DISS. ETH NO. 26970

Optimization Algorithms as Feedback Controllers for Power System Operations

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Adrian Robert Hauswirth

M.Sc., Eidgenössische Technische Hochschule Zürich (ETHZ)
B.Sc., École Polytechnique Fédérale de Lausanne (EPFL)
born on 12. Oktober 1988
citizen of Gsteig b. Gstaad, Schweiz

accepted on the recommendation of

Prof. Dr. Florian Dörfler, examiner
Prof. Dr. Gabriela Hug, co-examiner
Prof. Dr. Steven Low, co-examiner
Prof. Dr. Andrew Teel, co-examiner

2020

orcid.org/0000-0001-6205-0423

ISBN 978-3-907234-16-7
DOI 10.3929/ethz-b-000441303

To my family

In the beginner's mind there are many possibilities, but in the expert's there are few.

Zen Mind, Beginner's Mind
Shunryu Suzuki

The paradox is now fully established that the utmost abstractions are the true weapons with which to control our thought of concrete fact.

Science and the Modern World
Alfred North Whitehead

Abstract

In recent decades the design of feedback control systems that steer a physical system to an optimal state, rather than tracking a set-point, has taken many forms ranging from extremum seeking, to economic model predictive control, to internet congestion control.

One application emerging as a prime subject for this type of *closed-loop optimization* is the operation of power systems in the presence of intermittent renewable energy production and new consumption patterns. The vision shared by industry and academia is to implement *autonomous power grids* that can react intelligently to outages, sudden changes in renewable output, and other events. Simultaneously, the safety and reliability of power supply must not be jeopardized, and all available resources should be used to their full economic potential.

Even from a simplified viewpoint, this problem poses conceptual challenges that are not covered by existing approaches to closed-loop optimization: First, the flow of electric power in a grid can be highly nonlinear, especially under critical operating conditions. Second, running an electricity grid relies on the satisfaction of many constraints, such as voltage limits or thermal line ratings. Further, power system operation is subject to limited generation capacity and inelastic demand. Third, model information about large-scale, interconnected, and open power grids is generally limited to an approximate model for the steady-state power flow.

In summary, the safe, reliable, and optimal real-time operation of power grids requires a new scalable and model-light approach to closed-loop optimization that applies to nonlinear systems with complicated constraints. Towards this end, this thesis presents contributions to the recent field of *online feedback-based optimization* with a focus on non-

convex constrained optimization problems.

Projected gradient flows are the technical and mathematical center point of this thesis. We study their approximation and implementation as feedback controllers. We establish various convergence and robustness properties, and we show how they inform the design of real-time closed-loop optimization schemes for power systems operations.

In the first part, we study the general class of *projected dynamical system* (PDS). We extend their definition to arrive at a coordinate-free formulation, which makes them adequate for modeling physical systems. We further introduce and study *anti-windup approximation* (AWA) and *linearized output projection* (LOP) *discretizations* of PDS. AWAs result from high-gain anti-windup controllers and motivate the exploitation of physical input saturation for constraint enforcement in feedback-based optimization. LOP discretizations, on the other hand, offer a possibility to reliably enforce output constraints with a combination of first-order model information and output measurements.

The second part of the thesis emphasizes the convergence and robustness properties of projected gradient flows. In particular, we establish nominal convergence and the relation between stability and optimality for coordinate-free projected gradient flows. We show that these properties are preserved for anti-windup approximations and LOP discretizations of projected gradient flows.

Moreover, we study the stability requirements when interconnecting a controller that implements a projected gradient flow with a dynamic physical system. This type of robustness is particularly important for the envisioned application where feedback-based optimization schemes must not destabilize the power system dynamics.

To conclude the second part, we extend the applicability of projected gradient flows to time-varying setups by introducing *sweeping gradient flows* and quantifying their tracking performance in a special case. Again, this topic is essential for power systems, where electricity demand, available generation capacity, and other factors change over time, and feedback-based optimization schemes need to reliably track the optimal state.

Finally, in the third part, we illustrate these theoretical results for the closed-loop optimization of power systems. However, rather than presenting algorithms tailored to specific use cases or grid levels, we focus our discussion on the general control designs, possibilities, and fundamental limitations. For these designs, we also infer theoretical guarantees from the results of the first two parts of the thesis.

All in all, this thesis provides a fundamental look at the possibilities of implementing projected gradient flows in closed loop with a physical plant and points at the opportunities and potential shortcomings for the optimal real-time operation of power systems.

Zusammenfassung

Regelsysteme, die keine Führungsgröße benötigen, sondern direkt ein inhärentes Gütekriterium optimieren, haben in den letzten Jahrzehnten unterschiedliche Formen angenommen. Beispiele dafür reichen von Extremwertregelung über ökonomisch-modellprädiktive Regelung bis hin zur Engpasssteuerung in Computernetzwerken wie dem Internet.

Eine vielversprechende Anwendung für diese Art von *Optimierung im geschlossenen Regelkreis* ist der Betrieb von Energiesystemen unter dem Einfluss schwankender erneuerbar Energien und neuartiger Verbrauchsformen. Eine breit abgestützte Vision besteht darin, Stromnetze so weit zu automatisieren, dass sie intelligent auf Ausfälle, Schwankungen in der erneuerbaren Energieproduktion und auf andere Ereignisse reagieren können. Dabei sollen die Versorgungssicherheit nicht gefährdet und alle verfügbaren Ressourcen bestmöglich genutzt werden.

Allerdings wirft dieses Problem selbst unter vereinfachenden Annahmen konzeptionelle Herausforderungen auf, die von bestehenden Ansätzen nicht abgedeckt werden: Erstens kann der Leistungsfluss in einem Stromnetz hochgradig nicht-linear sein. Dies ist insbesondere unter kritischen Betriebsbedingungen der Fall. Zweitens hängt der sichere Betrieb eines Stromnetzes von der Erfüllung vieler Nebenbedingung wie der Spannungshaltung oder der Leitungskapazitäten ab. Dazu kommen die begrenzte Erzeugungskapazität und eine schwankende, aber inelastische Nachfrage. Drittens sind Netzmodelle für große, vernetzte Stromnetze im Allgemeinen nur beschränkt verfügbar und erlauben ausschließlich Aussagen über den stationären Leistungsfluss.

Zusammenfassend erfordert der sichere, zuverlässige und kostengünstige Echtzeitbetrieb von Stromnetzen einen neuen, skalierbaren und modellfreien Ansatz zur Optimierung im geschlossenen Regelkreis, der

auf nicht-lineare Systeme mit komplizierten Nebenbedingungen anwendbar ist. Zu diesem Zweck werden in dieser Arbeit Beiträge zum aktuellen Forschungsfeld der *regelungsbasierten Echtzeitoptimierung* für nicht-konvexe Probleme erarbeitet.

Der technische und mathematische Schwerpunkt dieser Arbeit sind *projizierte Gradientenflüsse*. Sowohl deren Approximation und Implementierung als Regler als auch grundlegende Konvergenz- und Robustheitseigenschaften sind Kernbestandteil dieser Dissertation. Daneben werden konkrete Reglerentwürfe für die Echtzeitoptimierung von Energiesystemen vorgestellt.

Der erste Teil dieser Arbeit befasst sich mit *projizierten dynamischen Systemen* (PDS). Eine neue, verallgemeinerte Definition dieser unstetigen Systeme führt zu einer koordinatenfreien Formulierung, die für die Modellierung physikalischer Systeme geeignet ist. Ferner werden *anti-Windup Approximationen* und *linearisierte Ausgangsprojektion* (LOP) als Näherungs- und Implementierungsverfahren von PDS eingeführt und untersucht. Anti-Windup Approximationen resultieren aus der Verwendung von hochverstärkten anti-Windup-Reglern und zeigen, wie physikalische Sättigung von Eingangsgrößen zur Erfüllung von Nebenbedingungen genutzt werden kann. Dem gegenüber erlauben LOP Diskretisierungen Nebenbedingungen auf Ausgangsgrößen zuverlässig zu erfüllen, indem Messungen mit Modellinformation erster Ordnung kombiniert werden.

Im zweiten Teil werden Konvergenz und Robustheit diskutiert. Das Augenmerk liegt dabei auf der Konvergenz koordinatenfreier projizierte Gradientenflüsse und der Beziehung zwischen Stabilität und Optimalität. Es wird gezeigt, dass anti-Windup-Approximationen und LOP Diskretisierungen diese Eigenschaften erhalten.

Außerdem werden Stabilitätsbedingungen für die Rückkopplung einer Regelstrecke und einem projizierten Gradientenfluss als Regler ausgearbeitet. Diese Robustheit ist besonders wichtig für die Anwendung in Stromnetzen, bei der eine Echtzeitoptimierung die schnelle Wechselstromdynamik nicht destabilisieren darf.

Zum Abschluss des zweiten Teils wird die Anwendbarkeit projizierter

Gradientenflüsse auf nicht-stationäre Probleme ausgeweitet. Zu diesem Zwecke werden *sweeping gradient flows* eingeführt und ihre dynamische Regelgüte in Spezialfällen quantifiziert. Dieses Thema ist wichtig im Zusammenhang mit Energiesystemen, wo sich der Energiebedarf, die verfügbare Erzeugungskapazität und andere Faktoren im Laufe der Zeit ändern und regelungsbasierte Optimierungsapplikationen dem optimalen Zustand zuverlässig folgen müssen.

Im dritten Teil werden diese theoretischen Ergebnisse anhand der regelungsbasierten Optimierung von Energiesystemen veranschaulicht. Anstatt Algorithmen vorzustellen, die auf spezifische Anwendungsfälle oder Netzebenen zugeschnitten sind, konzentriert sich die Diskussion auf allgemeine Reglerentwürfe, ihre Möglichkeiten und grundlegenden Limitationen sowie auf theoretische Garantien, die aus den Ergebnissen der ersten beiden Teile der Arbeit abgeleitet werden können.

Alles in allem bietet diese Arbeit einen grundlegenden Einblick in die Möglichkeiten der Implementierung projizierter Gradientenflüsse im geschlossenen Regelkreis mit einer Regelstrecke.

Acknowledgements

Looking back at the past years, I cannot help to be nostalgic about the countless great memories that I have collected and all the bright and inspiring people I have met during my Ph.D.

First and foremost, I would like to express my deepest gratitude to my advisors Prof. Florian Dörfler and Prof. Gabriela Hug for the fantastic opportunity to do my Ph.D. under their joint supervision. Their confidence, enthusiasm, and patience have continuously inspired me to tackle problems way beyond what I had thought I was capable of. I will never forget the amount of academic freedom, trust, and unflinching support I received from them — Florian and Gaby, you are phenomenal supervisors. Thank you for everything!

It was an incredible luck for me to work with Saverio Bolognani throughout my Ph.D. studies. I sincerely appreciate him for always having taken the time to brainstorm and discuss new ideas, but also help me figure out the most mundane tasks of everyday academic life.

I am especially grateful to Professor Andrew Teel for hosting me at UC Santa Barbara for half a year and being on my exam committee. Working with Andy has been an eye-opening experience. His vast knowledge of control theory and dedication to mathematical rigor are unparalleled and have helped me step up my game. Thank you, Andy, for being so patient with me.

Moreover, my sincere thanks go to Professor Steven Low. His work has greatly influenced my own thinking, and I am honored to have him on my examination committee.

Without Prof. Tyler Summers I probably would not have considered doing a Ph.D., nor would I have gotten a position at IfA nor PSL - and much less at both. Thanks, Tyler, for putting me on this track and

helping me succeed. Further, I wish to thank Dr. Joe Warrington, the first person I met at IfA. Ever since he supervised my semester project, our paths have crossed again and again. His sharp reasoning, honest opinions, and encouraging words never fail to impress me.

I further thank my senior colleagues Prof. Dominic Gross, Prof. John Simpson-Porco, Prof. Enrique Mallada, Dr. Andrey Bernstein, and Dr. Marcello Colombino for important technical inputs and advice in the past years. I owe special thanks to Prof. Emiliano Dall'Anese for his hospitality during my stay at NREL.

The work presented in this thesis would not have been possible without the insights gained from numerous student projects and my co-authors' inputs and hard work. I particularly wish to thank Dr. Felix Rey, Alessandro Zanardi, Sandeep Menta, and Ivo Caduff. And my very utmost gratitude goes to Irina Subotić, Verena Häberle, and Lukas Ortmann for their commitment to high-quality research; for not giving up, no matter how hard the problem; and for getting me out of my comfort zone once in a while. You are awesome, guys!

Further, I would like to thank all the professors at IfA and PSL, current and emeriti, for creating such a world-class research environment full of bright minds and genuinely friendly people. Well-organized seminars, coffee talks, group meetings, and social activities have made these years unforgettable.

I am indebted to Eve for being an awesome office mate and always having an open ear and time for ice cream. Under these circumstances, keeping my weight would not have been possible without the weekly "PSL runs." My special thanks go to Conor and Johanna helping to keep this tradition alive.

Thanks to all colleagues at PSL and IfA for creating such an enjoyable atmosphere, both in terms of a great work environment, and also during social activities. My special thanks go to Alex, Angeliki, Ben, Damian, Dario, Dmitry, Eva, Goran, Jack, Jeremy, Jo, Line, Marius, Matthias, Michael, Miguel, Nadezhda, Nicolas, Paul, Sandro, Stavros, Theo, Thierry, Tomas, Tommaso, Uros, and Yvonne for "early lunching", sharing teaching duties, being great office mates, helping me out with

the little things, and for all the fruitful research discussions.

I would also like to express my gratitude to Judith, Sabrina, and Tanja for all their administrative work, enabling the rest of us to focus on research and teaching.

On a personal level, I would like to thank Marc H., Michael M., Michael T., Andreas, and Marc Z. for long hikes, beers, and discussions on almost any topic. Thanks to Livio for being a great flatmate. My special gratitude goes to Valentine for all the great, past and future, adventures together.

Finally, my deepest thanks go to my parents, Margrit and Robert, and my sister Sandra. They have been supporting me through all these steps with their unconditional love and sacrifice. Without you, I would not be where I am today.

Adrian Robert Hauswirth Zurich, August 2020

Contents

Nomenclature

Basic Notation

$:=$	equal by definition
$\equiv$	identically equal
$\emptyset$	empty set
$x \mapsto f(x)$	mapping of the point x into $f(x)$
$f(\cdot, y)$	mapping $x \mapsto f(x, y)$ for a fixed y
$\mathbb{R}^n$ ($\mathbb{C}^n$)	real (complex) n-dimensional Euclidean space
$\mathbb{R}^n_{\geq 0}$	non-negative orthant of $\mathbb{R}^n$
$\mathbb{S}^n_+$	space of square symmetric, positive definite $n \times n$-matrices
$\mathbb{B}$ ($\operatorname{int} \mathbb{B}$)	closed (open) unit ball at 0 of appropriate dimension
$\mathbb{I}$	identity matrix (of appropriate dimension)
$\mathbb{0}$	zero matrix (of appropriate dimension)
A^T	transpose of matrix A
$\lambda_A^{\max}$, $\lambda_A^{\min}$	max and min eigenvalues of a square symmetric matrix A
$\sigma_B^{\max}$, $\sigma_B^{\min}$	max and min singular values of a matrix B
κ_A	condition number of a matrix $A \in \mathbb{S}^n_+$
$\operatorname{rank} A$	rank of a matrix A (or linear map)

$\langle \cdot, \cdot \rangle$	Euclidean inner product
$\langle \cdot, \cdot \rangle_A$	inner product induced by $A \in \mathbb{S}^n_+$
$D_x f$	differential of a map f at point x
$Df(x; v)$	directional derivative of a map f at point x in direction v
$\nabla_x f$	Jacobian ($m \times n$-matrix) of a map $f : \mathbb{R}^n \to \mathbb{R}^m$ at x
$\operatorname{grad}_G \Phi$	gradient of Φ in the metric G
$\mathcal{L}_F \Phi(x)$	Lie derivative of function Φ along vector field F at x
C^0	continuous maps (with appropriate domain/co-domain)
C^1	class of continuously differentiable maps
$C^{0,1}$	class of locally Lipschitz maps
$C^{1,1}$	class of C^1 maps with locally Lipschitz derivative
C^2	class of twice continuously differentiable maps
$\mathcal{K}_\infty$	class of $\mathcal{K}_\infty$ functions
$\mathcal{KL}$	class of $\mathcal{KL}$ functions
$\operatorname{cl} \mathcal{C}$	closure of a set $\mathcal{C}$
$\operatorname{co} \mathcal{C}$ ($\overline{\operatorname{co}}\, \mathcal{C}$)	(closed) convex hull of a set $\mathcal{C}$
$\operatorname{dom} F$	domain of a set-valued map $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$
$\operatorname{epi} f$	epigraph of a single-valued map $f : \mathbb{R}^n \to \mathbb{R}$
$\operatorname{gph} F$	graph of set-valued mapping
$d_{\mathcal{C}}(x)$	Euclidean point-to-set distance from x to the set $\mathcal{C}$
$P_{\mathcal{C}}(x)$	Euclidean projection of the point x onto the set $\mathcal{C}$
$P^A_{\mathcal{C}}(x)$	projection of x onto $\mathcal{C}$ with respect to $A \in \mathbb{S}^n_+$.
$T_x \mathcal{C}$	tangent cone of the set $\mathcal{C}$ at point x

$T_x^C\mathcal{C}$	(Clarke) regular tangent cone of the set $\mathcal{C}$ at point x
$N_x\mathcal{C}$	Euclidean normal cone of the set $\mathcal{C}$ at point x
$\bar{N}_x^{G,\alpha}\mathcal{C}$	cone of α-proximal normal vectors with respect to the metric G for the set $\mathcal{C}$ at the point x
$N_x^G\mathcal{C}$	normal cone of the set $\mathcal{C}$ at point x with respect to G
$\Pi_{\mathcal{C}}^G[w](x)$	projection of w onto $T_x\mathcal{C}$ at $x \in \mathcal{C}$ with respect to G
$\Pi_{\mathcal{C}}[w](x)$	Euclidean projection w onto $T_x\mathcal{C}$
$\Pi_{\mathcal{C}}^G[f](x)$	projection of $f(x)$ onto $T_x\mathcal{C}$ with respect to G
$\mathrm{K}[F]$	Krasovskii-regularization of the set-valued map F

List of Symbols for Part III

j	imaginary unit
$(\cdot)^*$	complex conjugate
$(\mathbf{V}, \mathbf{E})$	(directed) graph representing the power grid
$\mathbf{E}^{\mathrm{t}}(k), \mathbf{E}^{\mathrm{f}}(k)$	set of incoming and outgoing edges from node k, resp.
$\mathbf{N}(k)$	set of neighboring nodes of node k
N	number of nodes in power grid, i.e., $N := \lvert\mathbf{V}\rvert$
M	number of edgeds in power grid, i.e., $M := \lvert\mathbf{E}\rvert$
e, s	complex bus voltage phasors [p.u.] and apparent power injections [MVA]
v, θ	bus voltage magnitudes [p.u.] and angles [rad]
$\underline{v}, \overline{v}$	lower and upper limits on bus voltage magnitudes
p, q	nodal real and reactive power injections [MW]/[MVar]
$\underline{p}, \overline{p}$	lower and upper limits on real power injections

