation, the performance map will take the form $J(\cdot) = g_J(l(\cdot))$ and the barrier function $h(\cdot) = g_h(l(\cdot))$. We allow the perturbation frequencies ω_i to be scaled by a parameter ω_s. This is necessary to achieve results for dynamical systems, as the perturbation frequencies must adjusted to be slower than the plant. Set the ith component of the perturbation and demodulation signals as $\omega_s t$ as $S_i(\omega_s t) = a \sin(\omega_i \omega_s t)$, $M_i(\omega_s t) = (2/a) \sin(\omega_i \omega_s t)$.

We must redefine ζ using the measurements of safety and performance as y_h, y_J.

$$\zeta^T(\omega_s t, \theta, \xi, a) \coloneqq [(y_J - \xi_2) M(\omega_s t)^T, y_J, (y_h - \xi_4) M(\omega_s t)^T, y_h]. \tag{3.84}$$

The definition of ξ and F_0 are given by (3.42) and (3.41). Under the transformation $\theta = \hat{\theta} + S(\omega_s t)$

we have the proposed scheme for dynamical systems as

$$\begin{aligned}
\dot{x} &= f(x, \beta(x, \hat{\theta} + S(\omega_s t))), \\
\dot{\hat{\theta}} &= k\omega_{\mathrm{f}}\omega_s F_0(\xi), \\
\dot{\xi} &= -\omega_{\mathrm{f}}\omega_s(\xi - \zeta(\omega_s t, \hat{\theta} + S(\omega_s t), \xi, a)).
\end{aligned} \tag{3.85}$$

Letting $\sigma = \omega_s t$ in the new time scale we write the system as

$$\begin{aligned}
\omega_s \frac{dx}{d\sigma} &= f(x, \beta(x, \hat{\theta} + S(\sigma))), \\
\frac{d\hat{\theta}}{d\sigma} &= k\omega_{\mathrm{f}} F_0(\xi), \\
\frac{d\xi}{d\sigma} &= -\omega_{\mathrm{f}}(\xi - \zeta(\sigma, \hat{\theta} + S(\sigma), \xi, a)).
\end{aligned} \tag{3.86}$$

We construct a reduced system by letting $\omega_s = 0$ and taking $x = l(\hat{\theta} + S(\sigma))$. Using the time scale $\tau = \sigma\omega_{\mathrm{f}}$:

$$\begin{aligned}
\frac{d\theta_s}{d\tau} &= kF_0(\xi_s), \\
\frac{d\xi_s}{d\tau} &= -\left(\xi_s - \zeta\left(\frac{\tau}{\omega_{\mathrm{f}}}, \theta_s + S\left(\frac{\tau}{\omega_{\mathrm{f}}}\right), \xi_s, a\right)\right).
\end{aligned} \tag{3.87}$$

Also, we have that $y_h = g_h(l(\hat{\theta} + S(\tau/\omega_{\mathrm{f}})))$ and $y_J = g_J(l(\hat{\theta} + S(\tau/\omega_{\mathrm{f}})))$ appearing as terms in ζ.

With $J(\theta) \coloneqq g_J(l(\theta))$ and $h(\theta) \coloneqq g_h(l(\theta))$, the system (3.87) is the same as (3.46) - (3.47). The system was demonstrated to be SPA stable. Therefore we make the following conclusion based on similar arguments using averaging and singular perturbation results in Theorem 3.3, and the safety result in Theorem 3.4. Let $z_1(t) = \xi(t) - \mu(\hat{\theta}(t), a)$, $z_2(t) = x(t) - l(\hat{\theta}(t) + S(\omega_s t))$, $\tilde{\theta}(t) = \hat{\theta}(t) - \theta_c^*$.

Theorem 3.5 (Stability and Safety for Dynamical Systems)**.** *With $J(\theta) \coloneqq g_J(l(\theta))$ and $h(\theta) \coloneqq g_h(l(\theta))$ let Assumptions 3.1-3.3, 3.7-3.8 hold. Also, let either 3.4 and 3.5 hold or 3.6 hold.*

Then there exist $\beta_\theta, \beta_\xi, \beta_x \in \mathcal{KL}$ *such that for any positive triple* (Δ, ν, δ) *there exists* $a^* > 0$ *and* $\omega_{\mathrm{f}}^* > 0$ *such that for any* $a \in (0, a^*)$ *and* $\omega_{\mathrm{f}} \in (0, \omega_f^*)$ *there exists* $k^*(a) > 0$ *such that for any* $k \in (0, k^*(a))$*, there exists* $\omega_s^*(a, \omega_{\mathrm{f}}, k) > 0$ *such that for any* $\omega_s \in (0, \omega_s^*(a, \omega_{\mathrm{f}}, k))$ *the solutions to* (3.85) *satisfy*

$$\|\tilde{\theta}(t)\| \leq \beta_\theta \left(\|\tilde{\theta}(t_0)\|, k\omega_{\mathrm{f}}\omega_s(t - t_0) \right) + \nu \tag{3.88}$$

$$\|z_1(t)\| \leq \beta_\xi \left(\|z_1(t_0)\|, \omega_{\mathrm{f}}\omega_s(t - t_0) \right) + \nu \tag{3.89}$$

$$\|z_2(t)\| \leq \beta_x \left(\|z_2(t_0)\|, (t - t_0) \right) + \nu \tag{3.90}$$

and

$$h(\theta(t)) \geq h(\theta(t_0)) e^{-ck\omega_{\mathrm{f}}\omega_s(t - t_0)} + O(\delta) \tag{3.91}$$

for all $\|(\tilde{\theta}(t_0), z_1(t_0), z_2(t_0))\| \leq \Delta$*, and all* $t \geq t_0 \geq 0$*.*

The proof of (3.88) - (3.90) follows that of [64]. The proof of (3.91) follows naturally using the arguments in Theorem 3.4 (the argument is made in exactly the same manner when considering the bounds (3.88) - (3.89) in time scale σ). The only difference in the result compared to the static case follows from the fact that the parameter dynamics are scaled by the factor $k\omega_{\mathrm{f}}\omega_s$ resulting in the time constant in (3.91) as $ck\omega_{\mathrm{f}}\omega_s$ instead of $ck\omega_{\mathrm{f}}$ as in (3.64).

3.6 Simulation

We present an example which demonstrates the capability of the Safe ES algorithm as well as the geometry of the Lyapunov function which proves the main stability result of this work. The unknown CBF is $h(\theta) = \cos(\theta_1) - \theta_2$. This barrier function yields a nonconvex, semi-infinite safe set. The objective function is $J(\theta) = \theta_1^2 + (\theta_2 - 2)^2$ with an unconstrained optimum at $\theta = (0, 2)$. See Fig. 3.3 for the depiction of the problem with the safe set in the bottom half of the plot.

One can check, for this problem, that our assumptions hold. First, we check Assumption 3.3 by computing the expression for $\nabla h(\theta)^T \nabla J(\theta) = \|\nabla h(\theta)\| \|\nabla J(\theta)\|$ for $h(\theta) = 0$. After some

algebra, we arrive at the following condition in θ_1:

$$q(\theta_1) = 2\theta_1 \sin(\theta_1) + 2\cos(\theta_1) - 4 + 2\sqrt{\theta_1^2 + (\cos(\theta_1) - 2)^2}\sqrt{\sin(\theta_1)^2 + 1} = 0, \tag{3.92}$$

which can be shown to have a unique solution at $\theta_1 = 0$ and a corresponding value $\theta_2 = 1$ which is the single constrained minimizer of J on C.

Next, we check that Assumption 4 holds – perhaps one of the more unintuitive assumptions. After some algebra,

$$\frac{\nabla J(\theta)^T \nabla h(\theta)}{\|\nabla J(\theta)\| \|\nabla h(\theta)\|} = d_1(\theta_2) \frac{\theta_1 + d_2(\theta_2)}{\sqrt{\theta_1^2 + d_2^2(\theta_2)}}, \tag{3.93}$$

where $d_1(\theta_2) = -\sin(\theta_2)/\sqrt{\sin(\theta_2)^2 + 1}$ and $d_2(\theta_2) = \theta_2 - 2$. Recall we must check (3.93) is strictly less than 1 for $\rho \leq h(\theta) \leq 0$, for any $\rho < 0$, far away from constrained minimizer of J on C - this implies θ_2 must be bounded, but θ_1 can be large. This means that $d_1(\theta_2), d_2(\theta_2)$ must be bounded with $d_2(\theta_2)$ having a more strict bound - approximately $\|d_2(\theta_2)\| \leq 0.707$. As $\|\theta_1\|$ grows large it is clear than (3.93) tends towards 0.707, strictly less than 1. Other assumptions are straightforward to check.

In Fig. 3.3 we plot trajectories of the Safe ES algorithm, $\hat{\theta}(t)$, from various initial conditions. Design constants are chosen as $a = 0.3$, $c = 1$, $k = 0.0005$, $\omega_f = 10$, $M^+ = 10^4$, $\omega_1 = 10$, and $\omega_2 = 13$. We also show the shape of the level sets of a hypothesized Lyapunov function V_α for a value of $\alpha = 20$. (Note: we prove in this work that a sufficiently large α exists to demonstrate SPA stability, but it is complicated to check analytically if $\alpha = 20$ is sufficient for the specific region of interest shown in Fig. 3.3.) We see that in the unsafe region, the contribution of h in V_α becomes noticeable in the wiggles of the level of V_α. This is because the practical safety result guarantees that trajectories "climb" the levels of h in order to escape to a safe region as $t \to \infty$, and therefore h appears as a term in V_α in the unsafe region. Recall the main result of practical safety, $h(\theta(t)) \geq h(\theta(t_0))e^{-ck\omega_f(t-t_0)} + O(\delta)$.

The trajectories shown are typical of safety filtered control, with $\hat{\theta}(t)$ slithering smoothly

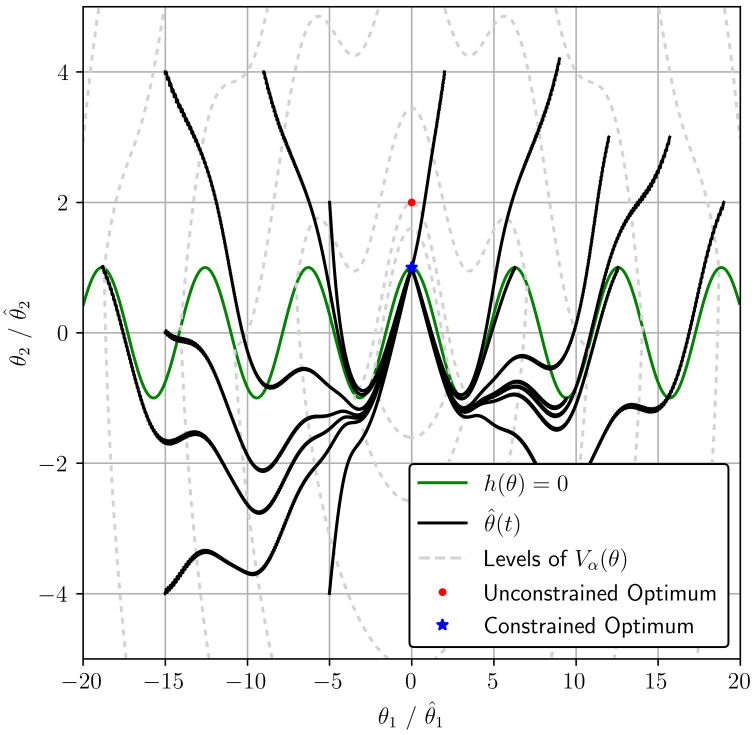

Figure 3.3. Various trajectories plotted of $\hat{\theta}$ along with the safe boundary $h(\theta) = 0$ and various level sets of V_α with $\alpha = 20$. The safe set is below the curve $h(\theta) = 0$.

around unsafe regions. The geometric conditions on the static maps of the problem also allow the constrained optimizer to be found, globally, with minor violations of safety along the way. Trajectories that start safe (or become safe) do not wander far out of the safe set for the rest of time, and all trajectories converge to a small region around the constrained optimum.

3.7 Conclusion

This chapter presents Safe ES, an ES controller designed to minimize an unknown objective function while ensuring the practical positivity of an unknown, yet measured CBF. In our approach, we add an approximation of a QP-based safety filter in the dynamics of the controller, advancing the current set of ES methodologies used to solve nonconvex constrained optimization problems. We use nonsmooth analysis tools and a key Lyapunov argument to help

us achieve our main theoretical results. The notion of practical safety, akin to practical stability, uses known results which tell us that the estimates of the unknown gradients of the problem converge relatively fast (with respect to the parameter).

Future work may include understanding the role of delays in the measurements on the safety of the system, and handing multiple constraints including equality constraints.

3.8 Chapter Appendix

3.8.1 Proof of Proposition 3.1

Proof. Let $J(\theta) = \frac{1}{2}\theta^T P\theta$ with $P > 0$ and $h(\theta) = h_1^T\theta + h_0$, with some constrained minimizer θ_c^*. We assume the the quadratic J has a global minimizer at 0 w.l.o.g. as we can always shift the coordinates to the form assumed.

Suppose Assumption 3.4 is not satisfied. Then $\exists\rho < 0$ for all $r > 0$ for all $f \in [0, 1)$ such that

$$f < \hat{h}_1^T \frac{P\theta}{\|P\theta\|} \leq 1, \tag{3.94}$$

for $\theta \in \{\rho \leq h_1^T\theta + h_0 \leq 0\} \cap \{\|\theta - \theta_c^*\| \geq r\}$. Note $\hat{h}_1 \coloneqq h_1/\|h_1\|$. Take $r = 1$ and $f = 1 - \frac{1}{8}(\frac{\lambda_{\min}}{\lambda_{\max}})^2$ where $0 < \lambda_{\min} \leq \lambda_{\max}$ are the smallest and largest eigenvalues of P. For two unit vectors u, v we can write $\|u - v\| = \sqrt{2(1 - u^T v)}$. Therefore, with the inner product bound in (3.94) and the choice of f, we have

$$\left\| \hat{h}_1 - \frac{P\theta}{\|P\theta\|} \right\| = \sqrt{2\left(1 - \hat{h}_1^T \frac{P\theta}{\|P\theta\|}\right)} \in [0, \lambda_{\min}/2\lambda_{\max}). \tag{3.95}$$

The closeness relation between the two unit vectors $\hat{h}_1$ and $P\theta/\|P\theta\|$ implies that, for any point θ there is a unit vector $\hat{d} = \hat{d}(\theta)$ such that

$$\hat{h}_1 = P\theta/\|P\theta\| + \epsilon\hat{d} \tag{3.96}$$

for $\epsilon \in [0, \lambda_{\min}/2\lambda_{\max})$.