$\underline{q}, \overline{q}$	lower and upper limits on reactive power injections
i^{t}_{km}	current flowing into line from k to m at node k
i^{f}_{km}	current flowing into line from k to m at node l
$\overline{i}$	vector of upper line current limits [p.u.]
$\mathbf{Y}$	complex-valued bus admittance matrix
$\mathbf{Y}^{\mathrm{t}}, \mathbf{Y}^{\mathrm{f}}$	complex-valued line admittance matrices
$\mathrm{g}_{km}, \mathrm{b}_{km}$	conductance and susceptance for line from k to m [p.u.]
$\mathrm{r}_{km}, \mathrm{x}_{km}$	resistance and reactance for line from k to m [p.u.]
$\mathrm{b}^{\mathrm{sh}}_{km}$	line charging susceptance for line from k to m [p.u.]
y_{km}	complex line admittance for line from k to m [p.u.]
$\mathrm{y}^{\mathrm{sh}}_{k}$	complex nodal shunt admittance at node k [p.u.]
$\mathrm{n}_{km}, \theta^{\mathrm{shift}}_{km}$	tap ratio and phaseshift at from-end of line from k to m
$\mathcal{M}$	power flow manifold (embedded in $\mathbb{R}^{4N}$)
$\mathcal{M}^{\mathrm{high}}$	high-voltage component of power flow manifold
$\mathcal{U}, \mathcal{U}^{\mathrm{high}}$	input constraint set, projection of $\mathcal{M}^{\mathrm{high}}$ onto input space
h	steady-state input-output map defined on $\mathcal{U}^{\mathrm{high}}$
$H(u)^T$	steady-state input-to-(output,input) sensitivites
$\mathcal{X}$	engineering constraint set
Φ	cost function in input and output variables
$\tilde{\Phi}$	reduced cost function, i.e., $\tilde{\Phi}(u) := \Phi(h(u), u)$
$\Psi_{\mathcal{X}}(u)$	quadratic penalty function for engineering constraint set
G^{E}	explicit Euclidean metric defined on $\mathcal{U}^{\mathrm{high}}$
G^{I}	implicit Euclidean metric on $\mathcal{U}^{\mathrm{high}}$
G^{N}	Newton metric on $\mathcal{U}^{\mathrm{high}}$

Acronyms

ACOPF AC optimal power flow

ACPF AC power flow

AWA anti-windup approximation

CQ constraint qualification

DVI differential variational inequality

ES extremum seeking

gas globally asymptotically stable

isc inner semicontinuous

KKT Karush-Kuhn-Tucker

las locally asymptotically stable

LICQ linear independence constraint qualification

LOP linearized output projection

lti linear time-invariant

MA modifier adaptation

MPC model predictive control

ODE ordinary differential equation

osc outer semicontinuous

PDS projected dynamical system

PFM power flow manifold

QP quadratic program

RTI real-time iteration

spas semiglobally practically asymptotically stable

SQP sequential quadratic programming

SSOSC strong second-order sufficiency condition

CHAPTER 1

Introduction & Survey

Most advances in mathematical optimization in the past decades have been geared towards computational implementations. The common viewpoint is that an optimization problem can be formulated, transformed, reduced, or relaxed. However, ultimately the necessary steps to solve the problem rely purely on numerical linear algebra, which can be implemented and run on a microprocessor. Then, this offline computed solution is used to reach a decision.

This paradigm is almost synonymous with the field of management science and operations research [34, 129], which has flourished ever since the inception of linear programming in the mid-20th century. Today, this kind of offline optimization is applied in various disciplines ranging from econometrics over statistical and machine learning [37] to optimal control [32].

However, from a control engineering perspective, solving an optimization problem offline (with known problem data) and implementing its output as a decision is a "feedforward" approach. In contrast, in this thesis, we consider *feedback* approaches to constrained nonlinear optimization to drive a (possibly nonlinear dynamical) system towards an optimal state. This type of closed-loop optimization has mainly been

This chapter is adapted from [Ha8]. The content is self-contained, and the notation and style are independent of the main part of the thesis.

pursued for four reasons:

(i) to increase robustness against disturbances and uncertainty,

(ii) to reduce dependence on model information, i.e., to make the control *model-free*,

(iii) to minimize computational effort, and

(iv) to eliminate the need for an exogenous set-point or reference.

These reasons resonate well with the general feedback and feedforward paradigms advocated in control textbooks [108, 14], and we will further dwell on them below:

Robustness is key in optimization: More often than not, practical problems lack precise data. Parameters and states based on measurements and statistical inference are inherently flawed, and so is the solution of an optimization problem based on such data. The earliest attempts at addressing this issue have resulted in the general theory of *sensitivity analysis* for optimization problems [219] which asks the question of how a solution changes as problem parameters vary. From a more practical perspective, *robust optimization* [27] and *stochastic programming* [46] offer ways to incorporate uncertainty in the problem. However, these approaches are inherently conservative and can entail massive computational cost because they need to take into account the full set of possible instances. In contrast, using feedback, one needs to react only to the actual realization, thereby naturally achieves disturbance rejection and increases robustness.

More radical than improving robustness is the idea of rendering optimization schemes *model-free*: A physical system defines a set of constraints (either algebraic or dynamic and often time-varying). Therefore, it is only natural to probe a system for information and learn its behavior from measurements instead of building a static model from first principles or previously acquired data.

These physical constraints imposed by the system cannot only be learned; they can also be exploited to enforce constraints directly: Physical laws naturally couple inputs and outputs of a plant. Hence, for

any input, the physical system will produce an output that satisfies this input-output relation. Similarly, saturation effects (such as actuator limitations) will naturally guarantee that states and inputs do not exceed their limits. Especially if one does not risk the plant's destruction, this type of natural constraint satisfaction can be used to *lighten the computational load* of a closed-loop optimization scheme because the physical system acts as a "constraint enforcer".

Finally, compared to a standard feedback control loop, a closed-loop optimization setup can run *without external set-point* or reference. Instead, an economic objective can be directly optimized as long as a cost function can be specified. This feature is particularly powerful in combination with the inherent constraint enforcement. Whereas in classical control setups pre-computed set-points have to be feasible (e.g., lie within actuator limits), convergence to a feasible and optimal steady state is a defining feature of closed-loop optimization.

The boundary between feedforward and feedback optimization, however, is not always clear cut. For instance, *model predictive control* (MPC) uses feedback to achieve robustness but relies on an accurate model based on which an optimal control problem is formulated and solved to optimality at every iteration. On the other hand, some of the control schemes presented in this thesis require the computationally cheap solution of a simple *quadratic program* (QP) to compute a feasible descent direction at every iteration.

Nevertheless, we focus on control schemes that result in closed-loop dynamics that mimic optimization algorithms such as gradient descent or saddle-point methods. The following simple example of a gradient system interconnected with a physical plant illustrates this idea.

Example 1.1. Consider the problem of minimizing an objective $\Phi(y)$ as a function of the measured output y of a plant

$$\dot{\zeta} = f(\zeta, u) \qquad y = g(\zeta) + d\,, \tag{1.1}$$

where d denotes an additive disturbance to the output.

Assume that, for any fixed input u, the plant is asymptotically stable with fast-decaying transients such that for every u there exists a unique

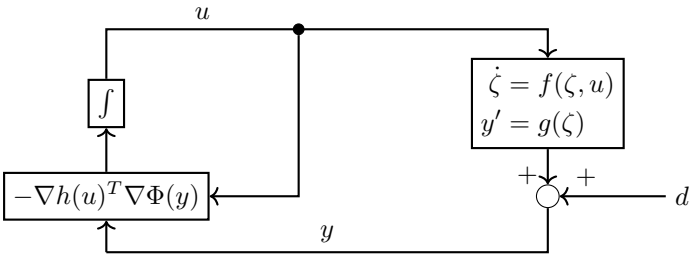

Figure 1.1: Simple Gradient Flow

output $y = h(u) + d$ where h is the steady-state input-to-output map, namely, $0 = f(h'(u), u)$ for some h' such that $h = g \circ h$. We assume h to be differentiable.

To minimize $\tilde{\Phi}(u) := \Phi(h(u) + d) = \Phi(y)$, we consider a gradient flow

$$\dot{u} = -\nabla \tilde{\Phi}(u)^T = -\nabla h(u)^T \nabla \Phi(h(u) + d)^T \,, \tag{1.2}$$

where $\nabla h(u)$ appears due to the chain rule applied to $\Phi(h(u) + d)$.

The gradient flow (1.2) is a closed system. However, recognizing $h(u) + d$ as the measurable output y, (1.2) can be easily transformed into an open system and interconnected with the plant (1.1), as shown in Figure 1.1 This yields the closed-loop dynamics

$$\begin{aligned} \text{plant} \begin{cases} \dot{\zeta} &= f(\zeta, u) \\ y &= g(\zeta) + d \end{cases} \\ \text{controller} \begin{cases} \dot{u} &= -\nabla h(u)^T \nabla \Phi(y)^T \,. \end{cases} \end{aligned} \tag{1.3}$$

It can be easily seen that any equilibrium point $(\zeta^\star, u^\star)$ of (1.3) is a steady-state of the plant and satisfies $\nabla \tilde{\Phi}(u^\star)^T = \nabla h(u^\star)^T \nabla \tilde{\Phi}(h(u^\star) + d)^T = 0$. Therefore, $u^\star$ is a critical point of $\tilde{\Phi}$ (and guaranteed to be a minimizer if $\tilde{\Phi}$ is convex).

Furthermore, concerning stability, if the plant dynamics are fast-decaying, we may assume that $y \approx h(u) + d$. In singular perturbation terminology, this approximation yields the so-called *reduced dynamics*, which, in the current case, are given by (1.2).

Crucially, the control law does not require explicit knowledge of h (nor of f, g). Instead, only the gradient of the cost function $\nabla\Phi(u)$ as well as the steady-state input-output sensitivities $\nabla h(u)$ are required. Moreover, the additive disturbance d does not need to be (explicitly) estimated and is fully rejected, i.e., an equilibrium is a critical point of $\Phi(h(u)+d)$, independently of the value of d. In practice, if $\nabla h(u)$ needs to be computed online, it is often beneficial to use the measurement y in the estimate of the input-output sensitivities $\nabla h(u)$. This is particularly the case when disturbances enter the plant non-additively.

Finally, if $f(\zeta, u) = A\zeta + Bu + w$ and $g(\zeta) = C\zeta$ form a *linear time-invariant* (lti) system and $\Phi(y) = \frac{1}{2}\|y - y^{\text{set}}\|^2$ is quadratic with a set-point y^{set}, we recover a simple linear integral controller $\dot{u} = -H^T(y - y^{\text{set}})$ where $H := -CA^{-1}B$ (assuming that A is Hurwitz). ■

While Example 1.1 captures some of the main conceptual ideas behind steering physical plants to an optimal steady state without exogenous inputs, it is not fully representative of some of today's challenging applications.

For instance, instead of simply optimizing an objective as a function of the plant output, the cost might also depend on the plant inputs. More importantly, however, in real-world applications constraints on the plant inputs and outputs often take center stage. Hence, the general optimization problem that one wishes to solve in closed loop can often be expressed as

$$\begin{aligned} \text{minimize} \quad & \Phi(u, y) \\ \text{subject to} \quad & y = h(u, d) \\ & u \in \mathcal{U} \\ & (u, y) \in \mathcal{X}\,, \end{aligned}$$

where h is the steady-state input-to-output map of a plant (as in Example 1.1) and d is a disturbance, albeit not necessarily additive. Further, $\mathcal{U}$ denotes a set of admissible inputs, and $\mathcal{X}$ is a set of additional engineering constraints on inputs and outputs.

One application where the full generality of (1.4) is required concerns the optimal real-time operation of power systems which will be discussed

in more detail in Part III of this thesis:

Example 1.2. The task of operating a power grid safely and efficiently may be cast as an *AC optimal power flow* (ACOPF) problem [106, 107, 134] which is given, in a simplified and stylized version, as

$$\begin{aligned}
\underset{v,\theta,p^{\mathrm{G}},q^{\mathrm{G}}}{\text{minimize}} \quad & \sum_{l=1}^{N} \Phi_l(p_l^{\mathrm{G}}, q_l^{\mathrm{G}}) && \text{(1.5a)}\\
\text{subject to} \quad & \forall l = 1, \dots, N: && \text{(1.5b)}\\
& p_l^{\mathrm{G}} - p_l^{\mathrm{L}} = v_l \sum_{k=1}^{N} v_k \left(\mathrm{g}_{lk} \cos(\theta_l - \theta_k) + \mathrm{b}_{lk} \sin(\theta_l - \theta_k)\right) && \text{(1.5c)}\\
& q_l^{\mathrm{G}} - q_l^{\mathrm{L}} = v_l \sum_{k=1}^{N} v_k \left(\mathrm{g}_{lk} \sin(\theta_l - \theta_k) - \mathrm{b}_{lk} \cos(\theta_l - \theta_k)\right) && \text{(1.5d)}\\
& \underline{p}_l \le p_l^{\mathrm{G}} \le \overline{p}_l && \text{(1.5e)}\\
& \underline{q}_l \le q_l^{\mathrm{G}} \le \overline{q}_l && \text{(1.5f)}\\
& \underline{v}_l \le v_l \le \overline{v}_l && \text{(1.5g)}\\
& (\mathrm{g}_{lk}^2 + \mathrm{b}_{lk}^2)\left(v_k^2 + v_l^2 - 4 v_k v_l \cos(\theta_k - \theta_l)\right) \le \bar{i}_{kl}^2\,, && \text{(1.5h)}
\end{aligned}$$

where $v, \theta, p^{\mathrm{G}}, q^{\mathrm{G}}, p^{\mathrm{L}}, q^{\mathrm{L}} \in \mathbb{R}^N$ denote voltage magnitude, voltage angles, real and reactive power generation and demand, respectively, at each of the N buses of the grid. The generation at each bus l incurs a cost $\Phi_l(p_l^{\mathrm{G}}, q_l^{\mathrm{G}})$ and the load $p^{\mathrm{L}}, q^{\mathrm{L}}$ is assumed to be fixed.

The physical flow of power at steady state is governed by the *AC power flow* (ACPF) equations (1.5c–d) [181, 105] where g_{lk} and b_{lk} denote the so-called *conductance* and *susceptance* of a transmission line connecting buses k and l. We adopt the convention that $\mathrm{g}_{lk} = \mathrm{b}_{lk} = 0$ if l and k are not connected by a line. Furthermore, (1.5e–g) denote simple constraints on power generation and voltages at each bus, and (1.5h) limits the current that flows through the line from k to ℓ.

Even in an offline setting, ACOPF problems like (1.5) are computationally demanding, because of the nonlinear power flow equations which

render the entire problem non-convex (although convex relaxations of the problem can yield global optimality certificates [181, 169]).

In an online setting, the purpose of (1.5) is to find adjustments to the power flow and power injections, given the actual realization of disturbances (as opposed to predictions used in offline *day/hour-ahead* calculations). Moreover, (1.5) can also be considered when reacting to unforeseen contingencies such as line outages (which modify g_{kl} and b_{kl}) or generator outages (which modify $\overline{p}_l, \overline{q}_l$).

In order to express (1.5) as a problem of the form (1.4), we need to identify inputs, outputs, and disturbances. For this purpose, for each generation unit, the real power output and either the reactive power or voltage magnitude at are assumed to be controllable (more precisely, one distinguishes between so-called *PQ*- and *PV*-*buses*). Further, the loads $p^{\mathrm{L}}, q^{\mathrm{L}}$ are considered be a disturbance in (1.5). All remaining quantities are treated as outputs (see [Ha14] for a more detailed discussion including the role of the *slack bus* in numerical simulations).

Consequently, (1.5) can be brought into the form (1.4) where the constraints (1.5e–h) are assigned to either $\mathcal{U}$ or $\mathcal{X}$ according to whether they apply only to controllable variables or not. Under normal operating conditions, the local existence and differentiability of the steady-state map h is guaranteed by the implicit function theorem. ■

We will see that solving the ACOPF problem in Example 1.2 in closed loop combines several aspects that require a new approach to feedback optimization that satisfied the following criteria:

(i) A large number of physical and engineering constraints of different nature have to be satisfied, some at all times (e.g., limits on generation) and others only asymptotically (e.g., thermal limits on line flows).

(ii) In general, only an approximate model in the form of the steady state AC power flow equations is available. A dynamical model for a large-scale power system is in practice not available, as parts of the system are owned by different stakeholders, models are proprietary, and the operating conditions (e.g., generation units

online) are time-varying. However, the power system dynamics may be assumed to be asymptotically stable.

(iii) The nonlinear nature, especially under critical operating conditions, of the power flow equations call for methods that work in the absence of convexity and for ill-conditioned problems.

These requirements lead us to consider *online feedback-based optimization* as the main topic of this thesis. This recent field builds upon the basic principle illustrated in Example 1.1 of interconnecting optimization algorithms with a physical system. In particular, for more general problems of the form (1.4) (such as the ACOPF problem) a variety of constraint enforcement mechanisms have been developed and different implementation and robustness aspects have been studied, yet many problems remain unsolved new algorithm designs keep emerging.

The remainder of this introductory chapter is structured as follows: In Section 1.1 we discuss several existing approaches, all of which can be considered *closed-loop optimization* techniques, each of them with specific use cases, advantages, and disadvantages. Because we aim at interconnecting off-the-shelf optimization algorithms with physical systems, in Section 1.2 we review recent advances in the study of numerical optimization algorithm as dynamical systems. This allows us in Section 1.3 to review recent work on constrained online feedback-based optimization and to illustrate different control designs which will be studied more rigorously in the remainder of the thesis. After this preview, Section 1.4 describes the structure of the remaining chapters.

1.1 Existing Approaches for Optimization in Closed Loop

The approaches to closed-loop optimization that we present in the following have emerged independently from each other and follow different philosophies. Each of methods is particularly suited for specific types of problems, but none comes without drawbacks.

(i) *Extremum seeking* puts an emphasis on being completely *model-free* and instead of probing the system with the help of a perturbation signal. However, the approach is limited to systems with low-dimensional inputs and methods for enforcing constraints are limited.

(ii) *Real-time iteration* schemes aim at solving classical receding horizon problems with limited resources. The focus is the stabilization of a plant under state and input constraints. For this purpose, a full model of the plant dynamics is generally required and computational burden scales with the planning horizon.

(iii) The primary merit of *modifier adaptation* is to mitigate the effects of model uncertainty when solving successive optimization problems. By itself, this method does not reduce the computational requirements compared to feedforward optimization. Furthermore, an auxiliary method to estimate input-output sensitivities is required.

(iv) Algorithms for *time-varying optimization* are specifically designed to track the optimizer of a non-stationary optimization problem and provide quantifiable performance guarantees. However, the need for a unique optimizer to track confines these schemes mostly to convex optimization problems.

(v) The control designs in the context of *network utility maximization* focus on a distributed implementation, but require a particular problem structure, and do not easily generalize to non-convex problems.

1.1.1 Extremum Seeking

Arguably, one of the oldest control methods to steer a plant to an extremum of a function rather than tracking a set-point is *extremum seeking* (ES) (see [238, 11] for historical accounts). Its popularity rose in the 1950's and 60's as part of adaptive control and later regained

momentum with [239, 152] which, for the first time, established rigorous stability guarantees.