Because $\theta \in \{\rho \leq h_1^T\theta + h_0 \leq 0\}$ we have

$$\begin{aligned}
&\rho - h_0 \leq \|h_1\|\|\theta\|\hat{\theta}^T\hat{h}_1 \leq -h_0, \\
&\rho - h_0 \leq \|h_1\|\|\theta\|\left(\frac{\hat{\theta}^T P\theta}{\|P\theta\|} + \epsilon\hat{\theta}^T\hat{d}\right) \leq -h_0,
\end{aligned} \tag{3.97}$$

substituting $\hat{h}_1$. The following quantity is strictly positive,

$$\frac{\hat{\theta}^T P\theta}{\|P\theta\|} + \epsilon\hat{\theta}^T\hat{d} > \frac{\lambda_{\min}}{\lambda_{\max}} - \frac{\lambda_{\min}}{2\lambda_{\max}} = \frac{\lambda_{\min}}{2\lambda_{\max}} > 0. \tag{3.98}$$

Therefore (3.97) can not hold for sufficiently large $\|\theta\|$ because $\theta \in \{\|\theta - \theta_c^*\| \geq 1\}$. □

3.8.2 Proof of Lemma 3.1

Proof. **Statement (1)** follows from that fact that $\max\{a,0\} + \max\{b,0\} = \max\{a,b,a+b,0\}$ for any real a, b. In our case $a = -\alpha h(\theta)$ and $b = J(\theta) - J^*$. We also have that $\max\{-\alpha h(\theta), J(\theta) - J^*, -\alpha h(\theta) + J(\theta) - J^*, 0\} = \max\{-\alpha h(\theta), J(\theta) - J^*, -\alpha h(\theta) + J(\theta) - J^*\}$ because $h(\theta) \geq 0$ implies $J(\theta) - J^* \geq 0$, by Assumption 3.1. By Proposition 2.3.12 in [21], V_α is regular because $h(\theta), J(\theta)$ each are regular because they are differentiable, and V_α can be expressed as a point-wise maximum of regular functions.

Statement (2a) follows directly from Proposition 2.1.2 and Theorem 2.4.7 in [21], if two conditions are met. The first condition is that the set $\mathcal{R}$ can be written as a sublevel set of a Lipschitz and regular function, as in $\mathcal{R} = \{W(\theta) \leq 0\}$. The second condition is that for $W(\theta) = 0$, $0 \notin \partial W(\theta)$.

First, let $\tilde{J}(\theta) = J(\theta) - J^*$. We can express $\mathcal{R}$ as

$$\mathcal{R} = \{\theta : W(\theta) \leq 0\}$$

where

$$W(\theta) = \max\{-h(\theta) + \rho^*, \tilde{J}(\theta) - \bar{V}, \tilde{J}(\theta) - \bar{V} - \alpha h(\theta)\},$$

and $\rho^* = \max\{\rho, -\bar{V}/\alpha\}$. So, W is Lipschitz and regular by Proposition 2.3.12 [21].

Now we must show for $W(\theta) = 0$, $0 \notin \partial W(\theta)$. We consider θ lying in the three sets spanning $\mathbb{R}^n$. We use Proposition 3.2 to compute the generalized gradients of W.

Case 1) $h(\theta) > 0$. Then $W(\theta) = \tilde{J}(\theta) - \bar{V}$ and $\partial W(\theta) = \{\nabla J(\theta)\} \neq \{0\}$ by Assumption 3.1.

Case 2) $h(\theta) = 0$. Then $W(\theta) = \tilde{J}(\theta) - \bar{V} = \tilde{J}(\theta) - \bar{V} - \alpha h(\theta)$. Therefore we can write $\partial W(\theta) = \text{co}\{\nabla J(\theta), \nabla J(\theta) - \alpha \nabla h(\theta)\}$. For some $p \in \partial W(\theta)$, $p = \nabla J(\theta) - \alpha\lambda \nabla h(\theta)$, for $0 \leq \lambda \leq 1$. Because $\theta \neq \theta_c^*$, by Assumption 3.3 we have that $p \neq 0$.

Case 3) $h(\theta) < 0$. So $\partial W(\theta) = \text{co}\{-\nabla h(\theta), \nabla J(\theta) - \alpha \nabla h(\theta)\}$. For some $p \in \partial W(\theta)$, we have the expression $p = -\lambda_1 \nabla h(\theta) + \lambda_2(-\alpha \nabla h(\theta) + \nabla J(\theta))$, with $0 \leq \lambda_1, \lambda_2 \leq 1$ and $\lambda_1 + \lambda_2 = 1$.

If Assumption 3.6 holds, clearly $p \neq 0$ for sufficiently large α because $\|\nabla h(\theta)\| > L$ by Assumption 3.2. Choosing $\alpha > L^{-1} \max_{\theta \in C_\rho} \|\nabla J(\theta)\|$, guarantees that $-\alpha \nabla h(\theta) + \nabla J(\theta) \neq 0$ and therefore $p \neq 0$.

If Assumptions 3.4 - 3.5 hold, then $\nabla h(\theta) \neq s \nabla J(\theta)$ for $s > 0$ outside some ball around θ_c^*. We can choose $\alpha > L^{-1} \sup_{\theta \in B_{r^*}(\theta_c^*)} \|\nabla J(\theta)\|$. Then $p \neq 0$. This demonstrates **Statement (2a)** and **Statement (2c)**.

Statement (2b): $\mathcal{R}$ is compact due to either Assumption 3.4- 3.5 holding, or 3.6 holding.

Now we prove $\{W(\theta) = 0\} \equiv \text{bd}(\mathcal{R})$. **Step 1**: show if $\theta \in \text{bd}(\mathcal{R})$ then $\theta \in \{W(\theta) = 0\}$. By contradiction, suppose $\theta \in \text{bd}(\mathcal{R})$ and $\theta \in \{W(\theta) < 0\}$. By continuity of W, there exists a $B_\epsilon(\theta)$ such that $W(y) < 0$ for all $y \in B_\epsilon(\theta)$. Therefore $B_\epsilon(\theta) \subseteq \{W(\theta) < 0\}$. But $B_\epsilon(\theta) \cap \mathcal{R}^c = B_\epsilon(\theta) \cap \{W(\theta) > 0\} = \emptyset$, which contradicts $\theta \in \text{bd}(\mathcal{R})$. **Step 2**: show if $\theta \in \{W(\theta) = 0\}$ then $\theta \in \text{bd}(\mathcal{R})$. By contradiction, suppose $\theta \in \{W(\theta) = 0\}$ and $\theta \notin \text{bd}(\mathcal{R})$. Then $\theta \notin \bar{\mathcal{R}}$ or $\theta \in \text{Int}(\mathcal{R})$. If $\theta \notin \bar{\mathcal{R}}$ then $\theta \notin \mathcal{R}$ because $\mathcal{R}$ is closed, which implies $W(\theta) > 0$ which is a contradiction. If $\theta \in \text{Int}(\mathcal{R})$ then there exists a $B_\epsilon(\theta)$ such that $B_\epsilon(\theta) \subseteq \text{Int}(\mathcal{R}) \subseteq \mathcal{R}$. Since $\theta \in \{W(\theta) = 0\}$ then

$W(\theta)$ achieves a local maximum on $B_\epsilon(\theta)$ which implies $0 \in \partial W(\theta)$ by Proposition 2.3.2 in [21], which is a contradiction as $0 \notin \partial W(\theta)$ by Statement (2a). □

3.8.3 Lemma 3.5

We group the quantities as the function Q,

$$Q(\theta)^T \coloneqq [\nabla J(\theta)^T, \nabla h(\theta)^T, h(\theta)]. \tag{3.99}$$

representing (3.16) as a function of the estimated variables $F = F(Q(\theta))$. And with a small disturbance $w(t)^T = [w_1(t)^T, w_2(t)^T, w_3(t)]$ we can consider a system, with disturbances to Q, written as

$$\dot{\theta} = F(Q(\theta) + w(t)). \tag{3.100}$$

Lemma 3.5 (Additive Disturbance)**.** *Under Assumptions 3.1 - 3.2, for any compact set* Ω *there exists a* $M^+, \epsilon^* > 0$ *such that for all* $\epsilon \in (0, \epsilon^*)$ *and* $\|w\| < \epsilon$*, the disturbed dynamics in* (3.100) *can be written as*

$$F(Q(\theta) + w(t)) = -\nabla J(\theta) + \frac{\nabla h(\theta)}{\|\nabla h(\theta)\|^2} \max\{\nabla J(\theta)^T \nabla h(\theta) - c h(\theta), 0\} + O(\epsilon). \tag{3.101}$$

for all $\theta \in \Omega$.

Proof. **Step 1 (Rewrite Dynamics, Grouping Terms):** Note that terms $w(t), w_i(t) = O(\|w\|)$

for $i \in \{1,2,3\}$. We write the perturbed dynamics $F(Q(\theta)+w(t)) = F_w$ as

$$\begin{aligned} F_w =& -\nabla J(\theta) - w_1(t) + \min\{(\nabla h(\theta)+w_2(t))^{-2}, M^+\}\cdot \\ & \max\{(\nabla J(\theta)+w_1(t))^T(\nabla h(\theta)+w_2(t)) - \\ & c(h(\theta)+w_3(t)),0\}(\nabla h(\theta)+w_2(t)), \\ =& -\nabla J(\theta) + \min\{(\nabla h(\theta)+w_2(t))^{-2}, M^+\}\cdot \\ & \max\{\nabla J(\theta)^T\nabla h(\theta) - ch(\theta),0\}\nabla h(\theta) + O(\|w\|), \end{aligned} \tag{3.102}$$

where we group terms proportional to w_i as $O(\|w\|)$ and note that the max term is proportional (or zero) to its' argument.

Step 2 (Find Large Enough M^+, Small Enough w): Next, we find a suitably large M^+ such that the max term is not active when $\min\{(\nabla h(\theta)+w_2(t))^{-2}, M^+\} = M^+$. First, Consider the infimum

$$G = \inf_{\theta\in\mathcal{S}} \|\nabla h(\theta)\| > 0 \tag{3.103}$$

$$\mathcal{S} = \Omega \cap \{\nabla J(\theta)^T\nabla h(\theta) - ch(\theta) \geq 0\}. \tag{3.104}$$

We have that $G > 0$ because for $\{h(\theta) \leq 0\} \cap \mathcal{S}$, $\|\nabla h(\theta)\| > L > 0$ (Assumption 3.2) and for $\mathcal{S} \setminus \{h(\theta) \leq 0\}$ we have that $\nabla J(\theta)^T\nabla h(\theta) > 0$ which implies $\|\nabla h(\theta)\| > 0$.

Now, we choose $\|w_2(t)\| < G/2$ and $M^+ > 4/G^2$. Therefore,

$$\frac{1}{\|\nabla h(\theta)+w_2(t)\|^2} \leq \frac{1}{(G/2)^2} = \frac{4}{G^2} < M^+ \tag{3.105}$$

for $\theta \in \mathcal{S}$. Taking $\epsilon^* = G/2$, we can express the dynamics without the min term as

$$\begin{aligned} F_w = -\nabla J(\theta) + \frac{\max\{\nabla J(\theta)^T\nabla h(\theta) - ch(\theta),0\}}{\|\nabla h(\theta)+O(\epsilon)\|^2}\nabla h(\theta) + \\ O(\epsilon), \end{aligned} \tag{3.106}$$

for $0 < \epsilon < \epsilon^*$. Note also the bound for the term

$$\frac{1}{\|\nabla h(\theta) + O(\epsilon)\|^2} < M^+ \tag{3.107}$$

holding for $\theta \in \mathcal{S}$, for all $0 < \epsilon < \epsilon^*$. In (3.106), we have now also taken the whole vector $w(t)$ as ϵ-small for simplicity. Note we also could have chosen $\epsilon^* = \mu G$ (and the bound on $\|w_2(t)\|$) for any $\mu \in (0, 1)$. If taking μ to be larger than $1/2$ (we took it as exactly $1/2$), this would have forced a larger choice of M^+.

Step 3 (Find an $\mathbf{O}(\epsilon)$ Approximation of $\|\nabla \mathbf{h}(\theta) + \mathbf{O}(\epsilon)\|^{-2}$): We now state two results. The first follows from the triangle inequality, and then squaring both sides of the resulting inequality:

$$\begin{aligned} &\|\nabla h(\theta) + O(\epsilon)\| \leq \|\nabla h(\theta)\| + \|O(\epsilon)\|, \\ \implies &\|\nabla h(\theta) + O(\epsilon)\|^2 \leq \|\nabla h(\theta)\|^2 + \|O(\epsilon)\|^2 + \\ &\qquad\qquad 2\|O(\epsilon)\|\|\nabla h(\theta)\|, \\ \implies &\|\nabla h(\theta) + O(\epsilon)\|^2 - \|\nabla h(\theta)\|^2 = O(\epsilon). \end{aligned} \tag{3.108}$$

The second next result uses the previous result, and determines a bound on $d(\theta)$, with

$$d(\theta) = \left\| \frac{1}{\|\nabla h(\theta) + O(\epsilon)\|^2} - \frac{1}{\|\nabla h(\theta)\|^2} \right\|. \tag{3.109}$$

We have,

$$d(\theta) = \left\| \frac{\|\nabla h(\theta)\|^2 - \|\nabla h(\theta) + O(\epsilon)\|^2}{\|\nabla h(\theta) + O(\epsilon)\|^2 \|\nabla h(\theta)\|^2} \right\|, \tag{3.110}$$

$$= \left\| \frac{\|\nabla h(\theta)\|^2 - \|\nabla h(\theta)\|^2 + O(\epsilon)}{\|\nabla h(\theta) + O(\epsilon)\|^2 \|\nabla h(\theta)\|^2} \right\|, \tag{3.111}$$

$$= \frac{\|O(\epsilon)\|}{\|\nabla h(\theta) + O(\epsilon)\|^2 \|\nabla h(\theta)\|^2},$$

$$\implies \frac{1}{\|\nabla h(\theta) + O(\epsilon)\|^2} - \frac{1}{\|\nabla h(\theta)\|^2} = O(\epsilon), \tag{3.112}$$

for $\theta \in \mathcal{S}$. We used the fact in the final line that $1/(\|\nabla h(\theta) + O(\epsilon)\|^2 \|\nabla h(\theta)\|^2) < M^+/G^2$ is bounded (see (3.107)) for $\theta \in \mathcal{S}$. We now use (3.112) to express $1/(\|\nabla h(\theta) + O(\epsilon)\|^2)$ in (3.106) as

$$\begin{aligned} F_w = &-\nabla J(\theta) + \max\{\nabla J(\theta)^T \nabla h(\theta) - c h(\theta), 0\} \cdot \\ &\left(\frac{1}{\|\nabla h(\theta)\|^2} + O(\epsilon) \right) \nabla h(\theta) + O(\epsilon), \end{aligned} \tag{3.113}$$

$$\begin{aligned} = &-\nabla J(\theta) + \frac{\max\{\nabla J(\theta)^T \nabla h(\theta) - c h(\theta), 0\}}{\|\nabla h(\theta)\|^2} \nabla h(\theta) + \\ &O(\epsilon). \end{aligned} \tag{3.114}$$

□

3.8.4 Averaging

Proposition 3.4 (Averaging in C^1)**.** *Suppose $Q : \mathbb{R}^n \to \mathbb{R}$ is differentiable with locally Lipschitz Jacobian, then under Assumption 3.7 the following averages can be computed as*

$$\frac{1}{\Pi} \int_0^\Pi Q(\theta + S(t)) M(t) dt = \nabla Q(\theta) + O(a), \tag{3.115}$$

$$\frac{1}{\Pi} \int_0^\Pi Q(\theta + S(t)) dt = Q(\theta) + O(a^2). \tag{3.116}$$

Proof. With $S_i(t) = a\sin(\omega_i' t)$ and $M_i(t) = \frac{2}{a}\sin(\omega_i' t)$, we take the period to be Π given as

$$\Pi = 2\pi \times \mathrm{LCM}\left\{\frac{1}{\omega_i'}\right\}, \quad i \in \{1, 2, \ldots, n\}, \tag{3.117}$$

where LCM denotes the least common multiple. For use later, we write the following useful averaging relations

$$\frac{1}{\Pi}\int_0^{\Pi} S_i(t)M_j(t)dt = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \tag{3.118}$$

$$\frac{1}{\Pi}\int_0^{\Pi} S_i(t)S_j(t)dt = \begin{cases} \frac{a^2}{2} & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \tag{3.119}$$

and note that $Q(\theta + S(t))$ can be Taylor expanded [31, Theorem B.6] around θ as

$$Q(\theta + S(t)) = Q(\theta) + S(t)^T DQ(\theta) + O(a^2). \tag{3.120}$$

Now let us expand the LHS of (3.115) as

$$\frac{1}{\Pi}\int_0^{\Pi} Q(\theta + S(t))M(t)dt = \frac{1}{\Pi}\int_0^{\Pi} \Big(Q(\theta)M(t) + S(t)^T DQ(\theta)M(t) + O(a^2)M(t)\Big)dt \tag{3.121}$$

into three terms. The term $Q(\theta)M(t)$ averages to zero because M has zero average. The term $S(t)^T DQ(\theta)M(t)$ averages to $DQ(\theta)$ by (3.118). The last term can be written as $O(a)$ because $M \propto 1/a$. Therefore

$$\frac{1}{\Pi}\int_0^{\Pi} Q(\theta + S(t))M(t)dt = DQ(\theta) + O(a). \tag{3.122}$$

Expanding (3.116) we have

$$\frac{1}{\Pi}\int_0^{\Pi} Q(\theta+S(t))dt = \frac{1}{\Pi}\int_0^{\Pi}\Big(Q(\theta)+S(t)^T DQ(\theta)+ O(a^2)\Big)dt \tag{3.123}$$

$$= Q(\theta)+O(a^2) \tag{3.124}$$

because $S(t)$ averages to 0. □

Acknowledgements

Chapter 3 has been submitted for publication of the material under the title "Safety-Filtered Extremum Seeking with Unknown CBFs" A. Williams, M. Krstic, A. Scheinker. Chapter 3 also contains a partial adaptation of the work contained in the conference paper "Semi-Global Practical Extremum Seeking with Practical Safety" A. Williams, M. Krstic, A. Scheinker, presented at the IEEE Conference on Decision and Control 2023. The dissertation author was the primary investigator and author of these two papers.

Chapter 4

Safe Extremum Seeking for Accelerators

We demonstrate the recent designs of Safe Extremum Seeking (Safe ES) on the 1 kilometer-long charged particle accelerator at the Los Alamos Neutron Science Center (LANSCE). Safe ES is a modification of ES which, in addition to minimizing an analytically unknown cost, also employs a safety filter based on an analytically unknown control barrier function (CBF) safety metric.

Tuning is necessitated by accelerators being large complex systems, with many drifting parameters due to thermal effects and degradation. At the same time, safe operation (the maintenance of state constraints) is crucial, as damage brings astronomical costs, both financially and in operation downtime.

Our measured (but analytically unknown) safety metric is the beam current. We perform multivariable Safe ES on three accelerator applications, in which we adapt 4, 6, and 3 magnet strength parameters, respectively. Two of the three applications are for validated simulation models of beamlines at LANSCE: the first for the Proton Radiography (pRad) beamline of 800 MeV protons for spot size tuning; the second on a high performance code, HPSim, for tuning the low energy beam transport (LEBT) region which contains a beam of 750 keV protons. The third is an experimental tuning of the steering magnets in the LEBT at LANSCE.

4.1 Introduction

4.1.1 Motivation

At Los Alamos Neutron Science Center (LANSCE) the 800-MeV proton linear accelerator requires weeks of tuning every year following an outage during a maintenance period, as well as hundreds of hours of tuning during operation, due to unknown drifting of components along the approximately 1 km beamline. It is often not possible to use a simulation tool or a model to tune the accelerator offline, due to the complexity of the system and because the system changes with time. Therefore, there is need for real-time optimization of parameters such as magnet strengths, radio-frequency (RF) cavity phases, RF cavity amplitudes, steering devices, etc. to correct the beam towards optimal performance. In this chapter we consider beam loss to be our measure of safety, as it arises often in scenarios relating to machine safety. Furthermore, it can easily be measured in the form of beam current. One may also consider other accelerator applications with a different measure of safety (or unknown constraint) like power draw or radiation level.

In this work, we use a novel modification of Extremum Seeking (ES) called Safe Extremum Seeking (Safe ES) to solve the problem of *safe optimization* or *constrained optimization*. The motivation for our use of Safe ES is the combination of complexity and lack of diagnostics in challenging charged particle beam tuning tasks in high energy accelerators. The need for a safe tuning algorithm is clearly demonstrated by the beam power of the LANSCE accelerator, which reaches 800 kW, a factor of 80 greater than a typical welding torch. Such a powerful charged particle beam, if not safely and carefully controlled, can instantly burn a hole in the beam pipe of a particle accelerator, destroying the high vacuum system and irradiating nearby components.

In this work we focus on two very important sections of the LANSCE accelerator, low energy beam transport (LEBT) near the front end of the accelerator and the Proton Radiography facility (pRad) experimental beam line at the very end of the accelerator.

Almost all large particle accelerators have LEBT sections in which the phase space (the collection of positions and velocities of all charged particles) of a relatively low energy beam

is first shaped and refined before it is accelerated to high energies in subsequent acceleration sections. The LEBT sections of accelerators are some of the most difficult to tune and control because low energy beam dynamics are dominated by complex collective effects such as space charge forces, which become much less relevant as the beams are accelerated to highly relativistic energies.

Our in-hardware demonstration of the Safe ES approach takes place in the LEBT section of the LANSCE accelerator. The LEBT is directly after the accelerator's beam source and transforms the 750 keV H^- ion beam from a continuous stream of particles into a roughly 600 μs long beam of individual bunches of ions that are each separated by approximately 5 ns. The bunched beam then enters the first resonant structures of the linear accelerator with the ~5 ns bunch-to-bunch spacing matching the period of the 201.25 MHz resonant electromagnetic fields for subsequent acceleration up to 800 MeV. Tuning of the LEBT region is crucial for LANSCE operations because it sets the initial conditions of the beam that define the rest of the beam transport. Tuning in the LEBT is also challenging because of a lack of diagnostics and because the beam has very low kinetic energy and is very space charge dominated, resulting in a halo of particles around the beam which intercept the beam pipes and accelerator components. We demonstrate experimentally that multi-variable Safe extremum seeking, adapting 3 steering magnet strengths, proved useful in recovering the safe operating condition of the LEBT region.

We also demonstrate the Safe ES approach via a simulation study of the pRAD experimental facility at LANSCE which is used to characterize the behavior of high explosives, materials under high strain rate, and probe the 3-dimensional density of objects [46]. Under normal operation of pRAD, beam operators view CCD cameras which are pointed at phosphor screens inside the beam pipe. The screens, located at various points along the beamline, illuminate when impacted by the beam, showing its spot size. Daily, several magnet strengths must be manually tuned such that the beam is appropriately sized at one or more of the screens. We show how Safe ES can shape the beam quickly and automatically, while remaining safe conditions as defined by low beam loss.

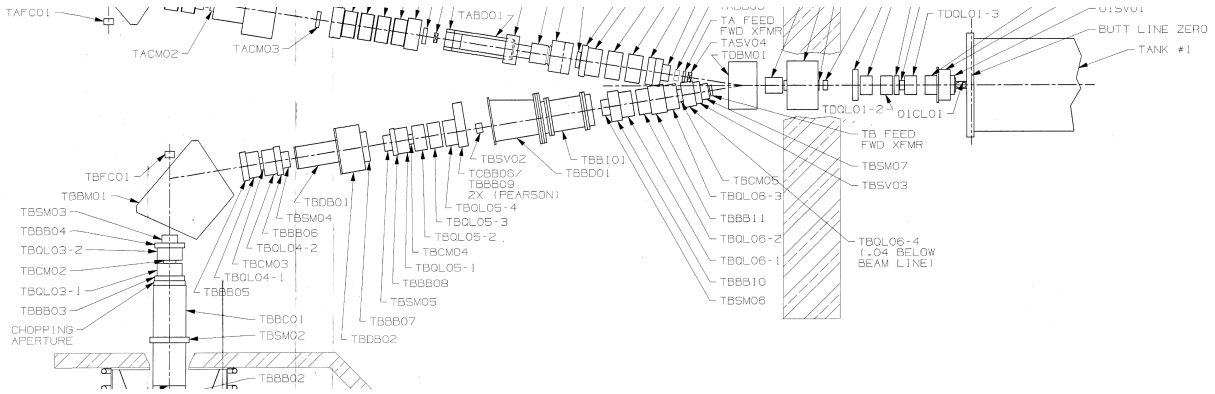

Figure 4.1. A diagram of Line B of the LEBT containing many components including current monitors, bending magnets, steering magnets, focusing magnets, bunchers, and more. Line B delivers H^- beam to the first tank of the drift tube linac. Line A, shown partially in the top quarter of the image, delivers H^+. In Sections 4.3.2 and 4.3.3 we tune several components in this section of the accelerator.

Safe ES, like classical ES, is 1) real-time 2) easily implementable 3) computationally simple. Yet, it can handle an unknown constraint which may be approximately maintained over the course of tuning. It minimizes an unknown objective function $J(\theta)$, over the parameters θ, while also ensuring that some measure of safety $h(\theta)$ is kept positive. Therefore, our algorithm approximately solves the following problem

$$\min_{\theta} J(\theta) \text{ subject to } h(\theta) \geq 0, \tag{4.1}$$

where h and J are directly measured signals - or are constructed from directly measured signals, but are analytically unknown. The algorithm also guarantees, approximately, invariance of the safe set and escape from unsafe initial conditions. We say that we solve the above problem "approximately" because ES (and Safe ES) fundamentally relies on approximating of the gradient of unknown functions. Violations of the safety of the system may be made arbitrarily small (or zero in some cases) by appropriately choosing design constants. The design of h is then especially useful in this regard as one may design it such that it provides a 'soft' limit on the safe signal of interest. For example, if some safe signal $s(t) \in [0,1]$ must never go below 0.5, then one may conservatively design $h = s(t) - 0.6$ if one is sure that the chosen small perturbation

signal and transient of the ES scheme will not cause a deep excursion of θ (a change in h more than 0.1) into unsafe territory. For preventing catastrophic failure to a system in practice, one would implement a higher-level controller to shut down the system once h is sufficiently negative. At LANSCE, high level systems like this will kill all beam if radiation levels spike or beam current is lost to a large degree.

We require the optimization to be done such that a trajectory remains practically safe for all time, or during the entire length of the tuning episode. Methods such as particle swam optimization and genetic algorithms cannot be applied to such scenarios unless a modification is made to ensure that the violation of the trajectory into an unsafe region does not occur.