The main idea behind ES is to inject a *dither signal* to locally explore the objective function and "learn" its gradient. This dither signal is generally sinusoidal, but other perturbations have been proposed [246]. Consequently, the objective can be optimized without recourse to any model information about the plant (and objective) and without any computation aside from the addition and multiplication of the dither signal. The following example illustrates this fact.

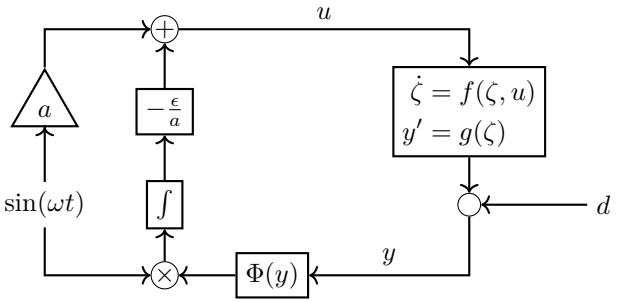

Figure 1.2: Simple Extremum Seeking to minimize $\Phi(h(u))$

Example 1.3. Consider the same setup as in Example 1.1. Namely, we wish to minimize $\Phi(y)$ where y is the output of a plant of the form (1.1) with a steady-state input-to-output map $y = h(u) + d$. We assume that the plant is single-input-single-output.

The reduced ES dynamics of the system in Figure 1.2 (i.e., replacing the fast plant dynamics by the steady-state map $y = h(u) + d$ as in Example 1.1) take the form

$$\dot{u} = f(u, t) := \Phi\left(h\left(-\tfrac{\epsilon}{a}u + a\sin(\omega t)\right) + d\right)\sin(\omega(t))\,. \tag{1.6}$$

The *averaged* dynamics are obtained by integrating $f(u, t)$ from 0 to $T = \frac{2\pi}{\omega}$ which, by using the Taylor expansion of $\Phi\left(h(-\frac{\epsilon}{a}u + a\sin(\omega t)) + d\right)$, yields

$$\tfrac{1}{T}\int_0^T \Phi\left(h\left(-\tfrac{\epsilon}{a}u + a\sin(\omega t)\right) + d\right)\sin(\omega(t)) \approx -\tfrac{\epsilon}{2}\nabla h(u)\nabla\Phi(h(u))\,.$$

Thus, the ES scheme approximates the gradient flow (1.2) from Example 1.1. However, in contrast to (1.1), ES merely requires measurements of $\Phi(y)$ and neither an estimate of ∇h nor of $\nabla\Phi$.

Moreover, ES systems generally evolve on three different timescales: the plant dynamics (which have been ignored for this example), the frequency range of the probing signal, and the slow averaged optimization dynamics. ■

Classically, averaging theory and singular perturbation analysis (for dynamic plants) are used to render the insights from Example 1.3 rigorous [239, 122, 152]. More recently, ES schemes have also been studied with the help of Lie bracket approximations which offer an alternative perspective [120, 86, 87].

While most works considered only finding extrema (i.e. minima or maxima) of unconstrained problems, constraints have been incorporated by submanifold constraints [88], barrier functions [74], and saddle-point formulations [89], and ES has been studied for Nash-equilibrium seeking [109, 231]. Further, ES has been studied for stochastic [233, 65] and discrete-time setups [97, 110, 232], and hybrid extensions have been proposed [202, 204].

ES has been applied in the automotive sector [145], process engineering [121], formation flight and obstacle avoidance [183, 36] and others. More modern applications concern problems in renewable energy such as maximum power point tracking in photovoltaic [116] or wind energy systems [115, 151], and Volt-VAR control in power systems [12].

Despite strong theoretical guarantees and being model-free, ES has been confined to relatively low-dimensional, mostly unconstrained, systems. This is due to the fact that plants with multidimensional input require probing signals at different, carefully chosen, frequencies that do not interfere with each other.

1.1.2 Real-Time Iterations

Historically speaking, MPC provides a method to control and stabilize a plant that is subject to input and state constraints. This is achieved

by numerically solving an optimal control problem with a finite receding horizon at every sampling time, but implementing only the first input of the computed optimal policy before solving the next problem with a shifted horizon and based on an updated state measurement.

The high computational requirements have long limited the application of MPC to relatively slow and low-dimensional plants in process engineering. For standard (linear) MPC, this issue has led to *explicit* MPC [7, 26] which exploits multi-parametric optimization [249] to solve the receding horizon problem ahead of time and implement the controller as a simple lookup table.

More interesting from our perspective are *real-time iterations* (RTIs) for *nonlinear MPC* [80, 40]. These methods have emerged as an approximation of *multiple shooting methods* [78] and have been proposed for various applications in process engineering [82], robotics [78], and for airborne kites [77].

The main idea of RTIs is to solve the optimal control problem only approximately at every iteration by performing only a single iteration of the underlying optimization algorithm (which is usually an *sequential quadratic programming* (SQP) scheme [206, 266, 191]). The first input of the approximate control policy is implemented and the optimization problem for the next sampling period is warm-started at a shifted version of the previous (approximate) solution. Example 1.4 below illustrates this procedure.

The underlying idea of RTIs is that the approximation error committed by performing only a single optimization iteration is offset by savings in computation time. In particular, because the receding horizon problem is solved more often, it changes less between samples. This feature allows one to prove stability and convergence of RTI schemes [81, 77, 79].

However, although they interleave optimization iterations with physical dynamics, RTIs have been developed for stabilization and require an exogenous set-point as well as a dynamic model of the plant. This is particularly reflected in the assumptions on the state cost function, which are, roughly speaking, required to be quadratic functions centered at the origin (see also Remark 1.1 further below on economic MPC).

Example 1.4. We consider a discrete-time plant $\zeta^+ = f(\zeta, u)$ for which the origin is a steady state. For simplicity, we do not consider any input or state constraints, although RTIs can incorporate them naturally.

Consider the following receding horizon problem at time l

$$\begin{aligned} \underset{r_{1:K-1}, s_{1:K}}{\text{minimize}} \quad & \sum_{k=1}^{K-1} \left[\begin{smallmatrix} s_k \\ r_k \end{smallmatrix}\right]^T Q \left[\begin{smallmatrix} s_k \\ r_k \end{smallmatrix}\right] + s_K^T R s_K \\ \text{subject to} \quad & 0 = s_1 - \overline{\zeta}_l \\ & 0 = s_{k+1} - f(s_k, r_k) \qquad \forall k \in \{1, \dots, K-1\}\,, \end{aligned} \tag{1.7}$$

where K denotes the horizon length, Q, R are positive definite stage and terminal cost matrices, and $\overline{\zeta}_l$ denotes the measured plant state at time l. Let $(\hat{r}^l, \hat{s}^l)$ denote the solution of (1.7) for the sampling instant l. Then the feedback control law at l is given by $u[l] = \hat{r}_1^l$. In other words, upon solving (1.7), the first control of the optimal policy is implemented at l. This is the key mechanism behind standard MPC.

RTI schemes approximate the solution of (1.7) by performing only a single iteration of an SQP method. Namely, let $z = (s, r, \lambda)$ and consider the Lagrangian at time l defined as

$$\begin{aligned} L^l(z) := \sum_{k=1}^{K-1} \left[\begin{smallmatrix} s_k \\ r_k \end{smallmatrix}\right]^T Q \left[\begin{smallmatrix} s_k \\ r_k \end{smallmatrix}\right] + s_K^T R s_K \\ + \lambda_1(s_1 - \overline{\zeta}_l) + \sum_{k=1}^{K-1} \lambda_k \left(s_{k+1} - f(s_k, r_k)\right)\,. \end{aligned}$$

An SQP iteration then takes the form

$$z^+ = z + \Delta z \qquad \text{where } \Delta z \text{ solves} \quad \nabla L^l(z) + \nabla_{zz}^2 L^l(z) \Delta z = 0\,, \tag{1.8}$$

where, in practice, the inverse of the Hessian $\nabla_{zz}^2 L^l$ is often approximated.

Crucially, the receding horizon problem (1.7) at the next sample $l+1$ is warm-started with the shifted approximate of the previous sample. This procedure leads, under additional assumptions, to local stability of the scheme. ■

Example 1.1 elucidates several differences with respect to the optimal steady-state control problem (1.4) we wish to solve: First, the purpose of RTIs is primarily to drive a plant to the steady state at the origin, not seeking out a steady-state with minimal cost (see Remark 1.1 below). Further, a full dynamic model f of the plant dynamics is required, and finally, the computational burden of solving the SQP iteration scales not only with the system dimension but also with the prediction horizon. Of course, compared to Example 1.1, while stabilizing the system RTI also seeks to minimize a quadratic running stage cost.

Remark 1.1. Traditionally, linear and nonlinear MPC have been considered with the goal of stabilizing a plant and to track a precomputed set-point or trajectory. The modern variation of *economic MPC* [93, 91, 209] studies the effects of incorporating an economic objective directly and thus not requiring an exogenous setpoint. This idea is very much in line with the topic of this thesis.

Economic MPC, however, still requires the solution of an optimal control problem at every iteration. Thus, it is computationally very expensive and does not seem to be amenable to an RTI implementation. Moreover, the optimal solution is not a-priori guaranteed to be a steady-state of the plant. This makes the stability analysis of economic MPC more involved and, in particular, still requires a full model of the plant dynamics. ■

1.1.3 Modifier Adaptation

In the context of *real-time optimization* in process engineering, the notion of *measurement-based optimization* [103, 56] has been used to collect several approaches towards mitigating the effects of model uncertainty in repetitive optimization applications. In the following we focus on so-called *modifier adaptation* (MA) methods [172, 173, 104, 112].

Given a model of a physical system, assume that we can solve an optimal steady-state problem like (1.4) numerically. When implementing this solution by setting the appropriate inputs of the plant to their precomputed optimal set-points, the mismatch between the model estimate (used for computing an optimal state) and the actual plant will invari-

ably lead to an system state that is suboptimal, and possibly violating constraints.

If the optimization of the optimal plant state is performed repeatedly, and at each step the solution is implemented on the physical system, MA provides a method to steadily reduce the discrepancy between model-based solution and physical plant by *modifying* the optimal steady-state problem at every iteration by incorporating plant measurements from the previous iteration. MA does not directly "learn" or identify a better model of the plant. Instead, MA corrects only the optimization problem by adding adaption terms to the cost and constraint functions. The following simple example illustrates this basic concept.

Example 1.5. Consider the same setup as in Example 1.1. Namely, we wish to minimize the function $\tilde{\Phi}(u) := \Phi(h(u) + d)$ where $y = h(u) + d$ is the steady-state input-to output map of a plant with fast-decaying dynamics. However, only an approximate model $\tilde{h}$ of h and an estimate $\tilde{d}$ of d are available.

Therefore, instead of minimizing $\tilde{\Phi}$, we repeatedly solve the problem

$$\underset{u}{\text{minimize}} \quad \tilde{\Phi}(\tilde{h}(u) + \tilde{d}) + \lambda_k^T u\,, \tag{1.9}$$

where λ_k is a *modifier* at iteration k that is adapted at every iteration based on the outcome of the previous iteration $u_k^\star$. In particular, λ is updated according to

$$\lambda_{k+1} = \nabla\tilde{\Phi}(u_k^\star) - \nabla\tilde{\Phi}'(u_k^\star)\,,$$

where $\nabla\tilde{\Phi}(u_k^\star)$ needs to be estimated and $\nabla\tilde{\Phi}'(u_k^\star) := \nabla\left(\Phi(\tilde{h}(u_k^\star) + \tilde{d})\right)$ is model-based. The particular structure of $\tilde{\Phi}$, however, lets us write

$$\nabla\tilde{\Phi}(u_k^\star) = \nabla h(u_k^\star)\nabla\Phi(h(u_k^\star) + d)\,,$$

where $h(u_k^\star) + d$ is the measured output of the plant. Thus, essentially, only $\nabla h(u_k^\star)$ needs to be estimated. If the scheme converges to some $u^\star$, we can easily verify that $u^\star$ is a critical point of $\tilde{\Phi}(u)$. ■

Example 1.5 is simplified to the point that a comparison with Example 1.1 is easily possible. However, MA methods are easily applied

to constrained problems where modifiers on constraints are introduced analogously [93, 69, 119].

Clearly, the tricky part about MA is the estimation of the (true) plant sensitivities $\nabla\tilde{\Phi}(u_k^\star)$. This can be achieved with finite differences [171], but ultimately, restricts the method to fairly low-dimensional setups. Moreover, MA does not reduce the computation burden nor does it aim to reduce the amount of model information required.

1.1.4 Time-Varying Online Optimization

A topic that is central to online optimization, either in open or closed loop, is the study of optimization problems that vary over time. In recent years, this topic has garnered significant interest because of its relevance to many applications in control, robotics, signal processing, and others.

The main focus has been the development of online algorithms that can track the solution of a time-varying optimization problem with tight performance guarantees. One can distinguish two different perspectives:

On the one hand, [163, 35, 123, 137, 269] and others frame time-varying optimization as an *iterative learning problem* in following sense: At every iteration k an agent chooses an action x_k and subsequently a convex function Φ_k is revealed. The agent's goal is to minimize her *regret*, i.e., some measure of accumulated suboptimality. In the simplest case, this may be the *static regret* $\sum_k \Phi_k(x_k) - \min_x \sum_k \Phi_k(x)$ comparing the player's actions against the best "constant" decision, or *dynamic regret* $\sum_k \Phi_k(x_k) - \sum_k \min_x \Phi_k(x)$ which compares the agent's actions against the sequence of minimizers of Φ.

On the other hand, [242, 207, 226, 208] are inspired more by control theory and describe time-varying optimization as a *tracking problem* whereby an optimization algorithm defines an time-varying solution map $t \mapsto x^\star(t)$ that needs to be followed as closely as possible by the online optimization scheme.

Both avenues share common points: First, convexity is generally assumed to guarantee the existence of a unique minimum and, in the case

of strong convexity, a unique minimizer [225]. Exceptions are [Ha10, 242] where results for non-convex optimization problems are presented. Second, to give meaningful performance guarantees, some sort of "bounded variation" in the optimization problem has to be assumed. For this purpose, assumptions about the time-variability and (uniform) boundedness of the problem components have to be made. A common assumption for this purpose is to assume that the rate of change of the optimizers is bounded by a known constant. This assumption is relaxed [Ha10] which shows how, in special cases, the rate of change of the optimizer can be bounded using information about the objective and the constraint functions only.

Depending on the application, the time-varying optimization problem might be unconstrained [201], or constrained to a stationary set [178, 123, 137, 269] or have time-varying constraints [Ha10, 242, 207, 208]. Time-varying constraints have been dealt with using barrier functions [96, 95] or perturbed sweeping processes [Ha10, Ha12, 242].

Roughly speaking, algorithms for time-varying optimization can be divided into *running* algorithms that do not incorporate any information about the evolution of the problem [29, 242, 201] and *predictive* schemes that exploit some knowledge or estimate about the change in the optimization problem [163, 96, 224, 226].

Example 1.6. Consider a time-varying objective function $\Phi(x,t)$ that is continuously differentiable in x. Assuming measurability of $\nabla_x \Phi(x,t)$ in t, we may consider the gradient system

$$\dot{x} = -\nabla_x \Phi(x,t) \tag{1.10}$$

to try to track minimizers of $\Phi(x,t)$ over time.

Assume, for simplicity, that $\Phi(x,t)$ is β-strongly convex in x for all t. Consequently, for every t, there exists a unique minimizer $x^\star(t)$ of $\Phi(x,t)$. Furthermore, assume that $x^\star(t)$ is ℓ-Lipschitz continuous, i.e., $\|x^\star(t') - x^\star(t)\| \leq \ell \|t' - t\|$. Namely, the rate of change of the optimizer is bounded.

Then, the distance between a trajectory $x(t)$ of (1.10) and $x^\star(t)$ sat-

isfies

$$\tfrac{d}{dt}\tfrac{1}{2}\|x(t)-x^\star(t)\|^2 \leq \langle \dot{x}(t), x(t)-x^\star(t)\rangle + \ell\|x(t)-x^\star(t)\|\,.$$

Consequently, $\|x(t)-x^\star(t)\|$ is decreasing as long as $\|x(t)-x^\star(t)\| > \ell/\beta$. It follows from standard invariance arguments that, as $t\to\infty$, $x(t)$ will be ℓ/β-close to $x^\star(t)$. A similar statement holds for the discrete-time case.

If, in addition, Φ is twice continuously differentiable and the time derivative of $\nabla_x\Phi$ is available, then trajectories of

$$\dot{x} = -\tau\nabla_x\Phi(x,t)^T - \tfrac{d}{dt}(\nabla_x\Phi)^T(x,t)$$

are guaranteed to converge to $x^\star(t)$ as $t\to\infty$. To see this, one may consider the time-varying Lyapunov function $W(x,t) := \frac{1}{2}\|\nabla_x f(x,t)\|^2$ whose time-derivative is negative for all t and all $x(t)\neq x^\star(t)$. ■

1.1.5 Distributed Network Utility Maximization

Arguably one of the largest man-made distributed feedback systems is formed by the congestion management mechanisms at the heart the internet [167]. These protocols manage the allocation of link capacities to individual connections in highly dynamic environments, for heterogeneous agents, and subject to real-world imperfections such as delays [240, 193, 261, 256].

Deterministic, continuous-time flow models have been particularly successful in explaining these protocols, by providing an optimization-based perspective in terms network utility maximization [168, 218, 166, 142], and enabling improved designs [260, 265, 259]. The following example illustrates the basic modeling approach.

Example 1.7. Given a communication network, let a set of N sources share M links. The set of links used by each source are collected in the routing matrix $R\in\mathbb{R}^{N\times M}$ in the sense that $R_{ij}=1$ if link j is used by node i and $R_{ij}=0$ otherwise.

Each link j in the network has an associated congestion measure μ_j (also referred to as *price*) which may be derived from queuing delays

or packet loss. Each link j has a finite capacity c_j. Each source i has a controllable *source rate* x_i, e.g., in the form of its *window size*. The different source rates define an aggregate flow for each line given by $y = R_j x$ where R_j is the row of R corresponding to link j. Conversely, sources are assumed to have access to the respective *aggregate price* $R^T \mu$. This is, for instance, satisfied by the fact that the aggregated probability of a packet loss (which can be estimated by the source) can be approximated by $\prod_{j \in L_i}(1 - p_j) \approx 1 - \sum_j p_j$ where L_i denotes the links used by source i and p_j is the probability of a packet being dropped on link j.

The simplest *link algorithm* to compute μ_j is given by the projected gradient law $\dot{\mu}_j = \Pi\left[y_j - c_j\right](\mu_j)$ where, for now, $\Pi\left[y_j - c_j\right](\mu_j) = y_j - c_j$ if $\mu_j = 0$ and $y_j > c_j$, and $\Pi\left[y_j - c_j\right] = 0$ otherwise.

Source controllers can often be modeled as $\dot{x}_i = -\nabla \Phi_i(x_i) - p_i y_i$ where $\Phi_i(x_i)$ is a specific type of cost function, reverse-engineered and depending on the particular protocol. (For consistency with the rest of the thesis, we consider the minimization of a cost, rather than the maximization of a utility function).

Hence, congestion control mechanisms interconnect source and link controllers into a feedback loop as illustrated in Figure 1.3 where $\Phi(x) := \sum_i \Phi_i(x_i)$. Importantly, each controller is fully distributed and requires only locally available information.

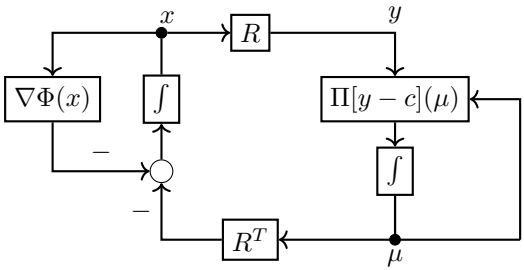

Figure 1.3: Simple Network Utility Maximization

Given the optimization problem

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad Rx \leq c\,, \tag{1.11}$$

note that the closed-loop dynamics

$$\begin{aligned} \dot{x} &= -\nabla_x L(x,\mu) = -\nabla\Phi(x) - R^T\mu \\ \dot{\mu} &= \Pi\left[\nabla_\mu L(x,\mu)\right](\mu) = \Pi\left[Rx - c\right](\mu) \end{aligned} \tag{1.12}$$

form a (projected) saddle-point flow of the Lagrangian $L(x,\mu) = \Phi(x) + \mu^T(Rx - c)$. Since Φ is convex, trajectories of (1.12) can be shown to converge to saddle-points of L, all of which correspond to optimizers of (1.11). ■

From a control and optimization perspective, source and link controllers form a closed feedback loop that implicitly tracks the solution of the underlying utility maximization problem [258]. From the perspective of closed-loop optimization, one can argue that the link controllers define a plant which implements the dual dynamics that have to be complemented by controllers that form the primal dynamics. This interpretation, however, does not quite fit the optimal steady-state control problem (1.4).

Another possibility, more in line with the optimal input-output problem (1.4), is to assume that $y = Rx$ forms a simple algebraic plant and that link and source controllers together form a two-stage dynamic controller where μ is an internal variable. This interpretation, however, clearly ignores much of the particular structure which allows for a very natural distributed implementation.

This same problem structure can also be observed more generally in optimization problems over networks where constraints are either local to each node or take the form of linear balancing constraints for which there exists an effective way to compute a price-signal.

Another important application where this structure can be identified and exploited is frequency control of power systems. Namely, in AC power grids, the deviation of the system frequency from its nominal value can be seen as a pricing signal that is available to all grid-connected components and indicates a mismatch between power generation and consumption. This and other insights have led to extensive work on

optimal frequency control schemes based on distributed averaging-based control [85, 10] and primal-dual methods [165, 267, 141] which identify already existing control mechanism as optimization dynamics and propose extensions or improved designs (see [179] for other pertinent references).

Although, this type of feedback-based optimization naturally admits a distributed implementation, it requires a special structure with either the physical system or an additional integrator implementing the dual dynamics. Furthermore, as we will see in Section 1.2.1, the primal-dual dynamics that form the closed-loop behavior do not easily generalize to non-convex problems with the same global stability and convergence guarantees as for convex setups.

1.2 Optimization Algorithms as Dynamical Systems

In this section, we survey advances in the study of optimization algorithms (without feedback control aspect), because these insights are important for closed-loop optimization since they may inspire new control designs and provide important theoretical guarantees.

In recent years, renewed attention has been paid to the fact that many numerical optimization algorithms can be in interpreted as dynamical systems (equipped with a stopping criterion). Hence, two questions have been asked: First, can dynamical systems inspire new optimization algorithms? And second, can existing algorithms be better understood using tools from dynamical systems analysis?

The former question has been supported by the fact that optimization dynamics such as gradient flows are well-defined and well-behaved on abstract spaces including smooth manifolds. Consequently, various works such as [38, 94, 51] culminating in [126] have shown how computational problems such as eigenvalue computations, sorting lists, linear programming, etc. can be framed as optimization problems on appropriate matrix manifolds and that numerical algorithms can be recovered

from discretizing the corresponding gradient flow. These insights have, in turn, led to the field of *optimization on manifolds* [3] which studies the design of numerical algorithms to minimize cost functions defined on Riemannian manifolds embedded in a vector space.

The latter question whether new insights for optimization can be gained using tools from control and dynamical system theory has been addressed in various ways. In the following we focus on saddle-point flows and accelerated first-order methods, both of which have been under intense investigation over the last years.

1.2.1 Primal-Dual Saddle-Point Dynamics

Simply speaking, under weak technical assumptions, solutions of a constrained optimization problem are saddle-points of the associated *Lagrangian*. For this reason, dynamical systems that seek out saddle-points rather than extrema of a function are of particular interest for constrained optimization.

Historically, *saddle-point flows* have primarily been studied in the context of nonlinear circuit analysis [228, 50], but their potential for optimization has been observed even before that [13, 149]. Although saddle-point flows have been studied throughout the years[39, 255], they have recently become a topic of intense study due their importance in the context of distributed network optimization (Section 1.1.5).

Generally, if a saddle function $L(x, \mu)$ is convex in x for every μ, concave in μ for all x, and either strictly convex in x or strictly concave in μ, then trajectories of the system

$$\dot{x} = -\nabla_x L(x, \mu)^T \qquad \dot{\mu} = \nabla_\mu L(x, \mu)^T \tag{1.13}$$

converge to a saddle-point of L, i.e., a point $(x^\star, \mu^\star)$ such that $L(x, \mu^\star) \geq L(x^\star, \mu^\star) \geq L(x^\star, \mu)$ for all x and all μ. In particular, (1.13) consists of a gradient descent in the *primal* variables x and a gradient ascent in the *dual* variables μ.

Differentiability of L is not generally required. In fact, (1.13) can be generalized to include projections on both x and/or μ [Ha7, 58–60] and,

even more generally, (1.13) can be defined in terms of subgradients if L is not differentiable [117, 255].

Convergence proofs for (1.13) and its generalization usually exploit monotonicity of the vector field $[-\nabla_x L(x,\mu)\ \nabla_\mu L(x,\mu)]$ and apply an invariance argument. However, without either strict convexity in x or strict concavity in μ, convergence to equilibria is not guaranteed and, instead, oscillations may occur [132] which can, however, be eliminated with a simple augmentation term [Ha7].

The following example illustrates the construction of a projected saddle-point flow to solve a nonlinear optimization problem with different types of constraints.

Example 1.8. Consider the problem

$$\begin{aligned} \text{minimize} \quad & \Phi(x) \\ \text{subject to} \quad & x \in \mathcal{X} \\ & g(x) \leq 0\,, \end{aligned} \tag{1.14}$$

where $\Phi : \mathbb{R}^n \to \mathbb{R}$ and $g : \mathbb{R}^n \to \mathbb{R}^m$ are convex and continuously differentiable, $\mathcal{X} \subset \mathbb{R}^n$ is non-empty and closed convex. Further, let the *partial* Lagrangian $L : \mathcal{X} \times \mathbb{R}^m_{\geq 0} \to \mathbb{R}$ of (1.14) be defined as

$$L(x,\mu) := \Phi(x) + \mu^T g(x)\,.$$

We define the projected saddle-point flow

$$\begin{aligned} \dot{x} &= \Pi_{\mathcal{X}} \left[-\nabla_x L(x,\mu)^T\right](x) = \Pi_{\mathcal{X}} \left[-\nabla\Phi(x)^T + \nabla g(x)^T \mu\right](x) \\ \dot{\mu} &= \Pi_{\mathbb{R}^m_{\geq 0}} \left[\nabla_\mu L(x,\mu)^T\right](\mu) = \Pi_{\mathbb{R}^m_{\geq 0}} \left[g(x)\right](\mu) \end{aligned} \tag{1.15}$$

where $\Pi_{\mathcal{X}}[w](x)$ and $\Pi_{\mathbb{R}^m_{\geq 0}}[w](x)$ project w and v onto the tangent cone of $\mathcal{X}$ and on the non-negative orthant $\mathbb{R}^m_{\geq 0}$ at x and μ, respectively. Consequently, trajectories of (1.15) cannot leave $\mathcal{X} \times \mathbb{R}^m_{\geq 0}$. This class of *projected dynamical systems* is well-studied, even for non-convex and non-Euclidean settings [Ha2, 125, 18, 68] and will be one of the main topics of this thesis.

Under adequate constraint qualifications and if Φ is strictly convex, trajectories of (1.15) are guaranteed to converge to a KKT point (and thereby to a global optimizer) of (1.14) [117].

If Φ is non-strictly convex, then trajectories of (1.15) converge to the optimizer of (1.15) if, instead of L, an augmented Lagrangian of the form

$$L^a(x,\mu) := \Phi(x) + \mu^T g(x) + \gamma(x)$$

is used, where γ is a convex penalty function for the set $\{x \,|\, g(x) \leq 0\}$. Namely, $\gamma(x) = 0$ for all $x \in \mathcal{X}$ and $\gamma(x) > 0$ for $x \notin \mathcal{X}$. This is, for example, the case for $\gamma(x) := \frac{\rho}{2}\|\max\{g(x), 0\}\|^2$ [Ha7].

If Φ or g are non-convex, convergence is not guaranteed unless an additional regularization of the form

$$L^b(x,\mu) := \Phi(x) + \mu^T g(x) + \gamma(x) + \frac{\rho'}{2}\|\mu\|^2$$

is added [29, 71, 242, 243]. However, this modification changes the equilibria of (1.15) which then do not necessarily coincide with KKT points of (1.14) anymore. ■

Remark 1.2. Augmented saddle-point flows exhibit strong parallels with classical PI-controllers: Consider the problem

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad y^{\text{set}} = Hx + d\,,$$

where y^{set} is set-point and H and d are a matrix and vector of appropriate size.

Assume that a plant exhibits the nonlinear dynamics

$$\dot{x} = -\nabla\Phi(x)^T + H^T u \qquad y = Hx + d\,.$$

Applying a proportional controller $u = -\rho(y - y^{\text{set}})$ realizes a closed-loop system of the form

$$\dot{x} = -\nabla\Phi(x)^T - \rho H^T(H + d - y^{\text{set}})$$

which is the gradient flow of the penalty-augmented cost $\Phi(x) + \frac{\rho}{2}\|Hx + d - y^{\text{set}}\|^2$.

An integral controller $\dot{u} = \eta(y - y^{\text{set}})$, on the other hand, leads to a closed-loop behavior that implements the saddle-point flow

$$\begin{aligned}\dot{x} &= -\nabla\Phi(x)^T - \rho H^T u\\ \dot{u} &= \eta(Hx + d - y^{\text{set}})\end{aligned}$$

for the Lagrangian $L(x,u) = \Phi(x) + u^T(Hx + d - y^{\text{set}})$ with a dual gain η.

Finally, combining both controllers into a PI-controller, leads to the system

$$\begin{aligned}\dot{x} &= -\nabla\Phi(x)^T - \rho H^T(u + (H + d - y^{\text{set}}))\\ \dot{u} &= \eta(Hx + d - y^{\text{set}})\end{aligned}$$

which is a saddle-point flow for the augmented Lagrangian

$$L^a(x,u) := \Phi(x) + u^T(Hx + d - y^{\text{set}}) + \tfrac{\rho}{2}\|Hx + d - y^{\text{set}}\|^2 . \qquad \blacksquare$$

All in all, saddle-point flows have proven to be very versatile and particularly useful for distributed optimization where agents share certain constraints. However, convergence and stability results for non-convex problems are not generally available. Moreover, even for convex problems, tuning can be difficult, especially for nonlinear problems because suboptimal choices of the dual gain η and penalty parameter ρ can lead to severely under- or over-damped transients, a problem that only gets more challenging for high-dimensional and ill-conditioned problems.

1.2.2 Accelerated First-Order Methods

Even though accelerated algorithms have yet to be successfully implemented as feedback controllers, we quickly survey recent results and discuss the challenges in deploying them in closed loop with a physical plant.

Ever since its inception, Nesterov's accelerated gradient method [190, 189] has remained enigmatic. Its convergence rate is provably optimal for a first-order method and is matched only by the heavy-ball method [114, 200] and the more recent "triple momentum" method [253], which was designed using tools from robust control [164]. Nevertheless, the intuition behind the phenomenon of acceleration remains elusive.

In [237] it was first noted that a continuous-time limit of Nesterov's method is given by the second-order ODE

$$\ddot{x} + \tfrac{3}{t}\dot{x} + \nabla\Phi(x) = 0\,, \tag{1.16}$$

where $\frac{3}{t}\dot{x}$ is commonly seen as time-varying damping term [262, 263, 15, 16].

However, for an application in (continuous-time) closed-loop optimization, (1.16) remains ill-suited because (1.16) is non-autonomous and thus not adequate as a running algorithm. Moreover, in [203] it was noted that (1.16) is not robust from control-theoretic perspective since the optimizer of Φ is not *uniformly* attractive and [Ha3] provides a counter-example showing the interconnection of (1.16) with a dynamical system is, in general, not asymptotically stable.

A case of Nesterov's method for μ-strongly convex objectives takes the form

$$x_{k+2} = x_{k+1} + \beta(x_{k+1} - x_k) - \alpha\nabla\Phi(x_{k+1} + \beta(x_{k+1} - x_k)), \quad (1.17)$$

where α, β are carefully chosen parameters based on μ and the Lipschitz constant of $\nabla\Phi$. Based on this special case, an alternative time-invariant ODE was established in [185]. However, despite its time-invariant nature, (1.17) is known for this lack of robustness against inexact gradient evaluations [177, 70, 176, 75].

Another topic of recent study are *adaptive resets* of the momentum in accelerated methods [192] which have been observed to increase the robustness towards uncertainty in the algorithm parameters. Attempts at making this heuristic more rigorous have led to new types of algorithms in the form of hybrid dynamical systems [135, 148, 203, 247, 147].

In summary, accelerated methods have received a lot attention in recent years and continue to inspire new work. However, their application to closed-loop optimization remains difficult because of their general lack of robustness.

1.3 Constrained Feedback Optimization

We now turn our attention to the main methods that have in recent years been developed to steer plants, such as electricity grids, to an optimal steady state subject to physical and engineering constraints.

Early designs for this type of online feedback-based optimization can be traced back to [53, 140]. Both works consider steering a dynamical plant towards the solution of a constrained convex optimization problem and resort to saddle-point formulations to enforce constraints. While [140] proposed control designs based on so-called "complementarity integrators", [53] considers a smooth saddle-flow variation and applies backstepping to design the feedback controller. Compared to the general problem (1.4), [53, 140] study a simple optimization problem in the output variable only, i.e., without input constraints and a cost function that does not depend on u.

More recent work has studied more complex (including non-convex) problem setups, different control designs, and has explored alternative approaches to establish stability and convergence properties. In the following, we first classify the various control papers by their problem setup and then provide examples of the particular methods studied in this thesis. Most of these works have been developed in the context of real-time power systems optimization. However, we limit ourselves to a discussion of the control-theoretic aspects and do not discuss application-focused works such as [144, 174, Ha14] and pertinent references in [179].

1.3.1 Convex Problem Setups

Assuming that (1.4) is a convex problem implies that the steady-state input-to-output map h is affine. Within this problem class, [187] studies an (otherwise) unconstrained setup whereas [159, 227, 57, 66, 158] also incorporate additional linear equality constraints on plant outputs. In [29] more general nonlinear convex inequality constraints on the outputs are used and [29, 57] also incorporate general convex constraints on plant inputs. However, it should be noted that the focus of [159, 227, 66, 158, 187] is on the interconnection with dynamical plants (see forthcoming Section 1.3.3) and not on constraint handling.

Time-varying problems are studied in [29, 66] which also provide bounds on the tracking performance under various assumptions (see also Section 1.1.4). Furthermore, [29] also models the effect of measurement noise.

Most of these papers consider a continuous-time setup [159, 227, 57, 67, 158, 187] while only [29] works in discrete time.

The control designs are, to a large degree, based on saddle-point algorithms (Section 1.2.1). For instance, [57] implements a double projection saddle-point flow (see Example 1.8), [66] is based on the proximal augmented Lagrangian method [76], and [29] implements a saddle-point flow with regularization in the dual variables.

The architectures in [159, 158, 187] are more difficult to categorize since they also include the design of estimators and stabilizing controllers.

1.3.2 Non-Convex Problem Setups

If the steady-state feedback optimization problem (1.4) is non-convex, the number of possible control designs for which stability is guaranteed is limited.

The source of non-convexity is often a nonlinear steady-state map or non-convex cost function whereas constraints are often assumed to be convex (especially if some sort of projection is involved). Hence, the continuous-time formulations in [Ha3–Ha6, Ha13, Ha15] and the discrete-time setups in [Ha1, 67, 71, 243, 241, 111] all deal with some form of non-convexity.

In terms of the underlying optimization dynamics, projected gradient schemes are particularly suited because of their convergence guarantees for non-convex problems as we will see in the rest of this thesis. Otherwise, quasi-Newton methods [241] have been considered, although they suffer from the fact that constraints cannot be easily enforced by projection. Yet another algorithm, based on regularized saddle-point flows, was proposed in [71, 243] (and [29] for the convex case). However, the stability conditions for non-convex regularized saddle-point flows are relatively involved.

The papers [71, 243, 241] further establish tracking results for time-varying problem setups. However, the bound in [71] is with respect to a linearly approximated problem and the bound in [241] is with respect to a penalty-relaxed problem.

1.3.3 Interconnection with Dynamic Plants

Most of the previously discussed papers in this section assume complete timescale separation between plant dynamics and optimization process and, consequently, model the plant as an algebraic steady-state map. Notable exceptions to this assumptions are [Ha3, 159, 227, 66, 158, Ha13, 187] which investigate the stability of interconnecting optimization dynamics with plant dynamics.

For this purpose, [Ha3, Ha13] pursue a singular perturbation approach and arrive at easy-to-compute sufficient conditions that guarantee overall stability. In particular, [Ha3] considers general nonlinear (but stable) plant dynamics and investigates a variety of optimization dynamics for convex and non-convex problems including (projected) gradient, saddle-point, and momentum methods. Using similar techniques, [227] provides stability guarantees for general low-gain integral controllers that satisfy an infinitesimal contraction property, but with a discussion of feedback-based optimization applications limited to LTI plant dynamics.

By restricting their attention to LTI plants, [66, 187] use tools from robust control to provide computational stability certificates in the form of LMIs. In particular, [187] studies the joint design of stabilizing control, estimator and optimization dynamics. Also limited to LTI plants, [159, 158] consider an output regulation framework and reduce the control design to a stabilization problem.

1.3.4 Enforcing Input Constraints with Saturation and Anti-Windup

Papers like [Ha3, 29, 57, 71, 111] consider constraints on plant inputs which are, from an optimization perspective, enforced by projection onto a set of feasible inputs. From an control viewpoint, however, one would like to exploit physical input saturation to enforce this type of constraints. For this purpose, [Ha4–Ha6] show rigorously how anti-windup control schemes can be used to implement this type of constraint enforcement. The following example illustrates the concept which will be studied in detail in Chapters 5 and 9.