Note, that we can indeed describe particle accelerators as two unknown, time varying static maps - $J(\theta,t)$ and $h(\theta,t)$. We imagine that for given parameters θ at time t, we instantly receive a reading of both J and h. This is true in many cases when tuning accelerators, as we are not able to do any feedback while the beam is traveling near the speed of light through the beampipe. We assume that only after the beam has completed its journey, will we have measurements of J and h, after which time we can perform the update to the parameters θ. Additionally, oscillations introduced by the ES scheme are therefore considerably slower compared to the dynamics of the plant/accelerator and do not interfere with its operation. The oscillations that we are introducing are changing the voltage or current set points of power supplies which result in changes in the magnetic fields of electromagnets, they are not mechanical oscillations and therefore do not wear out any components.

4.1.2 Accelerator Tuning and Optimization Literature

The Nelder-Mead simplex method [39] and ML based methods have been used to tune accelerators [5]. Extremum seeking has been studied in simulation for accelerator applications [89], and was first used experimentally in an accelerator tuning problem to minimize an objective function [87]. Since then, a bounded form of ES with guaranteed limits on parameter update rates [88], has been used extensively in several accelerator applications for free electron laser energy

maximization [78], electron beam trajectory control [84], real-time multi-objective optimization [82], and for beam loss minimization [83].

Recently, various machine learning (ML)-based methods have been developed for control and optimization of particle accelerator beams [12]. Bayesian optimization has become a popular tool in tuning and in some cases has been used to design safety aware tuning algorithms [48, 47, 26]. Bayesian optimization (and methods based on it), unlike ES, constructs a probabilistic estimate of the unknown functions, in the form of a Gaussian Process (GP), and determines a new point to sample based the fitted function. Neural networks (NN) have been used as surrogate models for magnet control [40] and for simulation-based optimization studies [29]. Neural networks are also being used for uncertainty aware anomaly detection to predict errant beam pulses [17], as virtual diagnostics for 4D tomographic phase space reconstructions [102], for predicting the transverse emittance of space charge dominated beams [58], and for high resolution longitudinal phase space virtual diagnostics [105]. Neural network-based deep reinforcement learning (RL) methods have been used for accelerator control [35], and in a sample efficient manner, which trains a policy based on data at two beam lines at CERN [41].

Although many ML-based tools have been developed they all suffer major limitations when it comes to time-varying systems. If a system changes then NN, GP, and RL methods all require new data for re-training of their models in order to be applicable for accelerator control. This major limitation is overcome by adaptive model-independent methods such as ES. Adaptive machine learning frameworks combining ES with neural networks have been demonstrated to extend the use of ML for time-varying systems. ES-based adaptive ML has been demonstrated for automatically shaping the longitudinal phase space of short intense electron beams in the LCLS FEL [81] and for creating virtual 6D diagnostics of time-varying charged particle beams [76, 80].

In all the methods described above, including those using ES, the approach to safety has been some expert-based combination of setting hard bounds on allowed parameter values, adding additional terms to the cost function, and algorithm hyper-parameter tuning. The Safe ES

method we demonstrate in this chapter reduces the amount of required hyper-parameter tuning and removes the manual design of a tradeoff between safety and optimization which depends on the weights given to safety-related cost function terms relative to objective-related cost function terms. For example, in the real-time multi-objective ES optimization application in [82], while the objective was beam spot size minimization, the safety-related term in the cost function was the beam's distance from a desired reference trajectory, and there was a tradeoff between the two depending on the weights. In the approach here, practical safety is always enforced.

4.2 Algorithm

4.2.1 Safe Extremum Seeking

The goal of our Safe ES algorithm is to approximately solve the problem in (4.1) where $\theta \in \mathbb{R}^n$ and $J, h : \mathbb{R}^n \to \mathbb{R}$. We will first recap the (semiglobal) Safe ES algorithm presented in Chapter 3. Given the static maps $J : \mathbb{R}^n \to \mathbb{R}$ and $h : \mathbb{R}^n \to \mathbb{R}$ we define the algorithm dynamics as:

$$\dot{\hat{\theta}} = k\omega_f(-G_J + \min\{||G_h||^{-2}, M^+\}\max\{G_J^T G_h - c\eta_h, 0\}G_h), \tag{4.2}$$

$$\dot{G}_J = -\omega_f(G_J - (J(\hat{\theta}(t) + S(t)) - \eta_J)M(t)), \tag{4.3}$$

$$\dot{\eta}_J = -\omega_f(\eta_J - J(\hat{\theta}(t) + S(t))), \tag{4.4}$$

$$\dot{G}_h = -\omega_f(G_h - (h(\hat{\theta}(t) + S(t)) - \eta_h)M(t)), \tag{4.5}$$

$$\dot{\eta}_h = -\omega_f(\eta_h - h(\hat{\theta}(t) + S(t))) \tag{4.6}$$

The maps is evaluated at θ, defined by

$$\theta(t) = \hat{\theta}(t) + S(t)\ . \tag{4.7}$$

The perturbation signal S and demodulation signal M are given by

$$S(t) = a\left[\sin(\omega_1 t),\ \ldots,\sin(\omega_n t)\right]^T, \tag{4.8}$$

$$M(t) = \frac{2}{a}\left[\sin(\omega_1 t),\ \ldots,\sin(\omega_n t)\right]^T. \tag{4.9}$$

4.2.2 Intuition Behind the Design

The tuning parameter vector, θ, is imparted with gradient descent dynamics acting on the objective function *and* gradient ascent dynamics acting on the safety function:

$$\dot{\theta} = k\omega_f\Big(\underbrace{-G_J}_{\text{Gradient } \textit{Descent} \text{ of the Objective}} + \overbrace{AG_h}^{\text{Gradient } \textit{Ascent} \text{ of the Safety}}\Big) + \underbrace{\dot{S}(t)}_{\text{Exploration Signal}}. \tag{4.10}$$

The quantity A is a non-negative, state-dependent, scalar function defined as

$$A := \min\{||G_h||^{-2}, M^+\}\max\{G_J^T G_h - c\eta_h, 0\} \geq 0, \tag{4.11}$$

$$\approx \frac{\max\{\nabla J(\theta)^T \nabla h(\theta) - ch(\theta), 0\}}{\nabla h(\theta)^2}. \tag{4.12}$$

It turns "on/off" to determine whether to consider safety and how much to consider it. Note that η_J, η_h, G_J, G_h are estimator states which are meant to converge close to the true quantities $J, h, \nabla J, \nabla h$.

To understand the design of A, consider the case when the true quantities are known and there is no need for the exploration/perturbation signal. So,

$$\dot{\theta}(t) = u, \tag{4.13}$$

for θ and u in n-dimensions. If we consider the gradient descent term $u(\theta) = -\nabla J(\theta)$ to be the nominal controller, we can employ the formulation given by a QP [9, 8] to find an additive safety term which, when combined with the nominal, yields an provably safe controller (i.e. the safe set

$\{h(\theta) \geq 0\}$ is said to be positively invariant). The safety aware control law is $u_s = -\nabla J(\theta) + \bar{u}$, and $\bar{u}$ is given by the QP

$$\bar{u} = \arg\min_{v \in \mathbb{R}^n} ||v||^2 \quad \text{subject to} \tag{4.14}$$

$$ch(\theta) + \nabla h(\theta)^T(-\nabla J(\theta) + v) \geq 0, \tag{4.15}$$

for some positive constant c. It was showed that there is an explicit solution for $\bar{u}$:

$$\bar{u}(\theta) = \frac{\nabla h(\theta)}{||\nabla h(\theta)||^2} \max\{\nabla h(\theta)^T \nabla J(\theta) - ch(\theta), 0\}. \tag{4.16}$$

It can now be seen that $\bar{u}(\theta) \approx A\nabla h(\theta) \approx AG_h$.

Safe ES has a number of desirable properties, the most notable of which is practical safety. Another desirable property is that the control law is smooth, and trajectories avoid crashing into an unsafe region without first trying to skirt around them while they are some distance away. Additionally the formulation of the QP in (4.14) - (4.15) guarantees that the modification to the nominal control always yields the final control law smaller than the nominal ($u_s \leq u$), for $h > 0$. This is a practical benefit for the control designer who must choose the design parameters such that the adaptation dynamics, given in the right-hand side of (4.2), must be made sufficient small for stability purposes. Yet another benefit is that when the trajectory starts in the unsafe region ($h < 0$), the value of h is guaranteed to monotonically increase in time (provided certain assumptions like the gradient of h does not vanish), escaping to more safe parts of the state space.

4.2.3 Implementation and Design Parameter Choices

To implement the extremum seeking algorithm we integrate the differential equations (4.2) - (4.6) numerically. Expressing the equations using $x^T = [\theta^T, G_J^T, \eta_J, G_h^T, \eta_h]$, we have $\dot{x} = f(x)$. With some initial condition $x(0) = x_0$, we compute $x_n = x_{n-1} + f(x)dt$ and the parameters θ can be set with new values iteratively. The value dt is chosen such that period of oscillation

$T = 2\pi/\omega_i \gg dt$. For all implementations of Safe ES shown in this chapter, we specify $M^+ = 10^4$.

Note that unlike the version of ES in [86], we rely on the explicit estimation of gradients. Therefore, in all the results and plots given we "warm up" the algorithm first by settings $k = 0$ for the first 1-3 periods of the perturbations $S(t)$, before turning on the algorithm and simulating with a nonzero value of k. This allows the estimated quantities to converge more closely to their true values before the adaptation of the parameter of θ begins. This technique can be thought of as a trick to more accurately initialize the estimator states G_J, η_J, G_h, and η_h.

We choose a value of a that causes the oscillatory response in the measurements of $J(\theta(t)), h(\theta(t))$ corresponding to the frequencies ω_i. We set ω_f roughly the same order of magnitude as ω_i - a smaller ω_f can be chosen if one wishes to add more smoothing to the estimator states. Finally, we gradually increase k from zero until sufficient performance is observed $J(\theta(t)), h(\theta(t))$. Choosing k too large may cause instability.

4.3 Accelerator Applications

4.3.1 Simulated pRad Tuning

For experimental studies on the tuning of Line C, which is the section of the accelerator responsible for delivering beam to pRad, we use the particle beam dynamics code TRANSPORT, which has been validated with measurement data and is shown to give accurate predictions of the LANSCE beam profile [75]. The code contains relevant quadrupole magnet strengths - which in practice are manually tuned by hand to achieve a spot size required for an optimal delivery of beam to the experiment.

The TRANSPORT model is a beam envelope model, that models the bulk behavior of the beam which is represented as an ellipse in 6 dimensions (3 spatial lengths and 3 spatial velocities). Even though it can model space charge effects, it cannot simulate particle loss. Therefore, we introduce a model of loss which is applied after TRANSPORT computes the beam dynamics solutions. This model is not based on data but gives us sufficient complexity to demonstrate our

algorithm. We model the percentage of beam remaining b [%] as a function of the integrated beam size of the TRANSPORT solution:

$$b = 1 - k_1 \int_0^{s_{end}} (X(s) + Y(s))ds - k_2 \int_0^{s_{end}} [((X(s) - r_p)^+)^2 + ((Y(s) - r_p)^+)^2]ds, \quad (4.17)$$

where $X^+ := \max\{X, 0\}$ and r_p is the radius of the beampipe. We choose $k_1 << k_2$ and the equation describes 1) a small linear loss on the size of the beam that always applies at any point s 2) a large quadratic loss applying only when the ellipse contacts the beam pipe. The small linear loss models the persistent loss which always occurs due to beam halo, and the large quadratic loss describes the sharp and sudden loss which only occurs when the central mass of the bunch comes close to the beam pipe. This model provides sufficient complexity which captures the dynamics of loss realistically, although it has not been thoroughly validated with data, as the TRANSPORT code itself has, which only describes the dynamics of the beam ellipse. We also provide the LANSCE magnet names used in this study: "AQM1","AQM2", "XQM3", "XQM4".

The strength of these four magnets correspond to the components of $\theta \in \mathbb{R}^4$. Given a the parameter θ, the function $b = b(\theta(t))$ because only changes in magnet strengths give rise to changes in the simulation.

The goal for this simulation study relates exactly to a tuning problem which must usually be performed in the control room. We desire to track a specific spot size of beam as a point in the beamline, directly preceding the pRad experimental dome. Operators usually are tasked with this job, but we will show Safe ES is capable of performing the task safely. We will track the size at the end of the simulation σ_x, σ_y [m] and define the objective as

$$J(\theta(t)) = (\sigma_x(\theta(t)) - 0.025)^2 + (\sigma_y(\theta(t)) - 0.025)^2, \quad (4.18)$$

and so the desired spot size is a 2.5 cm circle. We use values of $k_1 = 0.02$, $k_2 = 100$. We make

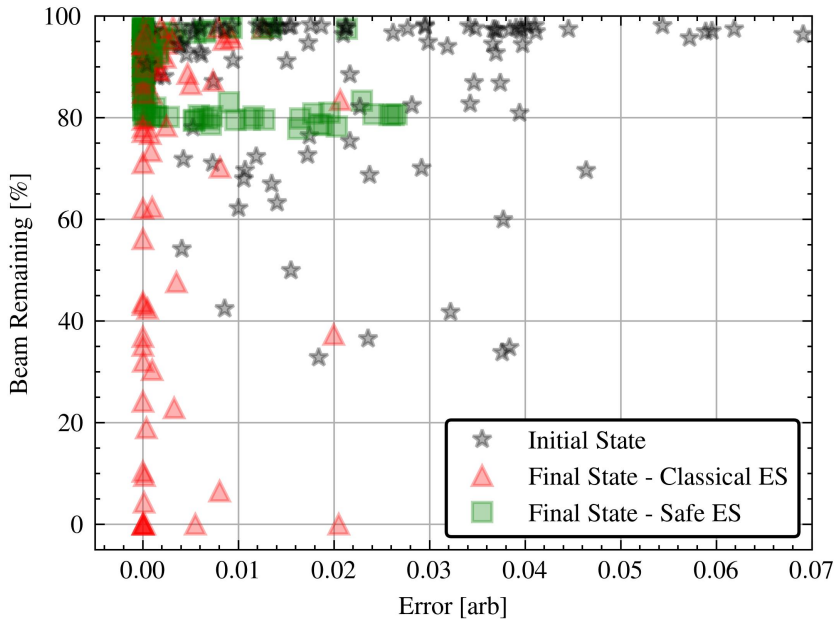

Figure 4.2. For 100 episodes, we plot both the initial and final error and beam remaining values for safe and unsafe extremum seeking.

the simple choice of

$$h((\theta(t)) = b(\theta(t)) - 0.80, \tag{4.19}$$

with no scaling to either h and J, which was done earlier in the HPSim study to achieve the same order of magnitude of J and h.