Example 1.9. Consider the same setup and controller as in Example 1.1. Namely, we want to drive a plant (with fast decaying dynamics and steady-state map $y = h(u) + d$) to an optimal steady state minimizing the cost function $\tilde{\Phi}(u)$.

In addition, we assume that the plant is subject to input saturation that acts like a projection onto a set $\mathcal{U}$ of admissible inputs. In order to mitigate the effects of integrator windup, a simple anti-windup scheme with a tuneable gain ϵ is in place as illustrated in Figure 1.4.

It turns out (which can be easily verified) that equilibrium points of this feedback loop correspond to KKT points of the optimization problem

$$\text{minimize} \quad \Phi(y) \qquad \text{subject to} \quad y = h(u) + d \qquad u \in \mathcal{U}\,.$$

Moreover, it can be shown that, under mild assumptions and high enough ϵ, the interconnection is asymptotically stable and as $\epsilon \to 0$ we recover a projected gradient flow. ■

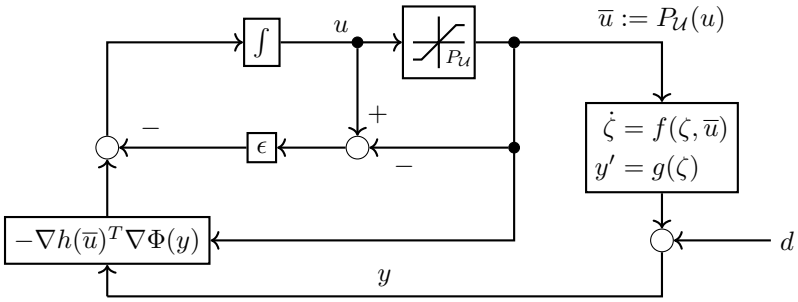

Figure 1.4: Anti-windup Gradient Flow

1.3.5 Enforcing Output Constraints by Linearized Output Projection

The following method has been proposed in [Ha1] and will be discussed in detail in Chapters 6 and 10. The controller combines strong convergence

guarantees with relatively weak assumptions and easy tuning.

Example 1.10. Consider the general steady-state optimization problem (1.4) and assume that $\mathcal{U} := \{u \,|\, Au \leq b\}$ and $\mathcal{X} := \{(x,u) \,|\, C_u u + C_y y \leq d\}$ are polyhedra.

To steer the plant to a solution of (1.4) we consider the feedback control law

$$u^+ = u + \alpha\, \widehat{\sigma}_\alpha(u,y) \quad y = h(u,d)\,, \tag{1.18}$$

where $\alpha > 0$ is a fixed step-size, $y = h(u,d)$ is the measured system output, and $\widehat{\sigma}_\alpha : \mathbb{R}^p \times \mathbb{R}^n \to \mathbb{R}^p$ is defined as

$$\begin{aligned} \widehat{\sigma}_\alpha(u,y) := \arg\min_{w \in \mathbb{R}^p} \; & \|w + G^{-1}(u) H(u,d)^T \nabla \Phi(u,y)^T\|^2_{G(u)} \\ \text{subject to} \quad & A(u+\alpha w) \leq b \\ & C_u u + C_y(y + \alpha \nabla_u h(u,d) w) \leq d\,, \end{aligned} \tag{1.19}$$

where $H(u,d)^T := \begin{bmatrix} \mathbb{I}_p & \nabla_u h(u,d)^T \end{bmatrix}$ and G assigns to every $u \in \mathcal{U}$ a positive definite matrix. While technically $\nabla_u h(u,d)$ depends on the input u and the disturbance d, in practice, it is often possible to estimate $\nabla_u h(u,d)$ from u and the output y without explicitly estimating d. Furthermore, simulations indicate that the scheme is very robust against uncertainty in $\nabla_u h$.

The feedback law (1.18) approximates a projected gradient descent, by computing a descent direction that is feasible with respect to $\mathcal{U}$ and approximately feasible (up to first order) with respect to $\mathcal{Y}$. However, any equilibrium point of (1.18) has to be feasible and a KKT point of (1.4).

Note that, computing $\widehat{\sigma}_\alpha(u,y)$ requires the solution of a quadratic program. However, in comparison to RTI schemes, the computational effort does not scale with a prediction horizon and no explicit model of the plant dynamics is required. Instead, it is enough to estimate $\nabla_u h(u,d)$. On the other hand, in general, (1.18) does not lend itself to a natural distributed implementation. Although, depending on the problem structure, one can solve (1.19) distributedly at every iteration.

Apart from the step-size α, the *metric* G is another design parameter. Fairly standard choices are the identity matrix or the Hessian of

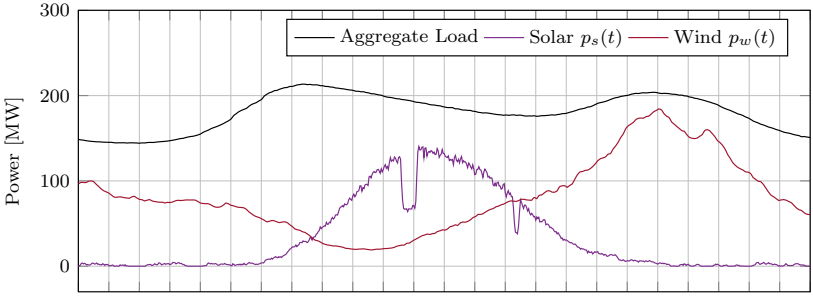

Figure 1.5: Load and generation profiles for the modified 30-bus test case in Example 1.11

the reduced cost function $\tilde{\Phi}(u) = \Phi(h(u,d),u)$, although the following example will feature a third option. ■

The following example is taken from Part III of this thesis and illustrates the performance of linearized output projection scheme described in Example 1.10 applied to the real-time power systems optimization problem from Example 1.2.

Example 1.11. To track the solution of ACOPF problem as explained in Example 1.2, we consider the output linearization approach in Example 1.10 and implement a feedback controller according to (1.18). For this purpose, we consider a modified version of the IEEE 30-bus power systems test case as considered in [Ha14]. In particular, in addition to three generators and three synchronous condensers, a wind and solar plant provide intermittent generation capacity. Their respective profiles are illustrated in Figure 1.5. The goal is hence to optimally use these renewable resources without violating engineering constraints such as voltage limits and thermal line ratings.

To implement (1.18) we evaluate $\tilde{H}(u,y) \approx \nabla_u h(u,d)$ based on input and output measurements and the power flow model (1.5). Further, we choose $G(u) := \tilde{H}(u,y)^T \tilde{H}(u,y)$.

Figure 1.6 illustrates the (almost perfect) performance achieved by this feedback-based optimization approach in terms of cost compared to

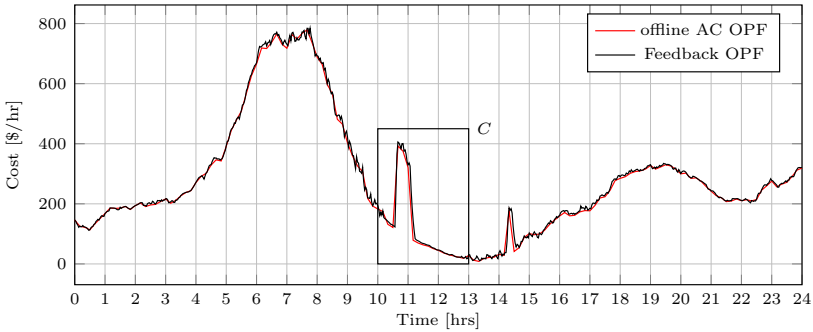

Figure 1.6: Cost value realized feedback-based optimization scheme compared to the offline a-posteriori solution for the modified 30-bus test case in Example 1.11

the a-posteriori sequential solution of the ACOPF problem.

More importantly, however, Figure 1.7 shows that over the entire simulation horizon constraint violations are very minor and only temporary. The controller achieves this by jointly managing real and reactive power infeed to manage voltage magnitudes and line currents. In particular, both solar and wind generation have to be curtailed to prevent line overloads. ■

1.4 Structure of the Thesis

The remainder of this thesis is structured into three parts, based on their level of abstraction and technical tools. Parts I and II share the same terminology, and mathematical style, whereas Part III is deliberately kept more high-level with a self-contained notation for power systems. Nevertheless, the three parts build upon each other.

Chapter 2 summarizes most of the mathematical background required for the results in the first two parts of the thesis. Where possible, the usefulness for subsequent chapters is explained for each of the tools to allow for more selective reading.

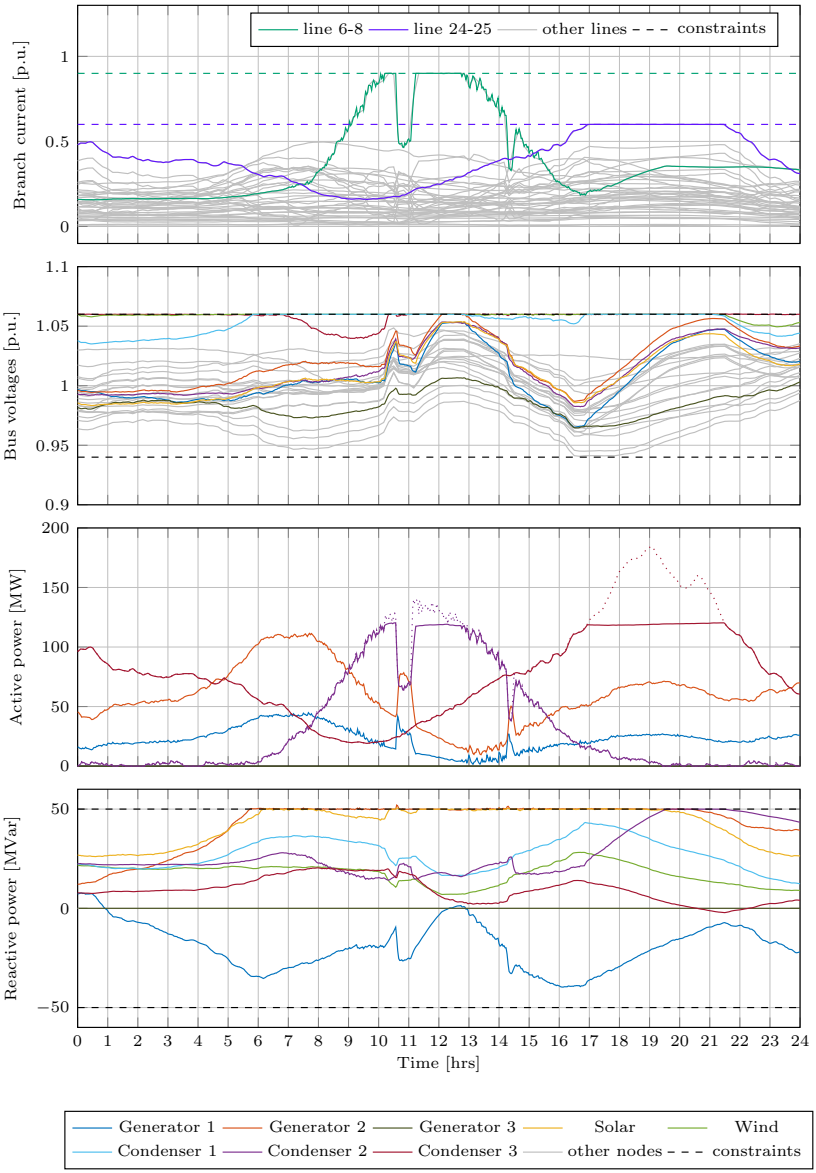

Figure 1.7: Feedback-based ACOPF optimization for 30-bus case

The definitions in Chapter 3 of *metrics* (Section 3.1) and *prox-regular sets* (Sections 3.2 and 3.3) are used throughout the thesis. Although, for simplicity, the reader may skip the latter and, instead, assume convexity (which implies prox-regularity) throughout the rest of the thesis where prox-regularity is imposed. The same simplification can also be made for Clarke regularity.

In Part I, we study the class of *projected dynamical system* (PDS) and how they can be implemented as the closed-loop behavior of feedback control loops. These insights lay the foundation to realize *constrained* feedback-based optimization in the subsequent parts.

To theoretically motivate this application, we first show in Chapter 4 that PDSs admit a coordinate-free definition, i.e., they are independent of any particular coordinate representation. This feature is desirable for any model of a physical system and, particularly, for steady-state power systems for which multiple, equivalent formulations exist.

Chapters 5 and 6 present two fundamentally new ways of approximating and implementing PDSs: One the one hand we consider *anti-windup approximations* (AWAs), which exploit physical saturation to enforce input constraints, and, on the other hand, we introduce *linearized output projection* (LOP) discretizations, which form a numerical integration scheme that can be implemented as a feedback controller. LOP discretizations are particularly suited to enforce constraints on plant outputs and turn out to be one of the most promising algorithms for constrained real-time power systems optimization, as explained in Example 1.10 and Chapter 15.

For both approximation schemes, we use the theory of perturbed differential inclusions to prove uniform convergence and semiglobal practical stability for the nominal PDS, thus motivating why we can base our control designs on PDSs.

Part II studies continuous-time optimization dynamics in the form of PDSs with a particular focus on *projected gradient flows* and their convergence, stability, and invariance properties.

Chapter 7 provides an introduction to the second part by summarizing well-known pathologies that can occur with unconstrained (i.e.,

un-projected) gradient flows and thus illustrating the natural limitations of gradient-based approaches which also extend to projected gradient flows. Furthermore, Chapter 7 also discusses the relation between projected gradient flows and existing offline optimization schemes such as *proximal algorithms*, which are recovered by particular discretization schemes.

Chapters 8, 9, and 10 mirror the structure of Part I by establishing convergence results for nominal projected gradient flows, their anti-windup approximations, and LOP discretizations, respectively. In particular, for the AWA and LOP schemes, we show *robust convergence* in the sense that trajectories of both approximations converge to the equilibria of the nominal projected gradient flow.

In Chapter 11, we establish conditions for the stable interconnection of projected gradient flows with dynamic plants. This problem is inherent to feedback-based optimization, particularly the real-time optimization of power systems with complicated, partially unknown, but stable dynamics. Using ideas from singular perturbation theory, we show that one can always choose a feedback gain small enough not to destabilize the interconnected system.

Chapter 12 proposes an extension of projected gradient flows for time-varying optimization problems. The modeling is based on so-called *perturbed sweeping processes* and establishes asymptotic tracking guarantees for convex problems. These insights are based on the general sensitivity results documented Appendix A. Although these results do not achieve the same level of generality as the previous chapters, they are informative for the local behavior of feedback-based optimization schemes for time-varying problems.

Finally, in Part III, we discuss the relevance of these theoretical results in the context of power systems operations. Chapter 13 revisits conventional steady-state AC power flow modeling and optimal power flow problems.

In Chapter 14, arguably the most critical chapter of this thesis, this power system model is cast into the framework of feedback-based optimization. What follows are Statements 1 and 2, which tie in all the

relevant theoretical results from the previous chapters. Further, the chapter also provides blueprints for numerical feedback controllers and highlights potential limitations concerning algebraic singularities.

Chapter 15 then explores three different numerical implementations and their performance in a 30-bus test grid with renewable energy infeed. These simulations help us illustrate the different characteristics of each algorithm. In particular, these computations show how well each algorithm enforces constraints such as voltage limits and line ratings in the presence of variable consumption and renewable infeed.

These simulations illustrate the robustness of the proposed schemes beyond the theoretical results by considering time delays between measurement and actuation and the impact of inaccurate model information.

Chapter 16 concludes the thesis. On the one hand, we summarize the state-of-the-art for feedback-based optimization designs based on projected gradient flows, and we list open theoretical problems. On the other hand, we propose practical use cases for feedback-based optimization in power system operations.

CHAPTER 2

Preliminaries

2.1 Notation

We consider only finite-dimensional spaces. Unless noted otherwise, we work in the usual Euclidean setup for $\mathbb{R}^n$ with inner product $\langle \cdot, \cdot \rangle$ and 2-norm $\| \cdot \|$. We write $\mathbb{R}^n_{\geq 0}$ for the non-negative orthant. The closed (open) unit ball of appropriate dimension is denoted by $\mathbb{B}$ ($\operatorname{int} \mathbb{B}$).

The closure, boundary, convex hull, and closed convex hull of a set $\mathcal{C} \subset \mathbb{R}^n$ are denoted by $\operatorname{cl} \mathcal{C}$, $\partial \mathcal{C}$, $\operatorname{co} \mathcal{C}$, and $\overline{\operatorname{co}}\, \mathcal{C}$, respectively. The set $\mathcal{C}$ is *locally compact* if it is the intersection of a closed and an open set. If $\mathcal{C}$ is non-empty, we use the notation $\|\mathcal{C}\| := \sup_{v \in \mathcal{C}} \|v\|$. A neighborhood $\mathcal{V} \subset \mathcal{C}$ of $x \in \mathcal{C}$ is understood to be a *relative neighborhood*, i.e., with respect to the subspace topology on $\mathcal{C}$. Given a convergent sequence $\{x_k\}$, the notation $x_k \underset{\mathcal{C}}{\to} x$ implies that $x_k \in \mathcal{C}$ for all k. If $x_k \in \mathbb{R}$, the notation $x \to 0^+$ means that $x_k > 0$ holds for all k and that x_k converges to 0.

Let V and W be vector spaces endowed with norms $\| \cdot \|_V$ and $\| \cdot \|_W$, respectively, and let $\mathcal{C} \subset V$. A map $F : \mathcal{C} \to W$ is *ℓ-Lipschitz* if

$$\|F(y) - F(x)\| \leq \ell \|y - x\| \tag{2.1}$$

holds for all $y, x \in \mathcal{C}$. When $\mathcal{C} = \mathbb{R}^n$, we call F *globally* ℓ-Lipschitz. The map F is *(globally) Lipschitz* if it is (globally) ℓ-Lipschitz for some $\ell > 0$.

The map F is *locally ℓ'-Lipschitz at $x' \in \mathcal{C}$* if for every $\ell \geq \ell'$ there

exists a (relative) neighborhood $\mathcal{N} \subset \mathcal{C}$ of x' such that (2.1) holds for all $x, y \in \mathcal{N}$. In other words, ℓ' is the largest lower bound on ℓ such that (2.1) is satisfied, i.e., $\ell' = \limsup_{y \to x} \|F(y) - F(x)\| / \|y - x\|$. The map F is *locally Lipschitz* (denoted by $C^{0,1}$) if, for every $x \in \mathcal{C}$, F is locally ℓ-Lipschitz for some $\ell > 0$.

Differentiability is understood in the sense of Fréchet. Namely, if $\mathcal{C}$ is open, then the map F is *differentiable at* x if there is a linear map $D_x F : V \to W$ such that

$$\lim_{y \to x} \frac{\|F(y) - F(x) - D_x F(y - x)\|_W}{\|y - x\|_V} = 0\,.$$

The map F is *differentiable* (C^1) if it is differentiable at every $x \in \mathcal{C}$. Further, F is $C^{1,1}$ if it is C^1 and $D_x F$ is $C^{0,1}$ (as function of x). Finally, given bases for V ($\dim V = m$) and W ($\dim W = n$), the *Jacobian of* F *at* x is denoted by the $n \times m$-matrix $\nabla F(x)$ of partial derivatives. If $n = 1$, then $\nabla^2 F(x) \in \mathbb{R}^{m \times m}$ denotes the *Hessian of* F *at* x.