To understand the design of h some additional information about the accelerator system in question is needed. Due to the underlying engineering physics of accelerators (of kind at LANSCE), the task of guaranteeing that near 100% of the beam survives is impossible. From experience, accelerator operators understand that if roughly 80% of the beam survives, then this is as close to 'safe' as one can hope for. This limitation is due to a combination of beam loss in the form of beam halo intercepting collimators and the limited acceptance of the resonant accelerating structures. Therefore, in our case we place the boundary of 'marginal' safety at 80% of beam remaining.

We randomly initialize the magnet settings at up to 20% of their default settings, keeping points which have at least 20% beam remaining. This is so that we do not assume we start

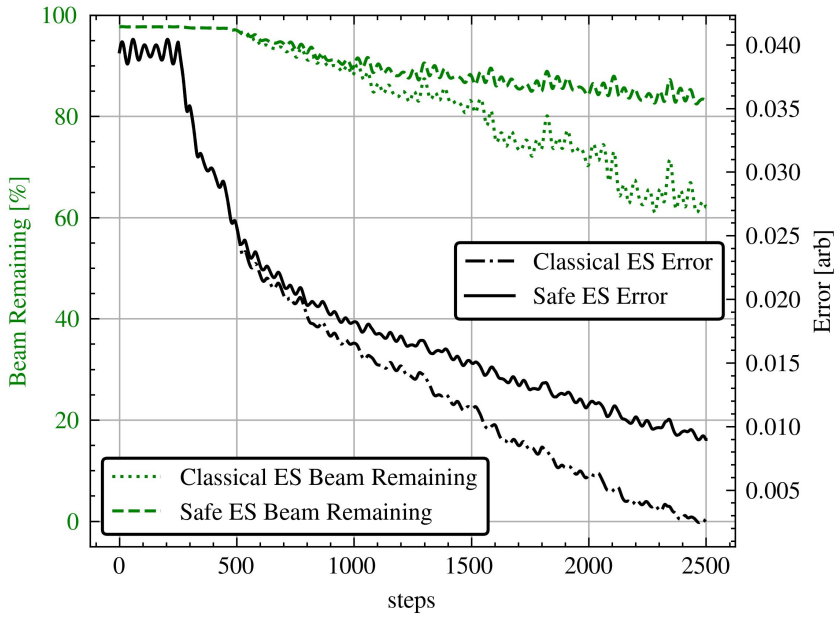

Figure 4.3. Percentage of beam remaining and error for safe and unsafe extremum seeking, for one given trajectory.

the simulation in an unrealistic initial condition (catastrophic beam loss) which may crash the simulation. Then we run both Classical ES and Safe ES (with coefficients above) for 2,500 steps and compute the objective error and beam remaining values at the end of the episodes. We do this for 100 different initial conditions and show the initial condition and final condition of both the Safe ES and Classical ES trajectories in Fig. 4.2. The Safe ES parameters used are $dt = 0.012$, $a = 0.005$, $k = 0.04$, $\omega_f = 10$, $\omega_1 = 5$, $\omega_2 = 7$, $\omega_3 = 11$, $\omega_4 = 13$ and $c = 1$.

To construct the Classical ES algorithm, we simply use the dynamics of (4.2) - (4.4) without the safety term in the dynamics of (4.2):

$$\dot{\hat{\theta}} = -k\omega_f G_J, \tag{4.20}$$

$$\dot{G}_J = -\omega_f (G_J - (J(\hat{\theta}(t) + S(t)) - \eta_J) M(t)), \tag{4.21}$$

$$\dot{\eta}_J = -\omega_f (\eta_J - J(\hat{\theta}(t) + S(t))). \tag{4.22}$$

We choose the parameters of the Classical ES scheme to be the same as those used in the Safe ES

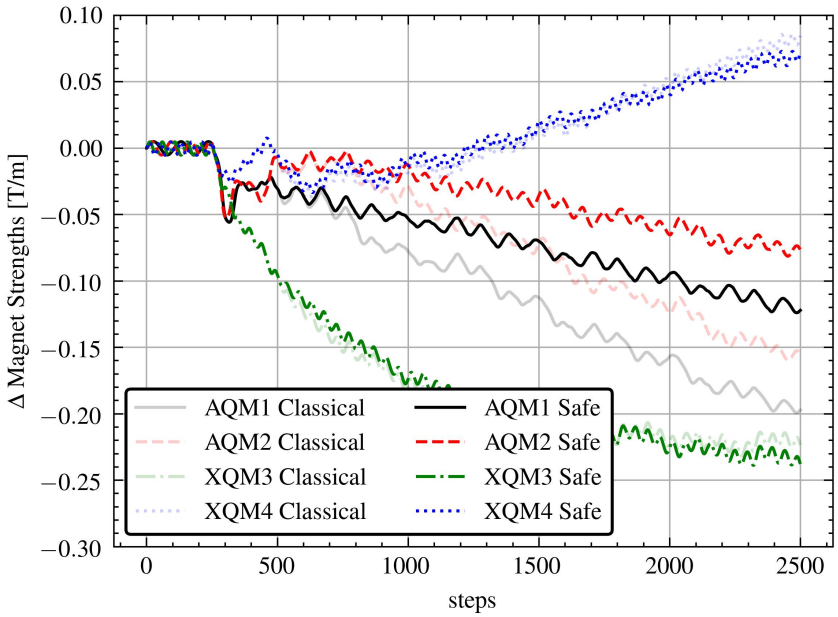

Figure 4.4. The change in magnet strength trajectories, $\hat{\theta} - \hat{\theta}_0$, of Safe ES compared with Classical ES.

scheme - apart from the irrelevant value c which does not appear in (4.20) - (4.22).

Fig. 4.2 depicts the effectiveness of the algorithm with respect to various initial conditions. We show for a 100 randomly chosen initial conditions, which have small and large values of both performance and safety, all trajectories of Safe ES eventually achieve approximately 20% or better beam remaining - albeit with various levels of performance. Classical ES shows better performance in general (more red arrows are grouped on the left-hand side of the plot) but various levels of safety.

In Fig. 4.3 and 4.4 we choose one trajectory of the 100 episodes and plot the magnet trajectories as well as the beam remaining signal and error signal along the trajectory. As expected, the Classical ES trajectory in Fig. 4.3 finds itself traveling into an unsafe region in favor of better performance, while the Safe ES trajectory still achieves good performance but does so while keeping the system safe. Running the simulation for longer than 2,500 steps may also indeed show that Safe ES in this case may ultimately achieve near perfect performance by taking a longer path through parameter space.

4.3.2 Simulated LEBT Tuning

HPSim [69] is a GPU accelerator particle tracking code that simulates all of the RF and magnetic devices used at LANSCE and accurately describes the physics of beams, including challenging collective effects like "space charge". It has been used to demonstrate tuning algorithms [87] and is currently being developed as an online digital twin of the LANSCE linear accelerator (linac) [38].

The LEBT section delivers the 750 keV beam from the source to the drift tube linac (DTL) accelerating structures of the linac. In the example presented here, we apply Safe ES to automatically tune 6 of the quadrupole magnets in LEBT which are responsible for confining the beam to the pipe as it is transported. We have specified that we want the smallest possible beam entering the DTL, yet we must attain approximately greater than 80% of the beam while doing so. Although this is not precisely the goal of the tuning the LEBT beamline, it gives us with a realistic scenario showing the benefits of Safe ES when we are not sure if gradient descent of the objective function leads the parameters in a safe direction or not.

In this simulation study we compare Safe ES with Classical ES with a cost function modified with a small weight on the safety of the system, as this is often a trick used to incorporated safety considerations in the Classical ES scheme. We will also demonstrate the qualitative behavior of the c parameter in the Safe ES design and show that Safe ES in general leads to a configuration of the accelerator which leads to less contact with the beam and sensitive structures in the beampipe.

Given the goals and safety concerns we have described, we choose the following cost function to minimize,

$$J(\theta(t)) = w_J(\sigma_x(\theta(t))^2 + \sigma_y(\theta(t))^2), \tag{4.23}$$

and we also choose the function h as

$$h(\theta(t)) = w_h(b(\theta(t)) - 0.80), \tag{4.24}$$

where b, σ_x, σ_y are signals measured in time and w_J, w_h are chosen such that J, h are approximately on the order of 10. Using HPSim we determine 6 quadrupoles along the beamline to use for tuning which are roughly equally spaced in s - the variable used to describe the longitudinal positive along the beamline. We provide the LANSCE magnet channel names used in this study: "TBQL005V01", "TBQL005V02", "TBQL005V03", "TBQL005V04", "TBQL006V03", "TBQL006V04" with corresponding magnet strengths $\theta \in \mathbb{R}^6$. The physical locations of the components can be shown in Fig. 4.1 as "TBQL05-1", "TBQL05-2", ..., "TBQL06-4".

It is a common strategy to handle the safety problem in Classical ES with a modified cost:

$$J(\theta(t)) = w_J(\sigma_x(\theta(t))^2 + \sigma_y(\theta(t))^2) - wb(t) \tag{4.25}$$

where w is chosen sufficiently large to (with some luck) yield a safe trajectory. We compare our algorithm with this strategy and show that in this case it is not possible to know exactly how to specify w without achieving some loss in performance or loss of safety. While for any choice of our hyper-parameter c (which is the most analogous hyper parameter to w and in some sense governs the safety-versus-performance tradeoff), we achieve approximate safety in all cases. This is because c does not change whether or not we maintain approximate safety, but simply how fast the trajectory is allowed to approach the unsafe region.

The Safe ES parameters used are $dt = 0.005$, $a = 0.1$, $k = 0.05$, $\omega_f = 10$, $\omega_1 = 5$, $\omega_2 = 7$, $\omega_3 = 11$, $\omega_4 = 13$, $\omega_5 = 17$, $\omega_6 = 19$. We choose $w_J = 50$ and $w_h = 100$ so that the functions h and J are roughly of the same order.

The Classical ES algorithm uses the same parameters (apart from c) and is implemented in the same way as in the pRad tuning example using (4.20) - (4.22).

In Fig. 4.5 we show that for various values of c, approximate safety is enforced. The smaller values of c dictate a slower rate of approach of the trajectory towards the barrier. Fig. 4.6 shows that Classical ES with a modified cost function, in (4.25), does not always remain safe and the weight w cannot be known ahead of time to guarantee best performance and safety. It turns

out that choosing $w \approx 0.07$ in the Classical ES scheme would deliver performance and safety comparable to that of the Safe ES controllers, but this choice would require excessive effort, or extra system knowledge, on the part of the control designer.

In Fig. 4.7 we compare the beam 'envelope' which was found with Safe ES versus Classical ES. The coordinate s runs longitudinally along the beam path, and Fig. 4.7 plots the transverse size of the beam, which is the actual height (σ_y) and width (σ_x) of the beam if one were to look down the beam pipe. The algorithm has no knowledge of this information, only the measurement of J and h. The figure also provides information about at which point beam loss is occurring as the sharp change in the size of the beam marks where the start of the first DTL modules lies. The Classical ES trajectory used was for that of $w = 0$ which yielded approximately a final loss value of about 30%. It was compared to the envelope found using Safe ES with $c = 1$. Fig. 4.7 demonstrates that the transverse size of the beam along the beam pipe was smaller in the case of Safe ES. This is physically what we expect given that large beam losses are expected to occur when the beam contacts the beampipe.

4.3.3 Experimental LEBT Tuning

For the in-hardware LEBT demonstration, the ES objective function, J, was a measure of beam remaining I_b (negatively proportional to beam loss), while the safety constraint, h, was a measure of how much beam is being lost at a collimator in the LEBT, I_c. The amount of beam that hits the collimator is proportional to the transverse beam size in this section of the accelerator and it is crucial to keep that beam size small enough to reduce how much beam is intercepted. At the same time, pinching the beam down to too small of a size at the collimator creates very large divergence of the beam, as particle become highly repellent to each other, and would cause downstream beam loss. The Safe ES method was set up to dynamically balance the trade-off between beam size at the collimator and overall beam loss.

The safety of the system is a function of the measured current I_c at the collimator and the performance of the system is a function of a measured current I_b farther downstream. The

signals I_c and I_b are the raw readings from sensors. Both signals in their raw form are negative as the H^- ion beam has a negative current. If the magnitude of I_c becomes too large, or too negative, then a larger fraction of beam is contacting the collimator leading to lack of safety. On the other hand, if I_b becomes large (more negative) then we have a large fraction of beam surviving at the end of the line - therefore we desire that I_b be minimized as much as possible. Note that in this tuning problem, it is not possible to tune I_c identically to zero. I_c should merely be constrained approximately within some small range, which is why Safe ES is an ideal choice of algorithm in this scenario.

The safety goal was to keep the measured current intercepted at the LEBT collimator, I_c, at a value greater than approximately -0.5 mA. The reading of $I_c(t)$ lies in the range $[-2.5, 0]$, and we define the safety function defined as

$$h(\theta(t)) = w_h(0.5 + I_c(\theta(t)), \tag{4.26}$$

with $w_h = 10$ so that $h \geq 0$ corresponds to the region of safety where $-0.5 \leq I_c \leq 0$. We implement a 10 point average of the raw data (shown in Fig. 4.10) to smooth the measurements of J and h - see the smoothed measurements in Fig. 4.8. The beam loss was minimized by simply defining $J(\theta(t)) = I_b(\theta(t))$. Minimizing loss in this case is the same as minimizing the beam current because I_b is always negative. The tuning parameters used are the three LEBT steering magnet strengths, given by the channels "TBSM402P01", "TBSM502P01", and "TBSM602P01" with the corresponding strengths $\theta \in \mathbb{R}^3$. The physical locations of the components can be shown in Fig. 4.1 as "TBSM04", "TBSM05", and "TBSM06". I_b and I_c have channel labels "TDCM001102" and "TDBA001101" respectively. The Safe ES parameters used are $dt = 0.359$, $a = 0.1$, $k = 0.03$, $\omega_f = 1$, $\omega_1 = 1$, $\omega_2 = 1.375$, $\omega_3 = 1.75$ and $c = 1$.

Results of tuning in the first 100 steps are are shown in Fig. 4.8 and 4.9. In Fig. 4.10 we show that initially the steering magnets are manually tuned to achieve a desirable operating condition. We then detune the magnets and run the algorithm around minute 4. The algorithm

demonstrates it can recover the original operating condition, minimizing I_b to -1.8 mA and driving I_c to near -0.5 mA. The data from Fig. 4.8 - 4.9 is generated using a 10-point average of the data during the times shown between the dotted lines in Fig. 4.10.

This is an example of multi-variable Safe ES where the parameters of the algorithm converge to the barrier of h, achieving practical safety, while also minimizing the objective. Additionally, we have demonstrated that this algorithm can save an accelerator system operating unsafely and drive the system towards a safe operating condition.

4.4 Conclusion

In this chapter we demonstrate that Safe ES has several uses in tuning various accelerator systems and subsystems. We have shown that in two validated simulations of beamlines at LANSCE, we can use Safe ES to perform tuning. Additionally, we present an in-hardware demonstration of tuning the LEBT section of the accelerator. The advantage of this method is that we require no knowledge or prior gathering of data of the underlying accelerator system and can guarantee practically safe operation. Furthermore, compared to other methods commonly used in particle accelerator systems, Safe ES is very simple to implement.

Acknowledgements

Chapter 4 is a reprint of the material as it appears in IEEE Transactions on Control System Technology 2024, under the title “Experimental Safe Extremum Seeking for Accelerators” A. Williams, A. Scheinker, E. Huang, C. Taylor, M. Krstic. The dissertation author was the primary investigator and author of this paper.

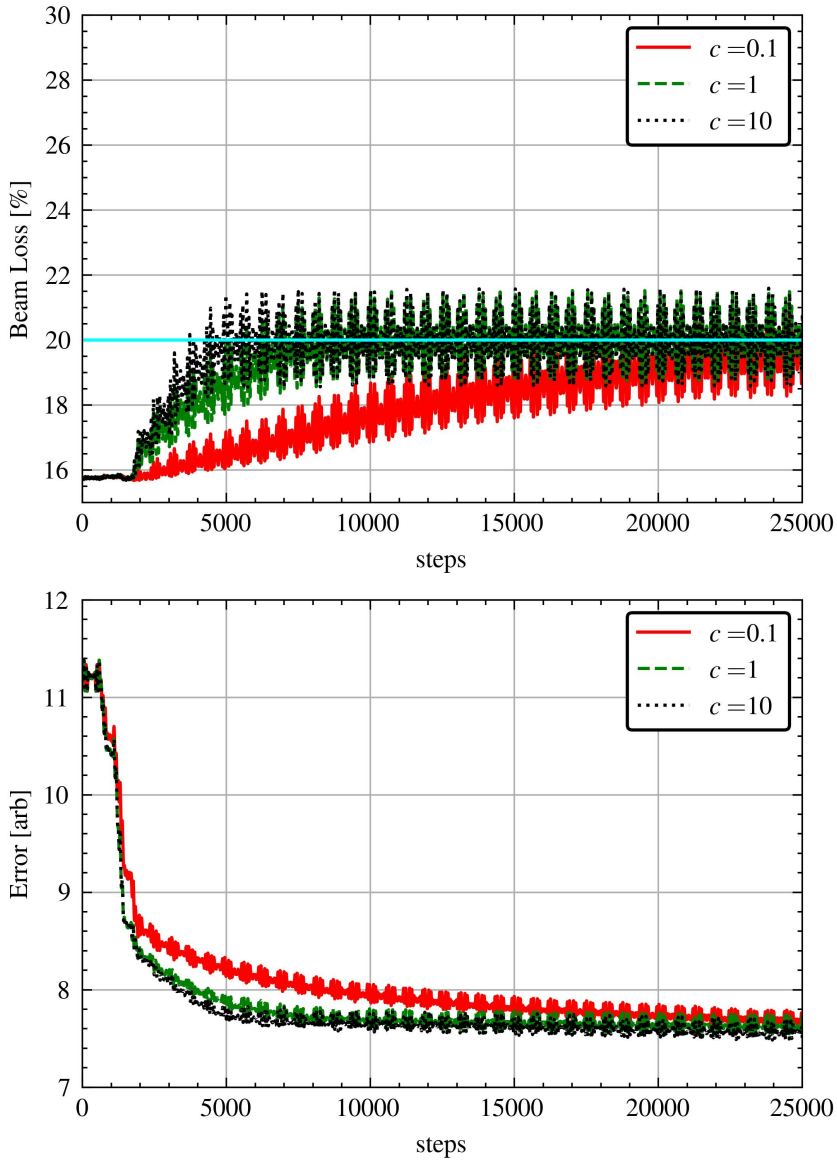

Figure 4.5. Percentage of beam loss and error for Safe ES with different values of c.

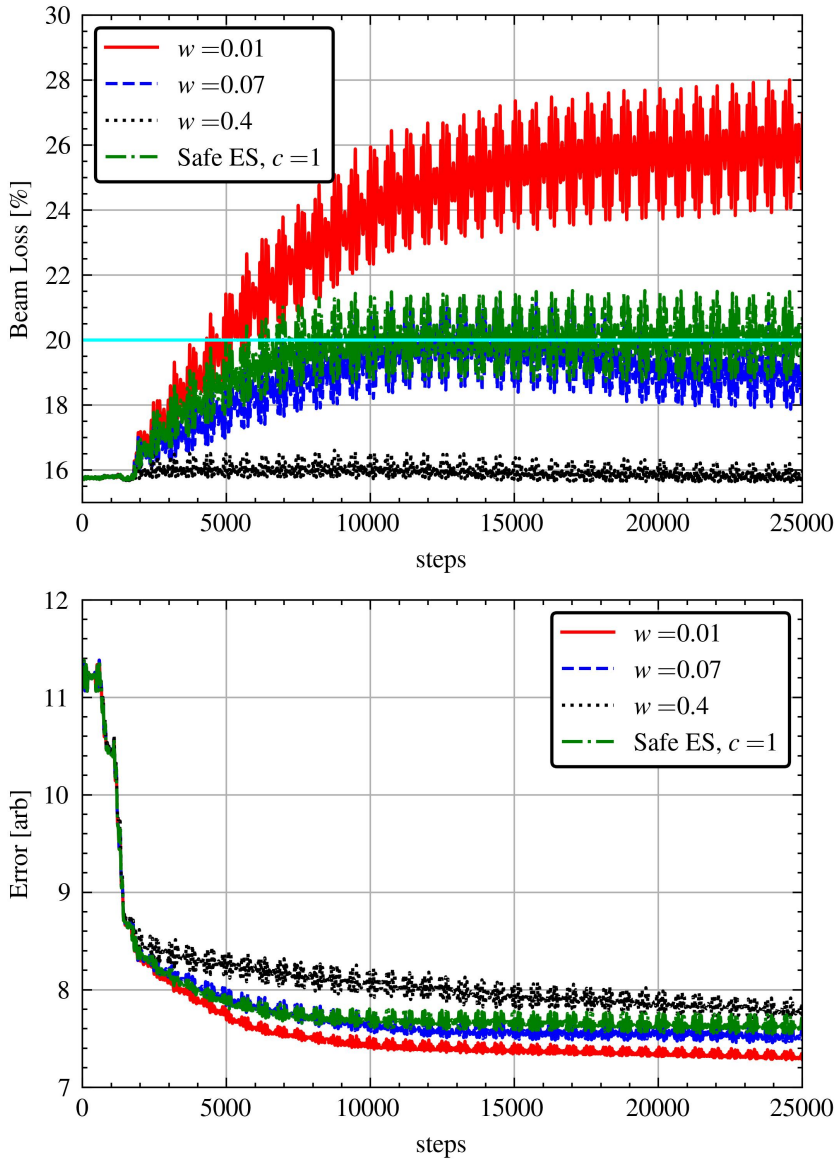

Figure 4.6. Percentage of beam loss and error for Classical ES with various weights w, plotted against one trajectory of Safe ES. Note that the Error values shown in the bottom plot is $w_J(\sigma_x(t)^2+\sigma_y(t)^2)$, for each of the 4 trajectories given.

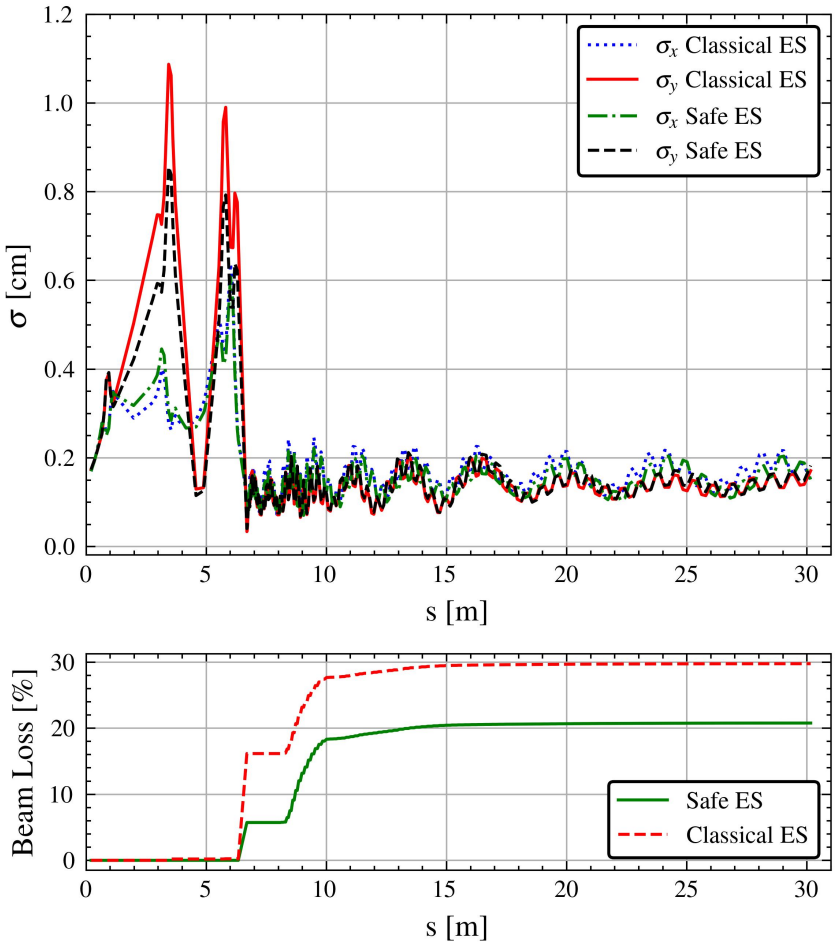

Figure 4.7. The beam envelope along s, depicting the transverse size of the beam throughout the LEBT region and through the first 4 DTL modules in Classical ES with $w = 0$ and Safe ES with $c = 1$.

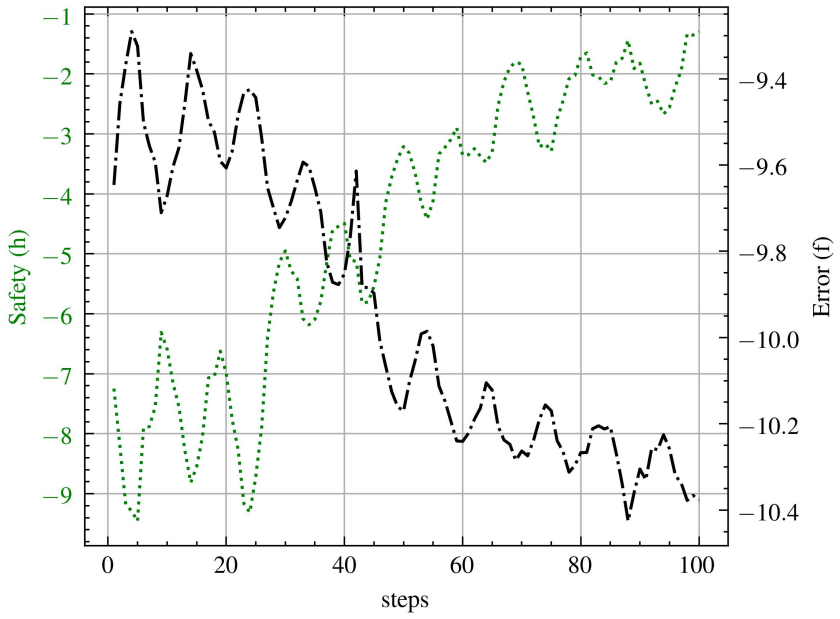

Figure 4.8. Barrier function and objective function over the first 100 steps using a 10 point average of the noisy measurements of I_b and I_c.