From differential geometry, we adopt the definition that a C^1 ($C^{1,1}$) *diffeomorphism* is a C^1 ($C^{1,1}$) map $\phi : \mathcal{V} \to \mathcal{W}$ with a C^1 ($C^{1,1}$) inverse between two open sets $\mathcal{V}, \mathcal{W} \subset \mathbb{R}^n$. In other words, ϕ is an invertible coordinate transformation between $\mathcal{V}$ and $\mathcal{W}$.

The *(Euclidean) distance to* $\mathcal{C} \subset \mathbb{R}^n$ is defined as $d_{\mathcal{C}}(x) := \inf_{\tilde{x} \in \mathcal{C}} \|x - \tilde{x}\|$, and the *(Euclidean) projection* $P_{\mathcal{C}} : \mathbb{R}^n \rightrightarrows \mathcal{C}$ is given by the set $P_{\mathcal{C}}(x) := \{\tilde{x} \in \mathcal{C} \,|\, \|x - \tilde{x}\| = d_{\mathcal{C}}(x)\}$.

We use the usual definition of convexity for sets and functions. Namely, a set $\mathcal{C} \subset \mathbb{R}^n$ is *convex* if for every $x, x' \in \mathcal{C}$ we have $(1 - \delta)x + \delta x' \in \mathcal{C}$ for all $\delta \in (0, 1)$. A C^1 function $f : \mathcal{C} \to \mathbb{R}$ on a convex set $\mathcal{C}$ is *(strictly) convex* (relative to $\mathcal{C}$) if, for every $y, x \in \mathcal{C}$, we have

$$f(y) \geq (>) f(x) + \nabla f(x)(y - x)\,,$$

and f is β*-strongly convex* (for some $\beta > 0$) if

$$f(y) \geq f(x) + \nabla f(x)(y - x) + \tfrac{\beta}{2}\|y - x\|^2\,.$$

In particular, if f is twice continuously differentiable in a neighborhood of $\mathcal{C}$, f is β-strongly convex if and only if $\lambda^{\min}_{\nabla^2 f(x)} \geq \beta$ for all $x \in \mathcal{C}$. A map $F : \mathcal{C} \to \mathbb{R}^n$ is convex if all its components $(F_1, \dots F_n)$ are convex.

The identity matrix (of appropriate size) is written as $\mathbb{I}$. We denote the set of square symmetric, positive definite matrices of size n by $\mathbb{S}^n_+$. Any $G \in \mathbb{S}^n$ induces an inner product $\langle \cdot, \cdot \rangle_G$ defined by $\langle u, v \rangle_G := u^T G v$ for all $u, v \in \mathbb{R}^n$. In addition, we define the 2-norm induced by G as $\|u\|_G := \sqrt{\langle u, u \rangle}$. The *maximum* and *minimum eigenvalues* of G are denoted by $\lambda_G^{\max} := \max\{\|v\|_G \,|\, \|v\| = 1\}$ and $\lambda_G^{\min} := \min\{\|v\|_G \,|\, \|v\| = 1\}$ respectively, and the *condition number* is defined as $\kappa_G := \lambda_G^{\max}/\lambda_G^{\min}$.

In this context, also recall that the 2-norms induced by any two matrices $A, B \in \mathbb{S}^n_+$ on a finite-dimensional vector space are equivalent. That is, for a vector space V with norms $\|\cdot\|_A$ and $\|\cdot\|_B$ there are constants $\ell, \ell' > 0$ such that for every $v \in V$ it holds that $\ell \|v\|_A \leq \|v\|_B \leq \ell' \|v\|_A$. For instance, $\ell = \lambda_B^{\min}/\lambda_A^{\max}$ and $\ell' = \lambda_B^{\max}/\lambda_A^{\min}$.

2.2 Nonlinear Optimization

We briefly revisit standard notions from nonlinear optimization which will be required throughout the thesis, but in particular in Parts II and III. The following definitions and results are standard and can be found, e.g., in [22, 31, 170] and others.

Consider a nonlinear optimization problem of the form

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \Phi(x) \\ \text{subject to} \quad & h(x) = 0 \\ & g(x) \leq 0\,, \end{aligned} \tag{2.2}$$

where the objective $\Phi : \mathbb{R}^n \to \mathbb{R}$ and the constraint functions $h : \mathbb{R}^n \to \mathbb{R}^m$ and $g : \mathbb{R}^n \to \mathbb{R}^p$ are continuously differentiable. Further, we define the *feasible set of* (2.2) as $\mathcal{C} := \{x \in \mathbb{R}^n \,|\, h(x) = 0,\, g(x) \leq 0\}$. Given $x \in \mathcal{C}$, let $\mathbf{I}(x) := \{i \,|\, g_i(x) = 0\}$ denote the *set of active inequality constraints at* x and $\bar{\mathbf{I}}(x) := \{i \,|\, g_i(x) < 0\}$ is the *set of inactive inequality constraints at* x.

A point $x \in \mathbb{R}^n$ is *feasible for* (2.2) if $x \in \mathcal{C}$. A feasible point $x^\star \in \mathcal{C}$ is a *(local) optimizer*[1] *of* (2.2) if $\Phi(\hat{x}) \geq \Phi(x^\star)$ for all $\hat{x} \in (x^\star + \epsilon \mathbb{B}) \cap \mathcal{C}$

[1] To avoid confusion we reserve the term *solution* for trajectories of dynamical

for some $\epsilon > 0$. If $\Phi(\hat{x}) > \Phi(x^\star)$ for all $\hat{x} \in (x^\star + \epsilon\mathbb{B}) \cap \mathcal{C}$ with $\hat{x} \neq x^\star$ we say that $x^\star$ is a *strict optimizer of* (2.2), and $x^\star$ is a *global optimizer of* (2.2) if $\Phi(\hat{x}) \geq \Phi(x^\star)$ for all $\hat{x} \in \mathcal{C}$.

Definition 2.1 (LICQ)**.** *Given a set* $\mathcal{C} := \{x \in \mathbb{R}^n \mid h(x) = 0,\, g(x) \leq 0\}$ *where* $h : \mathbb{R}^n \to \mathbb{R}^m$ *and* $g : \mathbb{R}^n \to \mathbb{R}^p$ *are continuously differentiable, the* linear independence constraint qualification *(LICQ)* holds at $x \in \mathcal{C}$ *if*

$$\operatorname{rank} \begin{bmatrix} \nabla h(x) \\ \nabla g_{\mathbf{I}(x)}(x) \end{bmatrix} = m + |\mathbf{I}(x)|$$

holds. The set $\mathcal{C}$ satisfies LICQ *if LICQ is satisfied at all* $x \in \mathcal{C}$.

The *Lagrangian* $L : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p_{\geq 0} \to \mathbb{R}$ *of* (2.2) is defined as

$$L(x, \lambda, \mu) := \Phi(x) + \lambda^T h(x) + \mu^T g(x)\,,$$

and λ, μ are referred to as the *dual variables* (or *Lagrange multipliers*) associated with the constraints h and g, respectively.

Definition 2.2 (KKT conditions)**.** *A point* $(x^\star, \lambda^\star, \mu^\star) \in \mathcal{C} \times \mathbb{R}^m \times \mathbb{R}^p_{\geq 0}$ *satisfies the Karush-Kuhn-Tucker* (KKT) conditions for (2.2) *if*

$$\nabla_x L(x^\star, \lambda^\star, \mu^\star) = \nabla\Phi(x^\star) + (\lambda^\star)^T \nabla h(x^\star) + (\mu^\star)^T \nabla g(x^\star) = 0$$

and $\mu_i^\star = 0$ *hold for all* $i \in \bar{\mathbf{I}}^\xi_{x^\star}$ *(or, equivalently,* $\mu_i^\star g_i(x^\star) = 0$ *for all* $i = 1, \ldots p$*).*

The point $x^\star \in \mathcal{C}$ *is* first-order optimal *or* critical[2] *for* (2.2) *if there exist* $(\lambda^\star, \mu^\star)$ *such that the KKT conditions are satisfied.*

Under LICQ every optimizer is a KKT point [22, Thm. 4.2.13]:

Theorem 2.1. *If* $x^\star$ *is a local optimizer of* (2.2) *and LICQ holds at* $x^\star$*, then there exist unique* $(\lambda^\star, \mu^\star) \in \mathbb{R}^m \times \mathbb{R}^p_{\geq 0}$ *such that the KKT conditions are satisfied at* $(x^\star, \lambda^\star, \mu^\star)$*.*

systems and use the term *optimizer* (or *minimizer/maximizer*) in the context of optimization problems.

[2] In Definition 2.8 we will replace this definition with a more general geometric notion of critical points.

Remark 2.1. There exist weaker *constraint qualifications* (CQs) than LICQ that guarantee that a minimizer satisfies the KKT conditions. For instance, in convex optimization it is common to require *Slater's condition* to hold which requires the feasible set to have a non-empty relative interior. Note that Slater's condition does not imply nor is it implied by LICQ [22, Ch. 5].

Another common CQ is the *Magasarian-Fromowitz CQ* which generalizes LICQ in so far that not all of the active inequality constraints need to be linearly independent [31, Prop. 3.3.8]. Even weaker CQs, such as the *Abadie CQ* [22, Thm. 5.1.3], take a more abstract geometric form but are less easy to manipulate algebraically.

Ultimately, however, all constraint qualifications serve the purpose to connect the "algebraic" KKT conditions of (2.2) with the more geometric optimality requirement that, at an optimizer, the gradient of the objective is necessarily contained in the normal cone of the feasible set (see Example 2.2).

Throughout the thesis we limit ourselves to LICQ for three reasons:

(i) LICQ is the only CQ that guarantees that dual multipliers are unique [257].

(ii) LICQ is easy to manipulate.

(iii) For large classes of parametrized optimization problems LICQ hold *generically*, i.e., for almost all parameter values [Ha11, 230].

Nevertheless, most results presented in this thesis presumably hold under weaker CQs. ■

Corollary 2.1. *If g and Φ are convex, h is linear, and LICQ holds at $x^\star$, $x^\star$ is a global optimizer of (2.2) if and only if it is a critical point.*

Definition 2.3 (SSOSC)**.** *Consider (2.2). The point $(x^\star, \lambda^\star, \mu^\star) \in \mathcal{C} \times \mathbb{R}^m \times \mathbb{R}^p_{\geq 0}$ satisfies the* strong second-order sufficiency condition *(SSOSC) if $y^T \nabla^2_{xx} L y > 0$ holds for all $y \neq 0$ for which (i) $\nabla_x h(x^\star) y = 0$ holds, and (ii) $\nabla_x g_i(x^\star, \xi) y = 0$ holds for all i for which $\mu_i^\star > 0$.*

SSOSC guarantees that a KKT point is a strict optimizer of (2.2):

Theorem 2.2. [22, Thm. 4.4.2] *If $(x^\star, \lambda^\star, \mu^\star)$ is a KKT point for (2.2) and satisfies SSOSC, and LICQ holds at $x^\star$, then $x^\star$ is a strict local minimizer of (2.2).*

We say that $x^\star \in \mathcal{C}$ is a *regular minimizer of* (2.2) if LICQ is satisfied at $x^\star$ and there exist $(\lambda^\star, \mu^\star)$ such that $(x^\star, \lambda^\star, \mu^\star)$ form a KKT point that satisfies SSOSC.

2.2.1 Sensitivity Analysis

In Chapters 6, 10, and 12 we will encounter parametrized problems of the form

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \Phi(x, \xi), \\ \text{subject to} \quad & h(x, \xi) = 0 \\ & g(x, \xi) \leq 0\,, \end{aligned} \tag{2.3}$$

where $\xi \in \Xi$ is a perturbation parameter and $\Xi \subset \mathbb{R}^o$. Furthermore, $\Phi : \mathbb{R}^n \times \Xi \to \mathbb{R}$, $h : \mathbb{R}^n \times \Xi \to \mathbb{R}^m$, and $g : \mathbb{R}^n \times \Xi \to \mathbb{R}^p$ are continuously differentiable in x.

Analogously to the unperturbed problem (2.2), we define the (parametrized) feasible set $\mathcal{C}(\xi) := \{x \,|\, h(x, \xi) = 0, g(x, \xi) \leq 0\}$ and Lagrangian $L(x, \lambda, \mu, \xi) := \Phi(x, \xi) + \lambda^T h(x, \xi) + \mu^T g(x, \xi)$ for all $(x, \lambda, \mu) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^p_{\geq 0}$ and all $\xi \in \Xi$. For a fixed ξ, the definitions of LICQ, KKT, SSOSC, as well as Theorems 2.1 and 2.2 and Corollary 2.1 apply accordingly.

Importantly, under assumptions on the differentiability of Φ, h and g, the following result based on [138, Thm. 2.3.2 & 2.3.3] states that a regular optimizer of (2.3) (and the associated Lagrange multipliers) are locally continuous functions of ξ:

Theorem 2.3. *Consider (2.3) and assume that Φ, h, and g are twice continuously differentiable in x, and that Φ, h, g, $\nabla_x \Phi, \nabla_x h, \nabla_x g, \nabla^2_{xx}$, $\Phi \nabla^2_{xx} h$, and $\nabla^2_{xx} g$ are continuous in ξ. Further, let $\bar{x}^\star$ be a regular minimizer of (2.2) for $\bar{\xi}$ with multipliers $(\bar{\lambda}^\star, \bar{\mu}^\star)$. Then, on a neighborhood $\mathcal{N} \subset \Xi$ of $\bar{\xi}$, there exist continuous maps $x^\star : \mathcal{N} \to \mathbb{R}^n$, $\lambda^\star : \mathcal{N} \to \mathbb{R}^m$ and $\mu^\star : \mathcal{N} \to \mathbb{R}^p$ such that*

(i) $x^\star(\bar{\xi}) = \bar{x}^\star$, $\lambda^\star(\bar{\xi}) = \bar{\lambda}^\star$, *and* $\mu^\star(\bar{\xi}) = \bar{\mu}^\star$,

(ii) for all $\xi \in \mathcal{N}$, *LICQ is satisfied at* $x^\star(\xi)$, *and* $(x^\star(\xi), \lambda^\star(\xi), \mu^\star(\xi))$ *is a KKT point for which SSOSC holds,*

(iii) for all $\xi \in \mathcal{N}$, $x^\star(\xi)$ *is a local minimizer for* (2.2) *and* $(\lambda^\star(\xi), \mu^\star(\xi))$ *are the corresponding Lagrange multipliers.*

If, in addition, Φ, h *and* g *be twice continuously differentiable in* (x, ξ)

(iv) $x^\star, \lambda^\star$ *and* $\mu^\star$ *are locally Lipschitz at* $\bar{\xi}$.

Under convexity and other assumptions that guarantee a unique (global) optimizer for every ξ, the maps $x^\star, \lambda^\star$, and $\mu^\star$ exist (and are single-valued) globally on Ξ:

Corollary 2.2. *Consider* (2.3) *and assume that* Φ, h, *and* g *are* C^2 *in* x, *and that* Φ, h, g, $\nabla_x\Phi, \nabla_x h, \nabla_x g, \nabla^2_{xx}\Phi$, *and* $\nabla^2_{xx} g$ *are continuous in* ξ. *Further, for all* $\xi \in \Xi$, *let*

(i) Φ *be strongly convex in* x, g *be convex in* x, *and* h *be linear in* x,

(ii) $\mathcal{C}(\xi)$ *be non-empty,*

(iii) LICQ be satisfied for all $x \in \mathcal{C}(\xi)$.

Then, there exist continuous maps $x^\star : \Xi \to \mathbb{R}^n$, $\lambda^\star : \Xi \to \mathbb{R}^m$ *and* $\mu^\star : \Xi \to \mathbb{R}^p_{\geq 0}$ *such that* $x^\star(\xi)$ *is the unique global optimizer of* (2.3) *for all* $\xi \in \Xi$ *and* $(\lambda^\star(\xi), \mu^\star(\xi))$ *are its Lagrange multiplier. If, in addition,* Φ, h *and* g *be twice continuously differentiable in* (x, ξ), *then* $x^\star, \lambda^\star$ *and* $\mu^\star$ *are locally Lipschitz at* $\bar{\xi}$.

Proof. By assumption, (2.3) is feasible for all $\xi \in \Xi$ and LICQ holds for all $x \in \mathcal{C}(\xi)$ and all $\xi \in \Xi$. Hence, by strong convexity of Φ and convexity of $\mathcal{C}(\xi)$, (2.3) admits a unique (global) optimizer for all $\xi \in \Xi$. Therefore, the solution map $\xi \mapsto x^\star(\xi)$ and the maps $\xi \mapsto \lambda^\star(\xi)\xi \mapsto \mu^\star(\xi)$ are single-valued. Moreover, for all $\xi \in \Xi$, the KKT conditions are satisfied and SSOSC holds (trivially) by (strong) convexity (of Φ)). It then follows from Theorem 2.3 that $x^\star$ and $\mu^\star$ are continuous around every $\xi \in \Xi$ and hence on all of Ξ. □

2.3 Set-Valued & Variational Analysis

We review basic notions from set-valued analysis and variational analysis. We limit ourselves to a very reduced set of concepts. For a comprehensive treatment of the following definitions and results see [131, 211, 17].

The definitions of set convergence and semicontinuity reviewed in Section 2.3.1 will be required explicitly only for Part I. The notions of tangent and normal cones, as well as Clarke regularity are important throughout the thesis. Example 2.2 is of particular importance in that respect, since it highlights the connections between these "geometric" concepts and the nonlinear optimization theory summarized in Section 2.2.

In our context, a *set-valued map* $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ assigns to every point $x \in \mathbb{R}^n$ a set $F(x) \subset \mathbb{R}^m$. The domain of F is denoted by $\operatorname{dom} F := \{x \in \mathbb{R}^n \,|\, F(x) \neq \emptyset\}$ and its graph is given by $\operatorname{gph} F := \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m \,|\, y \in F(x)\}$. The map F is *non-empty*, *closed*, *convex*, or *compact* if, for every $x \in \mathcal{C}$, the set $F(x)$ is non-empty, closed, convex, or compact, respectively. It is *locally bounded* if for every $x \in \mathbb{R}^n$ there exists $M > 0$ such that $\|F(y)\| \leq M$ for all y in a neighborhood of x. The map F is *bounded (on $\mathcal{C}$)* if there exists $M > 0$ such that $\|F(y)\| \leq L$ for all x (all $x \in \mathcal{C}$). The same definitions also apply to single-valued functions. In particular, every continuous function is non-empty, closed, convex, compact, and locally bounded. We slightly abuse notation by writing $F(x) = a$ instead of $F(x) = \{a\}$ if F is single-valued at x.