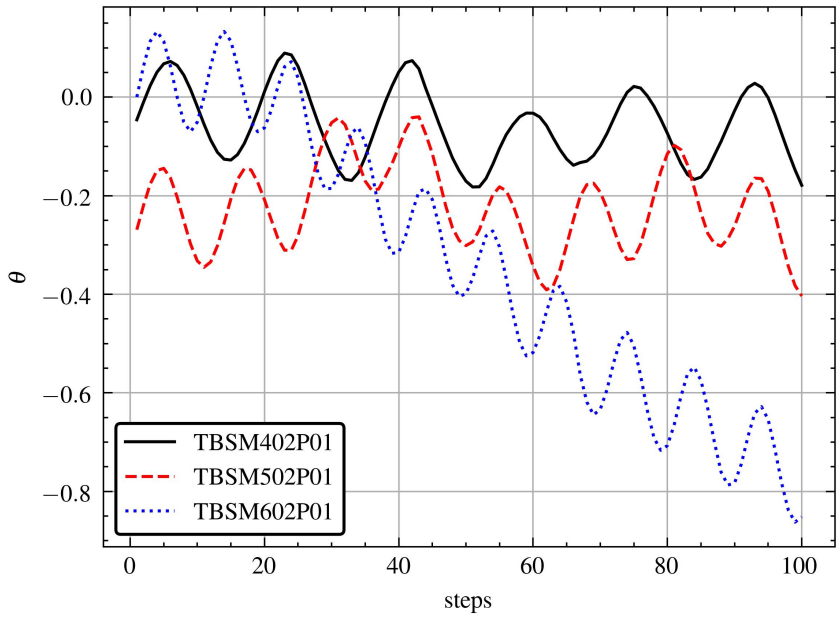

Figure 4.9. Parameter trajectories over the first 100 steps using a 10 point average of the noisy measurements of I_b and I_c.

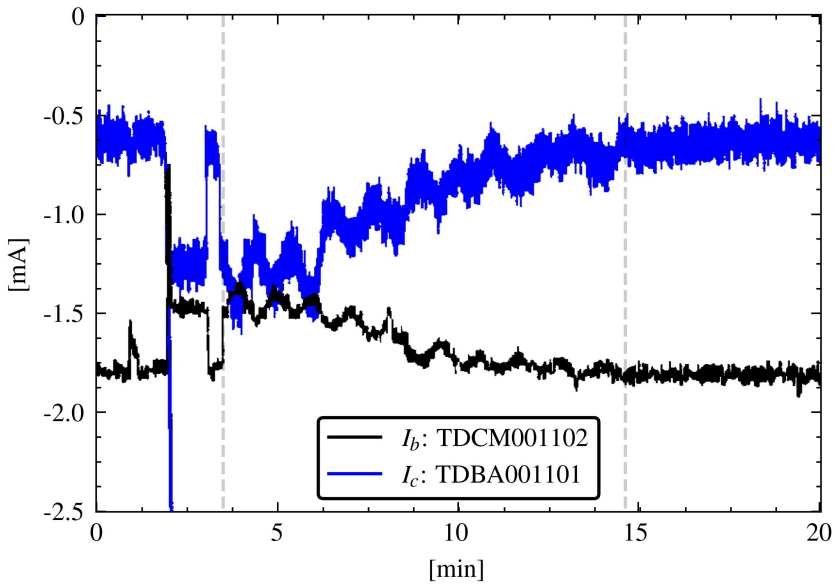

Figure 4.10. Raw plot of I_b and I_c taken from the LANSCE control room monitor for process variables. The first dotted vertical line marks the start of the algorithm, the second dotted line marks where recordings in Fig. 4.8 and Fig. 4.9 finish.

Chapter 5

Conclusions

5.1 Summary

In Chapter 1, we set the stage for this dissertation with an exploration of particle accelerators. The operation of accelerators is a challenging process, requiring continuous, precise control over the particle beam to prevent severe consequences, such as damage to the accelerator or safety hazards from unwanted radiation generation. Large accelerators, extending over miles with thousands of components, are subject to time varying drifts from aging, thermal effects, and environmental changes, necessitating constant adjustments to maintain beam quality. Often, operators must manually adjust control settings to fine-tune performance indicators such as beam intensity and energy distribution, all while ensuring the safety of both the accelerator itself and human operators.

To address these concerns, we introduce an algorithm known as "Assignably Safe ES" (ASfES) in Chapter 2. This controller, the first safe ES controller that we explore, is local in nature but includes several key innovations. One notable aspect is the high degree of control it offers over the safety behavior, allowing the user to assign the "attractivity rate" of the algorithm – a time constant that governs system's approach to and from the safe set, as measured by the safety metric h. Our analysis leverages classical averaging and singular perturbation theorems, and introduces the notion of "practical safety", which permits that the small safety violation can be made arbitrarily small with appropriate selection of design constants. We also extend

this algorithm to a Newton-based version of the ASfES algorithm, which we show should only be used in the case of a scalar parameter. The Newton-based version has an assignable rate of convergence of the parameter, only if the safety attractivity rate is not violated. Otherwise the convergence rate is modified to satisfy the attractivity rate requirement.

Another safe ES algorithm is offered next in Chapter 3 termed simply as "Safe ES", with guaranteed semiglobal stability and practical safety properties. The algorithm's design is different than that of ASfES, sacrificing user specified attractivity in favor of semiglobal convergence to an arbitrarily small neighborhood of the safe optimum. We make use of the current state of the art results on the semiglobal properties of ES combined with a Lypaunov based argument to achieve the key results of Chapter 3, and provide a general class of nonconvex constrained optimization problems which may be solved by Safe ES.

Finally we demonstrated the usefulness of Safe ES in Chapter 4, providing further rationale to real world problems in the operation of particle accelerators. The first two examples we present are conducted in validated simulators used to study beam dynamics of the linear accelerator at Los Alamos Neutron Science Center. In the simulated examples, we show that Safe ES can be used to perform spot size tuning (altering the size of the beam in the x and y directions, perpendicular to its direction of travel) in a safe manner, and compare its effectiveness relative to that of classical ES. In the third example of this chapter, we described a real experiment on the beam where Safe ES successfully tunes three bending magnets to steer the beam through a collimator safely. In this final example, we specify that the maximum beam survive through the collimator, and consider the measurement of the impact of the beam with the collimator to be a measure of safety, not allowing it to exceed a certain value too a large degree.

5.2 Future Work

There are several key areas surrounding the research presented in this dissertation which may be explored. The first is the consideration of delays into the measurements of the objective

and safety metrics. Delays have been studied in several ES settings although they have not been studied in an ES setting where safety is of concern. The need for the study of delays in systems where safety is required has many obvious applications – some of which arise in accelerator systems. For example, delays may arise where a diagnostic tool placed inside of the beam pipe requires several bunches (or passes of the beam through the tool) before it can produce a reading. Another common situation arises in large accelerator facilities where the reading of a measurement happens on a local wireless network, with a fixed update rate at which process variables can be read.

Another theoretical question concerns the semiglobality of the ASfES presented in Chapter 2. We showed that only local stability holds, although simulations show that indeed initial conditions (especially safe initial conditions) which are far away from the constrained optimum may still converge with the appropriate design constants. The design of Safe ES in Chapter 3 achieves semiglobality by slowing the movement of the parameter, which is achieved by choosing $k\omega_{\mathrm{f}}$ small. If this property can be demonstrated rigorously for the design of ASfES it may provide a way forward in demonstrating stronger convergence properties.

We also imagine that Safe ES can be utilized in more complex settings involving the simultaneous tuning of a machine learning model, which provides an estimate of a system metric, and the safe tuning of the the accelerator parameters in order to track this diagnostic metric. Here, both the tuning of this machine learning model and the beam parameters can be performed by classical ES and Safe ES on different time scales – the estimate provided by the diagnostic is provided fast, while the tuning of the accelerator parameters are performed slowly, in a safe manner. This work is presently ongoing and will be presented in 2024 at the International Particle Accelerator Conference (IPAC). Safe ES may prove a useful addition to the tool box for many other accelerator applications as accelerator systems are such complex machines, relationships between tuning parameters and the various system metrics are extremely difficult to know. We also hope to explore many other real life experiments applying Safe ES in accelerator facilities in the future. We are currently trying to make Safe ES more accessible as a tuning algorithm which

can be used by beam operators – individuals whose job is to continually monitor the performance of the beam. Developing a graphical user interface for this purpose is a currently ongoing project.

Bibliography

[1] I. Abel, M. Janković, and M. Krstić, "Constrained control of multi-input systems with distinct input delays," *International Journal of Robust and Nonlinear Control*, 2024.

[2] I. Abel, D. Steeves, M. Krstić, and M. Janković, "Prescribed-time safety design for strict-feedback nonlinear systems," *IEEE Transactions on Automatic Control*, pp. 1–16, 2023.

[3] V. Adetola and M. Guay, "Parameter convergence in adaptive extremum-seeking control," *Automatica*, vol. 43, no. 1, pp. 105–110, 2007.

[4] D. R. Agrawal and D. Panagou, "Safe and robust observer-controller synthesis using control barrier functions," *IEEE Control Systems Letters*, vol. 7, pp. 127–132, 2022.

[5] M. Aiba, M. Böge, N. Milas, and A. Streun, "Ultra low vertical emittance at sls through systematic and random optimization," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 694, pp. 133–139, 2012.

[6] A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, "Control barrier functions and input-to-state safety with application to automated vehicles," *IEEE Transactions on Control Systems Technology*, 2023.

[7] A. Allibhoy and J. Cortés, "Control barrier function based design of gradient flows for constrained nonlinear programming," *arXiv preprint arXiv:2204.01930*, 2022.

[8] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *IEEE Conference on Decision and Control*, 2014, pp. 6271–6278.

[9] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, pp. 3861–3876, 2017.

[10] ——, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.

[11] K. B. Ariyur and M. Krstic, *Real-time optimization by extremum-seeking control*. John Wiley & Sons, 2003.

[12] P. Arpaia, G. Azzopardi, F. Blanc, G. Bregliozzi, X. Buffat, L. Coyle, E. Fol, F. Giordano, M. Giovannozzi, T. Pieloni *et al.*, "Machine learning for beam dynamics studies at the cern large hadron collider," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 985, p. 164652, 2021.

[13] K. T. Atta, M. Guay, and R. Lucchese, "A geometric phasor extremum seeking control approach with measured constraints," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 1494–1500.

[14] A. Bacciotti and F. Ceragioli, "Stability and stabilization of discontinuous systems and nonsmooth lyapunov functions," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 4, pp. 361–376, 1999.

[15] P. Blackman, "Extremum-seeking regulators," in *An exposition of adaptive control*. Macmillan, 1962.

[16] F. Blanchini, S. Miani *et al.*, *Set-theoretic methods in control*. Springer, 2008, vol. 78.

[17] W. Blokland, K. Rajput, M. Schram, T. Jeske, P. Ramuhalli, C. Peters, Y. Yucesan, and A. Zhukov, "Uncertainty aware anomaly detection to predict errant beam pulses in the oak ridge spallation neutron source accelerator," *Physical Review Accelerators and Beams*, vol. 25, no. 12, p. 122802, 2022.

[18] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[19] J. Breeden and D. Panagou, "Guaranteed safe spacecraft docking with control barrier functions," *IEEE Control Systems Letters*, vol. 6, pp. 2000–2005, 2021.

[20] X. Chen, J. I. Poveda, and N. Li, "Continuous-time zeroth-order dynamics with projection maps: Model-free feedback optimization with safety guarantees," 2023.

[21] F. H. Clarke, *Optimization and nonsmooth analysis*. SIAM, 1990.

[22] M. H. Cohen and C. Belta, "Approximate optimal control for safety-critical systems with control barrier functions," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 2062–2067.

[23] J. Cortes, "Discontinuous dynamical systems," *IEEE Control Systems Magazine*, vol. 28, no. 3, pp. 36–73, 2008.

[24] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural lyapunov-barrier functions," in *Conference on Robot Learning*. PMLR, 2022, pp. 1724–1735.

[25] C. S. Draper and Y. T. Li, "Principles of optimalizing control systems and an application to the internal combusion engine," *The American Society of Mechanical Engineers*, 1951.

[26] J. Duris, D. Kennedy, A. Hanuka, J. Shtalenkova, A. Edelen, P. Baxevanis, A. Egger, T. Cope, M. McIntire, S. Ermon *et al.*, "Bayesian optimization of a free-electron laser," *Physical review letters*, vol. 124, no. 12, p. 124801, 2020.

[27] H.-B. Dürr and C. Ebenbauer, "A smooth vector field for saddle point problems," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 4654–4660.

[28] H.-B. Dürr, C. Zeng, and C. Ebenbauer, "Saddle point seeking for convex optimization problems," *IFAC Proceedings Volumes*, vol. 46, no. 23, pp. 540–545, 2013.

[29] A. Edelen, N. Neveu, M. Frey, Y. Huber, C. Mayes, and A. Adelmann, "Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems," *Physical Review Accelerators and Beams*, vol. 23, no. 4, p. 044601, 2020.

[30] A. Ghaffari, M. Krstić, and D. Nešić, "Multivariable newton-based extremum seeking," *Automatica*, vol. 48, no. 8, pp. 1759–1767, 2012.

[31] I. Griva, S. G. Nash, and A. Sofer, *Linear and nonlinear optimization*. Siam, 2009, vol. 108.

[32] M. Guay and D. Dochain, "A time-varying extremum-seeking control approach," *Automatica*, vol. 51, pp. 356–363, 2015.

[33] M. Guay and T. Zhang, "Adaptive extremum seeking control of nonlinear dynamic systems with parametric uncertainties," *Automatica*, vol. 39, no. 7, pp. 1283–1293, 2003.

[34] L. Hazeleger, D. Nešić, and N. van de Wouw, "Sampled-data extremum-seeking framework for constrained optimization of nonlinear dynamical systems," *Automatica*, vol. 142, p. 110415, 2022.

[35] S. Hirlaender and N. Bruchon, "Model-free and bayesian ensembling model-based deep reinforcement learning for particle accelerator control demonstrated on the fermi fel," *arXiv preprint arXiv:2012.09737*, 2020.