2.3.1 Set-Convergence & Semicontinuity

Given a sequence $\{x_k\}$ with $x_k \in \mathbb{R}^n$ and a set $\mathcal{C} \subset \mathbb{R}^n$, the notation $x_k \xrightarrow[\mathcal{C}]{sub} x$ denotes the existence of a subsequence $\{x_{k'}\}$ that converges to x and $x_{k'} \in \mathcal{C}$ for all k'. Similarly, $x_k \xrightarrow[\mathcal{C}]{ev} x$ implies that $x_k \in \mathcal{C}$ holds *eventually*, i.e., for all k larger than some K, and that $\{x_k\}$ converges to x. Given a sequence of subsets $\{\mathcal{C}_k\}$ of $\mathbb{R}^n$, its *outer limit* and *inner*

limit are given, respectively, as

$$\limsup_{k\to\infty} \mathcal{C}_k := \left\{ x \;\middle|\; \exists \{x_i\} : x_i \xrightarrow[\mathcal{C}_i]{sub} x \right\}$$

and

$$\liminf_{k\to\infty} \mathcal{C}_k := \left\{ x \;\middle|\; \exists \{x_i\} : x_i \xrightarrow[\mathcal{C}_i]{ev} x \right\}.$$

In other words, the outer limit of a sequence of sets $\mathcal{C}_k$ is the set of all points x for which there exists a subsequence $\{\mathcal{C}_{i_k}\}_k$ and points $x_k \in \mathcal{C}_{i_k}$ such that $x_k \to x$. The inner limit is the set of all x for which there exists points $x_i \in \mathcal{C}_i$ for all $i \geq N$ for some N such that $x_i \to x$. In particular, note that inner and outer limits of a sequence of sets are always closed sets.

As a simple example to distinguish between inner and outer limits, consider an alternating sequence of sets given by $\mathcal{C}_{2m} := \mathcal{A}$ and $\mathcal{C}_{2m+1} := \mathcal{B}$. Then, we have $\limsup_{k\to\infty} \mathcal{C}_k = \mathcal{A} \cup \mathcal{B}$ and $\liminf_{k\to\infty} \mathcal{C}_k = \mathcal{A} \cap \mathcal{B}$.

To see this, note that on the one hand any constant sequence $\{x_k\}$ with $x_k = c \in \mathcal{A} \cap \mathcal{B}$ for all k satisfies the requirement such that $c \in \liminf_{k\to\infty} \mathcal{C}_k$. On the other hand, any sequence $\{x_k\}$ with $x_{2m} = a \in \mathcal{A}$ for $m \in \mathbb{N}$ has a trivial (constant) subsequence converging to $a \in \mathcal{A}$ and hence $a \in \limsup_{k\to\infty} \mathcal{C}_k$.

If the inner and outer limits of a sequence $\{\mathcal{C}_k\}_k$ are equal, we say that $\{\mathcal{C}_k\}$ *converges* and $\lim_{k\to\infty} \mathcal{C}_k := \limsup_{k\to\infty} \mathcal{C}_k = \liminf_{k\to\infty} \mathcal{C}_k$ is its limit. In particular, given a sequence $F_k : \mathbb{R}^n \to \mathbb{R}^m$ of set-valued maps, we say that $\{F_k\}$ *converges graphically* if the graphs of F_k converge in the sense of set convergence, i.e. $\mathcal{S} = \lim_{k\to\infty} \operatorname{gph} F_k$. As a subset of $\mathbb{R}^n \times \mathbb{R}^m$, $\mathcal{S}$ uniquely defines a set-valued map $\hat{F}$ such that $\operatorname{gph} \hat{F} = \mathcal{S}$.

Example 2.1. Recall that a sequence of single-valued functions $f : I \to \mathbb{R}^n$ where $I \subset \mathbb{R}$ is a compact interval is said to *converge uniformly* to $f : I \to \mathbb{R}^n$ if for every $\epsilon > 0$ there exists $K > 0$ such that for all $k \geq K$ and all $x \in I$ we have $|f_k(x) - f(x)| < \epsilon$. It is easy to see that uniform convergence of $\{f_k\}$ implies graphical convergence of $\{f_k\}$. ■

Further, note that, as an generalization of the *Bolzano-Weierstrass Theorem*, we know that any sequence $\{\mathcal{C}_k\}$ with $\mathcal{C}_k \subset \mathbb{R}^n$ either has a sub-

sequence converging to a non-empty set, or escapes to the horizon (i.e., for every compact set $\mathcal{B} \subset \mathbb{R}^n$ there exists K such that $\mathcal{C}_k \cap \mathcal{B} \neq \emptyset$) [211, Thm 4.8] or [118, Thm 5.7].

Finally, for a set-valued map $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ we define

$$\limsup_{y \to x} F(y) := \bigcup_{x_k \to x} \limsup_{k \to \infty} F(x_k)$$

and

$$\liminf_{y \to x} F(y) := \bigcap_{x_k \to x} \liminf_{k \to \infty} F(x_k)\,.$$

Consequently, we say that F is *outer semicontinuous* (osc) *at* x if $\limsup_{y \to x} F(y) \subset F(x)$ and *inner semicontinuous* (isc) *at* x if it holds that $\liminf_{y \to x} F(y) \supset F(x)$ [211, Def. 5.4]. In other words, F is osc if for every sequence $x_k \to x$, the outer limit of $\{F(x_k)\}_k$ is a subset of $F(x)$ and F is isc if for every sequence $x_k \to x$, the inner limit of $\{F(x_k)\}_k$ is contains $F(x)$. The map F is osc (isc) if it is osc (isc) at every $x \in \mathbb{R}^n$. Importantly, a set-valued map $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ is outer semicontinuous if and only if its graph, $\operatorname{gph} F$, is a closed subset of $\mathbb{R}^n \times \mathbb{R}^m$ [118, Lem. 5.10].

2.3.2 Variational Geometry

The tangent cone of a set at a given point is, simply speaking, the set of direction in which one can leave the point and still remain in the set. This is captured by the following definition and illustrated in Figure 2.1a.

Definition 2.4. *Given a set $\mathcal{C} \subset \mathbb{R}^n$ and $x \in \mathcal{C}$, a vector $v \in T_x\mathbb{R}^n$ is a* tangent vector of $\mathcal{C}$ at x *if there exist sequences $x_k \underset{\mathcal{C}}{\to} x$ and $\delta_k \to 0^+$ such that $\frac{x_k - x}{\delta_k} \to v$. The set of all tangent vectors is the* tangent cone of $\mathcal{C}$ at x *and denoted by $T_x\mathcal{C}$.*

The tangent cone of $\mathcal{C}$ at x, also known as *(Bouligand's) contingent cone*, can also be equivalently defined in terms of set convergence as $T_x\mathcal{C} := \limsup_{\delta \to 0^+} \frac{1}{\delta}(\mathcal{C} - x)$. Also note that, by definition, $T_x\mathcal{C}$ is closed

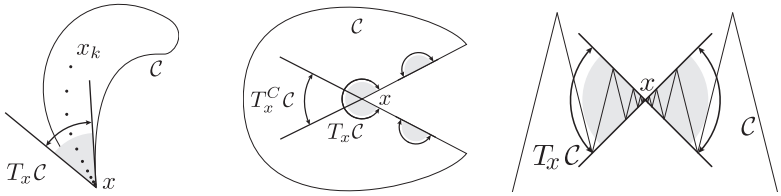

(a) (b) (c)

Figure 2.1: (a) Tangent cone construction according to Definition 2.4 by a sequence $\{x_k\}$ in $\mathcal{C}$ approaching the point x; (b) Example of a non-Clarke regular set $\mathcal{C}$, in particular, at x the (non-convex) tangent cone and the (convex) Clarke tangent cone do not coincide; (c) Example of a "non-derivable" tangent cone, where sequences $\{x_k\}$ in $\mathcal{C}$ that approach x do exist, but no continuously differentiable curve can be constructed.

and non-empty (namely, $0 \in T_x\mathcal{C}$) for any $x \in \mathcal{C}$. Furthermore, we use the convention that $T_x\mathcal{C} = \emptyset$ for all $x \notin \mathcal{C}$. Thus, we refer to $x \mapsto T_x\mathcal{C}$ as the (set-valued) *tangent cone mapping*.

Tangent cones for general sets can be fairly ill-behaved. For instance, a tangent vector $v \in T_x\mathcal{C}$ might not be *geometrically derivable*, i.e., there does not necessarily exist a continuously differentiable curve leaving x in direction v while remaining in $\mathcal{C}$, since the definition of $T_x\mathcal{C}$ calls only for the existence of a sequence of points. Moreover, the tangent cone map for a general set does not generally exhibit a particular type continuity.

Hence, a powerful property of a set is *Clarke* (or *tangential*) *regularity* which, roughly speaking, requires a set to have a boundary that is smooth, up to convex edges and corners. In this thesis, for simplicity, we consider only Clarke regularity for locally compact subsets of $\mathbb{R}^n$, i.e., sets that are the intersection of a closed and an open set. In particular, making the a-priori assumption of local compactness simplifies most definitions and results.

Definition 2.5. *For a locally compact set $\mathcal{C} \subset \mathbb{R}^n$ the* Clarke tangent cone at $x \in \mathcal{C}$ *is defined as the inner limit of the tangent cones, i.e.,* $T_x^C\mathcal{C} := \liminf_{y \to x} T_y\mathcal{C}$.

By definition of the inner limit, we have $T_x^C\mathcal{C} \subseteq T_x\mathcal{C}$. Furthermore, $T_x^C\mathcal{C}$ is closed, convex and non-empty for all $x \in \mathcal{C}$ [211, Thm. 6.26].

Definition 2.6. *We call a set $\mathcal{C} \subset \mathbb{R}^n$* Clarke regular *at x if it is locally compact and $T_x\mathcal{C} = T_x^C\mathcal{C}$. The set $\mathcal{C}$ is* Clarke regular *if it is Clarke regular for all $x \in \mathcal{C}$.*

A set that is not Clarke regular is illustrated in Figure 2.1b.

Clarke regularity of a set also implies that several notions of normal cones (such as *limiting*, *Clarke*, and *regular normal cones*) coincide and reduce to the following simple definition of the normal cone as the polar cone to the tangent cone:

Definition 2.7. *Let $\mathcal{C} \subset \mathbb{R}^n$ be Clarke regular, then the* normal cone *at $x \in \mathcal{C}$, denoted by $N_x\mathcal{C}$, is defined as the polar cone of $T_x^C\mathcal{C}$, i.e.,*

$$N_x^G\mathcal{C} := \left\{\eta \,\middle|\, \forall v \in T_x^C\mathcal{C} : \langle v, \eta\rangle \leq 0\right\} .$$

For $x \notin \mathcal{C}$, we adopt the convention that $N_x\mathcal{C} := \emptyset$. Note that, in contrast to $T_x\mathcal{C}$, the definition of the normal cone relies on an inner product defined on $\mathbb{R}^n$. An important property of Clarke regular sets (which will be generalized in the next chapter) is that the normal cone mapping $x \mapsto N_x\mathcal{C}$ is osc [211, Prop. 6.5].

As an important interlude, we next show how the concepts from nonlinear optimization in Section 2.2 can be interpreted from a geometric perspective:

Example 2.2. Consider the optimization problem

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \qquad x \in \mathcal{C}\,, \tag{2.4}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is non-empty, closed and Clarke regular, and $\Phi : \mathbb{R}^n \to \mathbb{R}$ is continuously differentiable in a neighborhood of $\mathcal{C}$.

Every local minimizer $x^\star \in \mathcal{C}$ of (2.4) satisfies $-\nabla\Phi(x^\star)^T \in N_{x^\star}\mathcal{C}$ [211, Thm. 6.12]. This statement is an abstraction of Theorem 2.1 as we explain next.

Consider a set $\mathcal{C} := \{x \in \mathbb{R}^n \,|\, h(x) = 0, g(x) \leq 0\} \subset \mathbb{R}^n$ where $h : \mathbb{R}^n \to \mathbb{R}^m$, $g : \mathbb{R}^n \to \mathbb{R}^p$ are continuously differentiable as before

in Section 2.2. Then, if LICQ is satisfied for all $x \in \mathcal{C} := \{x \,|\, h(x) = 0,\, g(x) \leq 0\}$, $\mathcal{C}$ is Clarke regular [211, Thm. 6.31]. In particular, the (Clarke) tangent cone at x is given by

$$T_x^C \mathcal{C} = T_x \mathcal{C} = \left\{ v \,|\, \nabla h(x) v = 0,\, \nabla g_{\mathbf{I}(x)}(x) v \leq 0 \right\},$$

where $\mathbf{I}(x) := \{i \,|\, g_i(x) = 0\}$. Moreover, the normal cone of $\mathcal{C}$ at x is given by

$$N_x \mathcal{C} = \left\{ \eta \,\middle|\, \eta = \nabla h(x)^T \lambda + \nabla g_{\mathbf{I}(x)}(x)^T \mu_{\mathbf{I}},\, \mu_{\mathbf{I}} \in \mathbb{R}_{\geq 0}^{|\mathbf{I}(x)|} \right\}. \tag{2.5}$$

In particular, for a KKT point $(x^\star, \lambda^\star, \mu^\star) \in \mathcal{C} \times \mathbb{R}^m \times \mathbb{R}_{\geq 0}^p$ according to Definition 2.2 we have

$$\nabla \Phi(x^\star)^T + \underbrace{\nabla h(x^\star)^T \lambda + \nabla g(x^\star)^T \mu}_{\in N_{x^\star} \mathcal{C}} = 0$$

since $\mu_i = 0$ for all $i \notin \mathbf{I}(x)$. Conversely, if $-\nabla \Phi(x^\star)^T \in N_{x^\star} \mathcal{C}$, then $x^\star$ is a first-order optimal point of Definition 2.2.

Thus, in summary, constraint qualifications such as LICQ guarantee that a set defined by constraints h and g form a Clarke regular set, and their normal cone takes an explicit form (2.5) and allows us to recover the classical "algebraic" KKT conditions. ■

In view of Example 2.2, we generalize the notion of critical point from Definition 2.2 as follows:

Definition 2.8. *Consider a Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$, $\Phi : \mathbb{R}^n \to \mathbb{R}$ continuously differentiable in a neighborhood of $\mathcal{C}$. The point $x^\star \in \mathcal{C}$ is a* critical point *of*

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad x \in \mathcal{C}$$

if and only if it holds that $-\nabla \Phi(x^\star)^T \in N_{x^\star} \mathcal{C}$.

The next lemma will be useful to eliminate equality constraints from optimization problems.

Lemma 2.1. *Let $\mathcal{C} \subset \mathbb{R}^m$ be Clarke regular and let $h : \mathbb{R}^m \to \mathbb{R}^p$ and $\Phi : \mathbb{R}^p \times \mathbb{R}^m$ be continuously differentiable in a neighborhood of $\mathcal{C}$ and $\mathbb{R}^p \times \mathcal{C}$, respectively. Then, $(y^\star, x^\star) \in \mathbb{R}^p \times \mathcal{C}$ is a critical point (minimizer) of*

$$\begin{aligned} \underset{y,x}{\text{minimize}} \quad & \Phi(y, x) \\ \text{subject to} \quad & y = h(x) \\ & x \in \mathcal{C} \end{aligned} \tag{2.6}$$

if and only if $x^\star$ is a critical point (minimizer) of

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \tilde{\Phi}(x) \\ \text{subject to} \quad & x \in \mathcal{C} \end{aligned} \tag{2.7}$$

where $\tilde{\Phi} : \mathbb{R}^m \to \mathbb{R}$ is defined as $\tilde{\Phi}(x) := \Phi(h(x), x)$ and $y^\star = h(x^\star)$.

Proof. Define $\mathcal{X} := \{(y, x) \,|\, y - h(x) = 0\} \subset \mathbb{R}^p \times \mathbb{R}^m$ which is a Clarke regular by Example 2.2 since LICQ is satisfied. Furthermore, we have

$$N_{(y,x)}\mathcal{X} = \left\{ \eta \,\middle|\, \eta = \begin{bmatrix} \mathbb{I}_p \\ -\nabla h(x)^T \end{bmatrix} \lambda,\ \lambda \in \mathbb{R}^p \right\}.$$

From [211, Thm. 6.14 & 6.42] it thus follows that the feasible set of (2.6) given by $\tilde{\mathcal{X}} := \mathcal{X} \cap (\mathbb{R}^p \times \mathcal{C})$ is Clarke regular. Moreover, we have

$$\begin{aligned} N_{(y,x)}\tilde{\mathcal{X}} &= N_{(y,x)}\mathcal{X} + N_{(y,x)}(\mathbb{R}^p \times \mathcal{C}) \\ &= \left\{ \eta \,\middle|\, \eta = \begin{bmatrix} \mathbb{I}_p \\ -\nabla h(x)^T \end{bmatrix} \lambda + \begin{bmatrix} 0 \\ v \end{bmatrix},\ \lambda \in \mathbb{R}^p,\ v \in N_x\mathcal{C} \right\}. \end{aligned}$$

Hence, if $(y^\star, x^\star) \in \tilde{\mathcal{X}}$ is a critical point of (2.6) then, by definition, we have that $-\nabla\Phi(y^\star, x^\star)^T \in N_{(y^\star,x^\star)}\tilde{\mathcal{X}}$, or equivalently

$$-\nabla\Phi(y^\star, x^\star)^T = \begin{bmatrix} \mathbb{I}_p \\ -\nabla h(x^\star)^T \end{bmatrix} \lambda + \begin{bmatrix} 0 \\ v \end{bmatrix} \tag{2.8}$$

for some $\lambda \in \mathbb{R}^p$ and $v \in N_{x^\star}\mathcal{C}$.

Next, note that $\begin{bmatrix} \nabla h(x^\star)^T & \mathbb{I}_m \end{bmatrix} \begin{bmatrix} \mathbb{I}_p \\ -\nabla h(x^\star)^T \end{bmatrix} = 0$, and therefore (2.8) implies that

$$-\underbrace{\begin{bmatrix} \nabla h(x^\star)^T & \mathbb{I}_m \end{bmatrix} \nabla \Phi(y^\star, x^\star)^T}_{\nabla \tilde{\Phi}(x^\star)^T} = \begin{bmatrix} \nabla h(x^\star)^T & \mathbb{I}_m \end{bmatrix} \begin{bmatrix} 0 \\ v \end{bmatrix} = v \in N_{x^\star}\mathcal{C}, \tag{2.9}$$

and therefore $x^\star$ is a critical point of (2.7).

Conversely, let $x^\star \in \mathcal{C}$ be a critical point of (2.7) and $y^\star = h(x^\star)$. Reversing the line of arguments above, it is easy to see that $\nabla \Phi(y^\star, x^\star)^T \in N_{(y^\star, x^\star)}\tilde{\mathcal{X}}$ and thus $(y^\star, x^\star)$ is a critical point of (2.6).

The equivalence of minimizers of (2.6) and (2.7) follows from the topological equivalence of the two problems: Because there is a continuous one-to-one mapping between feasible points of (2.6) and (2.7), if $x^\star$ is a minimizer $\tilde{\Phi}$ on $\mathcal{C}$, then $(h(x^\star), x^\star)$ minimizes Φ on $\tilde{\mathcal{X}}$. □

The next lemma shows that a C^1 diffeomorphism (i.e., a C^1 coordinate transformation) maps (Clarke) tangent cones to (Clarke) tangent cones. Hence, C^1 diffeomorphisms preserve Clarke regularity [211, Ex. 6.7].