[36] R. A. Horn and C. R. Johnson, *Matrix analysis*. Cambridge university press, 2012.

[37] S.-C. Hsu, X. Xu, and A. D. Ames, "Control barrier function based quadratic programs with application to bipedal robotic walking," in *2015 American Control Conference (ACC)*. IEEE, 2015, pp. 4542–4548.

[38] E.-C. Huang, C. E. Taylor, and P. M. Anisimov, "Enhanced beam diagnostics with existing bppms via gpu-powered multi-particle simulation," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2021.

[39] X. Huang, "Robust simplex algorithm for online optimization," *Physical Review Accelerators and Beams*, vol. 21, no. 10, p. 104601, 2018.

[40] J. S. John, C. Herwig, D. Kafkes, J. Mitrevski, W. A. Pellico, G. N. Perdue, A. Quintero-Parra, B. A. Schupbach, K. Seiya, N. Tran *et al.*, "Real-time artificial intelligence for accelerator control: A study at the fermilab booster," *Physical Review Accelerators and Beams*, vol. 24, no. 10, p. 104601, 2021.

[41] V. Kain, S. Hirlander, B. Goddard, F. M. Velotti, G. Z. Della Porta, N. Bruchon, and G. Valentino, "Sample-efficient reinforcement learning for cern accelerator control," *Physical Review Accelerators and Beams*, vol. 23, no. 12, p. 124801, 2020.

[42] V. V. Kazakevich, "Technique of automatic control of different processes to maximum or to minimum," USSR Patent Patent 66 335, 1943.

[43] V. Kazakevich, "On extremum seeking," *Moscow High Technical University*, 1944.

[44] H. K. Khalil, *Nonlinear systems; 3rd ed.* Upper Saddle River, NJ: Prentice-Hall, 2002.

[45] S. Z. Khong, D. Nešić, Y. Tan, and C. Manzie, "Unified frameworks for sampled-data extremum seeking control: Global optimisation and multi-unit systems," *Automatica*, vol. 49, no. 9, pp. 2720–2733, 2013.

[46] N. King, E. Ables, K. Adams, K. Alrick, J. Amann, S. Balzar, P. Barnes Jr, M. Crow, S. Cushing, J. Eddleman *et al.*, "An 800-mev proton radiography facility for dynamic experiments," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 424, no. 1, pp. 84–91, 1999.

[47] J. Kirschner, M. Mutny, N. Hiller, R. Ischebeck, and A. Krause, "Adaptive and safe bayesian optimization in high dimensions via one-dimensional subspaces," in *International Conference on Machine Learning*. PMLR, 2019, pp. 3429–3438.

[48] J. Kirschner, M. Mutnỳ, A. Krause, J. C. de Portugal, N. Hiller, and J. Snuverink, "Tuning particle accelerators with safety constraints using bayesian optimization," *Physical Review Accelerators and Beams*, vol. 25, no. 6, p. 062802, 2022.

[49] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE control systems letters*, vol. 3, no. 1, pp. 108–113, 2018.

[50] M. Krstic, "Inverse optimal safety filters," *IEEE Transactions on Automatic Control*, 2023.

[51] M. Krstic and M. Bement, "Nonovershooting control of strict-feedback nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1938–1943, 2006.

[52] M. Krstic and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000.

[53] C. Labar, E. Garone, M. Kinnaert, and C. Ebenbauer, "Constrained extremum seeking: a modified-barrier function approach," *IFAC-PapersOnLine*, vol. 52, no. 16, pp. 694–699, 2019.

[54] M. Leblanc, "Sur l'electrification des chemins de fer au moyen de courants alternatifs de frequence elevee," *Revue générale de l'électricité*, vol. 12, no. 8, pp. 275–277, 1922.

[55] C.-K. Liao, C. Manzie, A. Chapman, and T. Alpcan, "Constrained extremum seeking of a mimo dynamic system," *Automatica*, vol. 108, p. 108496, 2019.

[56] S.-J. Liu and M. Krstic, "Stochastic averaging in discrete time and its applications to extremum seeking," *IEEE Transactions on Automatic control*, vol. 61, no. 1, pp. 90–102, 2015.

[57] J. C. Luxat and L. H. Lees, "Stability of peak-holding control systems," *IEEE Transactions on Industrial Electronics and Control Instrumentation*, vol. IECI-18, no. 1, p. 11 – 15, 1971.

[58] F. Mayet, M. Hachmann, K. Floettmann, F. Burkart, H. Dinter, W. Kuropka, T. Vinatier, and R. Assmann, "Predicting the transverse emittance of space charge dominated beams using the phase advance scan technique and a fully connected neural network," *Physical Review Accelerators and Beams*, vol. 25, no. 9, p. 094601, 2022.

[59] S. Meerkov, "Asymptotic methods for investigating a class of forced states in extremal systems," *Automation and Remote Control*, vol. 28, no. 12, pp. 1916–1920, 1967.

[60] K. S. Miller, "On the inverse of the sum of matrices," *Mathematics magazine*, vol. 54, no. 2, pp. 67–72, 1981.

[61] T. G. Molnar, A. K. Kiss, A. D. Ames, and G. Orosz, "Safety-critical control with input delay in dynamic environment," *IEEE transactions on control systems technology*, 2022.

[62] I. Morosanov, "Method of extremum control," *Automatic and Remote Control*, vol. 18, pp. 1077–1092, 1957.

[63] M. Nagumo, "Über die lage der integralkurven gewöhnlicher differentialgleichungen," *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, vol. 24, pp. 551–559, 1942.

[64] D. Nesić, Y. Tan, W. H. Moase, and C. Manzie, "A unifying approach to extremum seeking: Adaptive schemes based on estimation of derivatives," in *49th IEEE conference on decision and control (CDC)*. IEEE, 2010, pp. 4625–4630.

[65] D. Nesic, A. Mohammadi, and C. Manzie, "A framework for extremum seeking control of systems with parameter uncertainties," *IEEE Transactions on Automatic Control*, vol. 58, no. 2, pp. 435–448, 2012.

[66] V. Obabkov, "Theory of multichannel extremal control systems with sinusoidal probe signals," *Automation and Remote Control*, vol. 28, pp. 48–54, 1967.

[67] T. R. Oliveira, M. Krstić, and D. Tsubakino, "Extremum seeking for static maps with delays," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1911–1926, 2016.

[68] Y. Ostrovsky, "Extremum regulation," *Avtomat. i Telemechan.*, vol. 18, pp. 852–858, 1957.

[69] X. Pang *et al.*, "Advances in proton linac online modeling," *IPAC'15*, pp. 2423–2427, 2015.

[70] J. I. Poveda, M. Krstić, and T. Başar, "Fixed-time nash equilibrium seeking in time-varying networks," *IEEE Transactions on Automatic Control*, vol. 68, no. 4, pp. 1954–1969, 2022.

[71] J. I. Poveda and N. Quijano, "Shahshahani gradient-like extremum seeking," *Automatica*, vol. 58, pp. 51–59, 2015.

[72] J. I. Poveda and A. R. Teel, "A framework for a class of hybrid extremum seeking controllers with dynamic inclusions," *Automatica*, vol. 76, pp. 113–126, 2017.

[73] Y. Rahman, M. Jankovic, and M. Santillo, "Driver intent prediction with barrier functions," in *2021 American Control Conference (ACC)*, 2021, pp. 224–230.

[74] R. T. Rockafellar, *Convex Analysis*. Princeton university press, 1970, vol. 18.

[75] P. K. Roy, C. E. Taylor, C. Pillai, and Y. K. Batygin, "Comparison of profile measurements and transport beam envelope predictions along the 80-m lansce prad beamline," *Journal of Physics: Conference Series*, vol. 1067, no. 6, p. 062002, sep 2018. [Online]. Available: https://dx.doi.org/10.1088/1742-6596/1067/6/062002

[76] A. Scheinker, "Adaptive machine learning for time-varying systems: low dimensional latent space tuning," *Journal of Instrumentation*, vol. 16, no. 10, p. P10008, 2021.

[77] ——, "100 years of extremum seeking: A survey," *Automatica*, vol. 161, p. 111481, 2024.

[78] A. Scheinker, D. Bohler, S. Tomin, R. Kammering, I. Zagorodnov, H. Schlarb, M. Scholz, B. Beutner, and W. Decking, "Model-independent tuning for maximizing free electron laser pulse energy," *Physical review accelerators and beams*, vol. 22, no. 8, p. 082802, 2019.

[79] A. Scheinker, F. Cropp, and D. Filippetto, "Adaptive autoencoder latent space tuning for more robust machine learning beyond the training set for six-dimensional phase space diagnostics of a time-varying ultrafast electron-diffraction compact accelerator," *Physical Review E*, vol. 107, no. 4, p. 045302, 2023.

[80] ——, "Adaptive autoencoder latent space tuning for more robust machine learning beyond the training set for six-dimensional phase space diagnostics of a time-varying ultrafast electron-diffraction compact accelerator," *Physical Review E*, vol. 107, no. 4, p. 045302, 2023.

[81] A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, "Demonstration of model-independent control of the longitudinal phase space of electron beams in the linac-coherent light source with femtosecond resolution," *Physical review letters*, vol. 121, no. 4, p. 044801, 2018.

[82] A. Scheinker, S. Hirlaender, F. M. Velotti, S. Gessner, G. Z. Della Porta, V. Kain, B. Goddard, and R. Ramjiawan, "Online multi-objective particle accelerator optimization of the awake electron beam line for simultaneous emittance and orbit control," *AIP Advances*, vol. 10, no. 5, 2020.

[83] A. Scheinker, E.-C. Huang, and C. Taylor, "Extremum seeking-based control system for particle accelerator beam loss minimization," *IEEE Transactions on Control Systems Technology*, vol. 30, no. 5, pp. 2261–2268, 2021.

[84] A. Scheinker, X. Huang, and J. Wu, "Minimization of betatron oscillations of electron beam injected into a time-varying lattice via extremum seeking," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 1, pp. 336–343, 2018.

[85] A. Scheinker and M. Krstić, "Minimum-seeking for clfs: Universal semiglobally stabilizing feedback under unknown control directions," *IEEE Transactions on Automatic Control*, vol. 58, no. 5, pp. 1107–1122, 2012.

[86] ——, "Extremum seeking with bounded update rates," *Systems & Control Letters*, vol. 63, pp. 25–31, 2014.

[87] A. Scheinker, X. Pang, and L. Rybarcyk, "Model-independent particle accelerator tuning," *Phys. Rev. ST Accel. Beams*, vol. 16, p. 102803, Oct 2013. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevSTAB.16.102803

[88] A. Scheinker and D. Scheinker, "Bounded extremum seeking with discontinuous dithers," *Automatica*, vol. 69, pp. 250–257, 2016.

[89] E. Schuster, C. Xu, N. Torres, E. Morinaga, C. Allen, and M. Krstic, "Beam matching adaptive control via extremum seeking," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 581, no. 3, pp. 799–815, 2007.

[90] Y. Tan, Y. Li, and I. M. Mareels, "Extremum seeking for constrained inputs," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2405–2410, 2013.

[91] Y. Tan, D. Nešić, and I. Mareels, "On non-local stability properties of extremum seeking control," *IFAC Proceedings Volumes*, vol. 38, no. 1, pp. 550–555, 2005.

[92] ——, "On non-local stability properties of extremum seeking control," *Automatica*, vol. 42, no. 6, pp. 889–903, 2006.

[93] Y. Tan, D. Nešić, I. M. Mareels, and A. Astolfi, "On global extremum seeking in the presence of local extrema," *Automatica*, vol. 45, no. 1, pp. 245–251, 2009.

[94] A. R. Teel, L. Moreau, and D. Nesic, "A unified framework for input-to-state stability in systems with two time scales," *IEEE Transactions on Automatic Control*, vol. 48, no. 9, pp. 1526–1544, 2003.

[95] A. R. Teel and D. Popovic, "Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming," in *Proceedings of the 2001 American Control Conference.(Cat. No. 01CH37148)*, vol. 3. IEEE, 2001, pp. 2394–2399.

[96] H. S. Tsien, "Engineering cybernetics," *The Macmillan Company*, 1954.

[97] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.

[98] Y. Wang, S. S. Zhan, R. Jiao, Z. Wang, W. Jin, Z. Yang, Z. Wang, C. Huang, and Q. Zhu, "Enforcing hard constraints with soft barriers: Safe reinforcement learning in unknown stochastic environments," in *International Conference on Machine Learning*. PMLR, 2023, pp. 36 593–36 604.

[99] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," *IFAC Proceedings Volumes*, vol. 40, no. 12, pp. 462–467, 2007, 7th IFAC Symposium on Nonlinear Control Systems.

[100] A. Williams, M. Krstić, and A. Scheinker, "Practically safe extremum seeking," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 1993–1998.

[101] A. Williams, M. Krstic, and A. Scheinker, "Semi-global practical extremum seeking with practical safety," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 6774–6779.

[102] A. Wolski, M. A. Johnson, M. King, B. L. Militsyn, and P. H. Williams, "Transverse phase space tomography in an accelerator test facility using image compression and machine learning," *Physical Review Accelerators and Beams*, vol. 25, no. 12, p. 122803, 2022.

[103] G. Wu and K. Sreenath, "Safety-critical and constrained geometric control synthesis using control lyapunov and control barrier functions for systems evolving on manifolds," in *2015 American Control Conference (ACC)*. IEEE, 2015, pp. 2038–2044.

[104] W. Xiao, C. G. Cassandras, and C. A. Belta, "Bridging the gap between optimal trajectory planning and safety-critical control with applications to autonomous vehicles," *Automatica*, vol. 129, p. 109592, 2021.

[105] J. Zhu, Y. Chen, F. Brinker, W. Decking, S. Tomin, and H. Schlarb, "High-fidelity prediction of megapixel longitudinal phase-space images of electron beams using encoder-decoder neural networks," *Physical Review Applied*, vol. 16, no. 2, p. 024005, 2021.