Lemma 2.2. *Let $\mathcal{V}, \mathcal{W} \subset \mathbb{R}^n$ be open and consider a C^1 diffeomorphism $\phi : \mathcal{V} \to \mathcal{W}$. Given $\mathcal{C} \subset \mathbb{R}^n$ and $\tilde{\mathcal{C}} := \mathcal{C} \cap \mathcal{V}$, for every $x \in \tilde{\mathcal{C}}$ it holds that*

$$T_{\phi(x)}\phi(\tilde{\mathcal{C}}) = D_x\phi(T_x\tilde{\mathcal{C}}) \tag{2.10}$$

$$T^C_{\phi(x)}\phi(\tilde{\mathcal{C}}) = D_x\phi(T^C_x\tilde{\mathcal{C}})\,. \tag{2.11}$$

Hence, $\phi(\tilde{\mathcal{C}})$ is Clarke regular at $\phi(x)$ if and only if $\tilde{\mathcal{C}}$ is Clarke regular at $x \in \tilde{\mathcal{C}}$.

Proof. We need to show only that $T_{\phi(x)}\phi(\tilde{\mathcal{C}}) \subset D_x\phi(T_x\tilde{\mathcal{C}})$. Since ϕ is a C^1 diffeomorphism the other direction follows by applying the same arguments to ϕ^{-1}.

Let $v \in T_x\tilde{\mathcal{C}}$. Then, by definition there exist $x_k \to x$ with $x_k \in \tilde{\mathcal{C}}$ and $\delta_k \to 0^+$ such that $(x_k - x)/\delta_k \to v$. Furthermore, $\|x_k - x\|/\delta_k$ converges to $\|v\|$. According to the definition of the derivative of ϕ, for the same sequence $\{x_k\}$ we have $\lim_{k\to\infty} \|\phi(x_k) - \phi(x) - D_x\phi(x_k - x)\|/\|x_k - x\| = 0$.

Since the limit of the element-wise product of convergent sequences equals the product of its limits we can write

$$\lim_{k\to\infty} \frac{\|\phi(x_k)-\phi(x)-D_x\phi(x_k-x)\|}{\|x_k-x\|}\frac{\|x_k-x\|}{\delta_k} = 0$$

which, using the fact that $D_x\phi$ is linear, simplifies to

$$\lim_{k\to\infty} \left\| \frac{\phi(x_k)-\phi(x)}{\delta_k} - D_x\phi\left(\frac{x_k-x}{\delta_k}\right)\right\| = 0\,.$$

This implies that $(\phi(x_k)-\phi(x))/\delta_k \to D_x\phi(v)$, and hence $D_x\phi(v)$ is a tangent vector of $\phi(\tilde{\mathcal{C}})$ at $\phi(x)$. This proves (2.10).

To show (2.11) we use (2.10) together with the definition of the Clarke tangent cone as the inner limit of the surrounding tangent cones (Definition 2.5). We can write

$$T^C_{\phi(x)}\phi(\tilde{\mathcal{C}}) = \liminf_{\hat{y}\to\phi(x)} T_{\hat{y}}\phi(\tilde{\mathcal{C}}) = \liminf_{y\to x} D_y\phi\left(T_y\tilde{\mathcal{C}}\right)\,.$$

Since $D_x\phi$ is continuous in x, we have

$$\liminf_{y\to x} D_y\phi(T_y\tilde{\mathcal{C}}) = \liminf_{y\to x} D_x\phi(T_y\tilde{\mathcal{C}})\,.$$

Further, Lemma 4.5 implies that

$$\liminf_{y\to x} D_x\phi(T_y\tilde{\mathcal{C}}) \supset D_x\phi(\liminf_{y\to x} T_y\tilde{\mathcal{C}}) = D_x\phi(T^C_x\tilde{\mathcal{C}})$$

and therefore we have $T^C_{\phi(x)}\phi(\tilde{\mathcal{C}}) \supset D_x\phi(T^C_x\tilde{\mathcal{C}})$. Again, since ϕ is a diffeomorphism, the opposite inclusion holds by applying the same argument to ϕ^{-1}. This shows (2.11) and completes the proof. □

2.4 Dynamical Systems & Stability

2.4.1 Continuous-Time Systems

In general, we understand a continuous-time dynamical systems to be defined by a differential inclusion (e.g.,[18, 100, 17] or [118, Ch. 5]). Namely, given a set-valued map $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and a set $\mathcal{C} \subset \mathbb{R}^n$, we

say that $x : [0, T] \to \mathcal{C}$ for some $T > 0$ is a *solution* of the (constrained) differential inclusion

$$\dot{x} \in H(x)\,, \qquad x \in \mathcal{C} \tag{2.12}$$

if x is absolutely continuous, x is *viable* (i.e., $x(t) \in \mathcal{C}$ for all $t \in [0, T]$), and $\dot{x}(t) \in H(x(t))$ holds for almost all $t \in [0, T]$.

A map $x : [0, \infty) \to \mathcal{C}$ is a *complete* solution of (2.12) if its restriction to any compact subset $[0, T]$ is a solution of (2.12). A solution is *maximal* if it is complete or cannot be extended, i.e., there does not exist a solution $x' : [0, T'] \to \mathcal{C}$ with $T' > T$ such that $x(t) = x'(t)$ for all $t \in [0, T]$. A solution is *unique* if for every other solution $x' : [0, T'] \to \mathbb{R}^n$ with $x'(0) = x(0)$ it holds that $x(t) = x'(t)$ for all $t \in [0, \min\{T, T'\}]$.

Definition 2.9. *An inclusion* (2.12) *is* well-posed *if $\mathcal{C}$ is closed, H is osc*[3] *and locally bounded relative to $\mathcal{C}$, and $H(x)$ is non-empty and convex for all $x \in \mathcal{C}$.*

Definition 2.10. *Given a set-valued map $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, a closed set $\mathcal{C} \subset \operatorname{dom} H$ is a* viability domain of H *if and only if $F(x) \cap T_x\mathcal{C} \neq \emptyset$ for all $x \in K$.*

Existence of solutions to constrained differential inclusions is guaranteed by the following result, which also states that solutions evolving on a viability domain are either complete or escape to the horizon in finite time.

Theorem 2.4 (Viability Theorem [18, Thm. 3.3.4 & 3.3.5])**.** *Let* (2.12) *be well-posed and $\mathcal{C}$ be a viability domain for H. Then, for any initial condition $x(0) \in \mathcal{C}$,* (2.12) *admits either a complete solution or a maximal solution on $[0, T)$ such that $\limsup_{t \to T^-} \|x(t)\| = \infty$.*

Further, if there exists $c > 0$ such that $\|H(x)\| \leq c(\|x\| + 1)$ holds for all $x \in \operatorname{dom} H$, then there exists a complete solution to (2.12) *for every initial condition $x_0 \in K$.*

[3] Recall that under local boundedness, outer semicontinuity and upper semicontinuity are equivalent [118, Lem. 5.16].

The following result is particularly useful for the study of viable trajectories.

Lemma 2.3. *Given a set $\mathcal{C} \subset \mathbb{R}^n$, an absolutely continuous map $x : [0,T] \to \mathcal{C}$ with $T \in [0,\infty]$ satisfies $\dot{x}(t) \in T_{x(t)}\mathcal{C} \cap -T_{x(t)}\mathcal{C}$ for almost all $t \in [0,T]$, where $-T_{x(t)\mathcal{C}} := \{v | -v \in T_{x(t)}\mathcal{C}\}$.*

Proof. Let $t \in (0,T)$ be such that $\dot{x}(t)$ exists. This implies that by definition

$$\dot{x}(t) = \lim_{\tau \to 0^+} \frac{x(t+\tau)-x(t)}{\tau} = \lim_{\tau \to 0^+} \frac{x(t)-x(t-\tau)}{\tau}.$$

Thus, by choosing any sequence $\tau_k \to 0$ with $\tau_k > 0$, the sequence $\frac{x(t+\tau_k)-x(t)}{\tau_k}$ converges to a tangent vector and $\frac{-x(t-\tau_k)+x(t)}{\tau_k}$ converges to a vector in $-T_{x(t)}\mathcal{C}$ by definition of $T_{x(t)}\mathcal{C}$ and the fact that $x(t) \in \mathcal{C}$ for all $t \in [0,T)$. □

For convenience, we introduce the following notion of truncated solution for unconstrained inclusions which will be used in Chapter 5 to study anti-windup approximations of projected dynamical systems:

Definition 2.11. *Consider (2.12), let $T, \epsilon > 0$ and $x_0 \in \mathbb{R}^n$. A solution $x : [0,T'] \to \mathcal{C}$ of (2.12) with initial condition $x(0) = x_0$ is (T,ϵ)-truncated if $x(t) \in x_0 + \epsilon\mathbb{B}$ for all $t \in [0,T']$ and either $T' = T$ or $\|x(T') - x_0\| = \epsilon$ holds.*

As stated above, Theorem 2.4 implies that maximal solutions of (2.12) are either complete or escape to the horizon. Therefore, by considering an augmented inclusion with $\hat{H}(x) := (H(x), 1)$, initial condition $\hat{x}(0) := (x(0), 0)$, $\hat{\mathcal{C}} = \mathbb{R}^n \times \mathbb{R}$, and $\hat{\mathcal{A}} = \mathcal{A} \times [0,T]$, Theorem 2.4 guarantees the existence of truncated solutions for every T and every ϵ:

Corollary 2.3. *Let (2.12) be well-posed and $\mathcal{C}$ be a viability domain of H. Then, for every $T, \epsilon > 0$ and every $x(0) \in \mathbb{R}^n$ there exists a (T,ϵ)-truncated solution.*

Hence, truncated solutions are convenient if finite escape times cannot be precluded, since their graph is always a compact subset of $[0,T] \times (x(0) + \epsilon\mathbb{B})$.

We also require the notion of σ-*perturbation* of an inclusion which will be convenient to prove uniform convergence and semiglobal practical robust asymptotic stability for anti-windup approximations and linearized output projection discretizations in Chapter 5 and Chapter 6, respectively.

Definition 2.12 ([118, Def. 6.27])**.** *Given $\sigma > 0$, the* σ-perturbation of (2.12) *is given by*

$$\dot{x} \in H_\sigma(x) \qquad x \in \mathcal{C}_\sigma \,,$$

where

$$\mathcal{C}_\sigma := \mathcal{C} + \sigma\mathbb{B} \quad \textit{and} \quad H_\sigma := \overline{\mathrm{co}}\, H((x + \sigma\mathbb{B}) \cap \mathcal{C}) + \sigma\mathbb{B} \quad \forall x \in \mathcal{C}_\sigma \,.$$

Note in particular that for $\sigma' \geq \sigma$ we have $\mathcal{C}_\sigma \subset \mathcal{C}_{\sigma'}$, $H_\sigma(x) \subset H_{\sigma'}(x)$ for all $x \in \mathcal{C}_\sigma$, and every solution of the σ-perturbation is a solution of the σ'-perturbation.

2.4.2 Notions of Stability

Consider (2.12) and a compact[4] set $\mathcal{A} \subset \mathcal{C}$. We say that $\mathcal{A}$ is

(i) *stable* (for (2.12)), if for every $\epsilon > 0$ there exists $\delta > 0$ such that every solution $x : [0, T] \to \mathcal{C}$ with $d_\mathcal{A}(x(0)) \leq \delta$ satisfies $d_\mathcal{A}(x(t)) \leq \epsilon$ for all $t \in [0, T]$,

(ii) *locally asymptotically stable* (las) if it is stable and there exists $\gamma > 0$ such that every complete solution of (2.12) with $d_\mathcal{A}(x(0)) \leq \gamma$ converges to $\mathcal{A}$, i.e., $\lim_{t\to\infty} d_\mathcal{A}(x(t)) = 0$,

(iii) *globally asymptotically stable* (gas) it is stable and every complete solution of (2.12) with $d_\mathcal{A}(x(0)) \leq \gamma$ converges to $\mathcal{A}$, i.e., $\lim_{t\to\infty} d_\mathcal{A}(x(t)) = 0$,

[4]Throughout the thesis, we limit ourselves to stability of compact sets and definitions of (asymptotic) stability are automatically "uniform".

Furthermore, we say that $\hat{x} \in \mathcal{C}$ is a *weak equilibrium* of (2.12) if the constant map $t \mapsto \hat{x}$ is a solution of (2.12). Namely, $\hat{x}$ is a weak equilibrium if and only if $0 \in F(\hat{x})$. The point $\hat{x}$ is a *strong equilibrium* of (2.12) if $t \mapsto \hat{x}$ is the only solution of (2.12) starting at $\hat{x}$.

Next, recall that $\omega : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a $\mathcal{K}_\infty$*-function* (denoted by $\omega \in \mathcal{K}_\infty$) if ω is continuous, strictly increasing, unbounded, and it holds that $\omega(0) = 0$. In particular, a compact set $\mathcal{A}$ is stable if and only if there exists $\omega \in \mathcal{K}_\infty$ such that for every solution $x : [0, T] \to \mathcal{C}$ of (2.12) we have $d_{\mathcal{A}}(x(t)) \leq \omega(d_{\mathcal{A}}(x(0)))$ for all $t \in [0, T]$.

Furthermore, recall that a $\mathcal{K}_\infty$-function can serve as a *modulus of continuity* for a uniformly continuous function. Namely, every map $f : \mathcal{V} \to \mathbb{R}^m$ for some $\mathcal{V} \subset \mathbb{R}^n$ that is uniformly continuous on $\mathcal{V}$ admits an $\omega \in \mathcal{K}_\infty$ such that

$$\|f(y) - f(x)\| \leq \omega(\|y - x\|)$$

for all $y, x \in \mathcal{V}$.

A function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a $\mathcal{KL}$*-function* (denoted by $\beta \in \mathcal{KL}$) if it is non-decreasing in its first argument, non-increasing in its second argument, $\lim_{r \to 0^+} \beta(r, s) = 0$ for each $s \in \mathbb{R}_{\geq 0}$, and $\lim_{s \to \infty} \beta(r, s) = 0$ for each $t \in \mathbb{R}_{\geq 0}$.

In particular, we have that a compact set $\mathcal{A} \subset \mathbb{R}^n$ is gas for (2.12) if and only if there exists $\beta \in \mathcal{KL}$ such that for every complete solution $x : [0, \infty) \to \mathcal{C}$ of (2.12) it holds that

$$d_{\mathcal{A}}(x(t)) \leq \beta(d_{\mathcal{A}}(x(0)), t) \qquad \forall t \in [0, \infty)\,.$$

Remark 2.2. Let (2.12) be well-posed and let $\mathcal{C}$ be a viability domain for H. Further, assume that $\mathcal{A} \subset \mathcal{C}$ is compact and gas. Any solution $x : [0, T] \to \mathcal{C}$ of (2.12) does not leave the (compact) set $\{x \,|\, d_{\mathcal{A}}(x) \leq \beta(d_{\mathcal{A}}(x(0)), t)\}$ and thus cannot escape to the horizon. By Theorem 2.4 it follows that for any initial condition $x(0) \in \mathcal{C}$, (2.12) admits a complete solution. ■

Consider a class of continuous-time dynamical systems $\{(H_\alpha, \mathcal{C}_\alpha)\}_\alpha$ for $\alpha \geq 0$ defined by the constrained inclusions

$$\dot{x} \in H_\alpha(x)\,, \qquad x \in \mathcal{C}_\alpha\,.$$

We say that a set $\mathcal{A} \subset \mathbb{R}^n$ is *semiglobally practically asymptotically stable* (spas) *for* $\{(H_\alpha, \mathcal{C}_\alpha)\}_\alpha$ if for every $\zeta > 0$ and every compact set $\mathcal{B}$, there exists $\alpha^\star > 0$ such that for all $\alpha < \alpha^\star$, every solution of $\dot{x} \in H_\alpha(x), x \in \mathcal{C}_\alpha$ starting in $\mathcal{B} \cap \mathcal{C}_\alpha$ converges to $\mathcal{A} + \zeta\mathbb{B}$. In other words, for every compact set of initial conditions, every neighborhood of $\mathcal{A}$ can be rendered attractive for a small enough α.

Invariance

The *ω-limit set* of a complete solution of (2.12), denoted by $\Omega(x)$, is the set of all points $\hat{x}$ for which there exists a sequence $\{t_k\}$ with $\lim_{k\to\infty} t_k = \infty$ and $\lim_{k\to\infty} x(t_k) = \hat{x}$.

A set $\mathcal{A} \subset \mathcal{C}$ is *weakly invariant*, if and only if for every initial condition $x_0 \in \mathcal{A}$, there exists a complete solution starting at x_0 that remains in $\mathcal{A}$ for all $t \in [0, \infty)$. The union of weakly (invariant subsets is again weakly invariant, hence the notion of *largest weakly invariant set* is well-defined. A set $\mathcal{A} \subset \mathcal{C}$ is *invariant*, if and only if for every initial condition $x_0 \in \mathcal{A}$, every complete solution starting at x_0 remains in $\mathcal{A}$ for all $t \in [0, \infty)$.

We will make us of the following invariance principle for differential inclusions. The result is a special case of [118, Thm. 8.2] which applies to hybrid systems. For a similar result for differential inclusion see [214].

Theorem 2.5. *Consider a continuous function $V : \mathbb{R}^n \to \mathbb{R}$, any functions $u : \mathbb{R}^n \to [-\infty, \infty]$, and a set $\mathcal{U} \subset \mathbb{R}^n$ such that $u(x) \leq 0$ for every $x \in \mathcal{U}$ and such that the growth of V along solutions of a well-posed inclusion $\dot{x} \in F(x)$ is bounded by u on $\mathcal{U}$, i.e., any solution $x : [0, T] \to \mathcal{U}$ of (2.12) satisfies*

$$V(x(t_1)) - V(x(t_0)) \leq \int_{t_0}^{t_1} u(x(\tau))d\tau\,.$$

Let a complete and bounded solution x of (2.12) be such that $x(t) \in \mathcal{U}$ for all $t \in [0, \infty)$. Then, for some $r \in V(\mathcal{U})$, x approaches the nonempty set that is the largest weakly invariant subset of $V^{-1}(r) \cap \mathcal{U} \cap \operatorname{cl} u^{-1}(0)$.

2.4.3 Discrete-Time Systems

In Chapters 6 and 10 we consider a class of iterative algorithms which we analyze as discrete-time dynamical systems. For this purpose, and for simplicity, we introduce only discrete-time systems with unique and complete solutions, namely, we consider only recursions of the form

$$x^+ = H(x) \qquad x \in \mathcal{C}\,, \tag{2.13}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is non-empty and $H : \mathcal{C} \to \mathcal{C}$.

A solution of (2.13) is a map $x : \mathbb{N}_0 \to \mathcal{C}$ such that $x[k+1] = H(x[k])$ for all $k \in \mathbb{N}_0$. From the fact that $\mathcal{C}$ is non-empty and the definition of H it follows immediately that (2.13) admits a unique complete solution for every initial condition $x[0] \in \mathcal{C}$.

The notions of *global asymptotic stability, semiglobal practical asymptotic stability* and *invariance* carry over from continuous-time systems. Furthermore, we will make use of the following invariance theorem (which is again a special case of [118, Thm. 8.2]). The original result for discrete-time systems can be found in [156, Thm. 6.3].

Theorem 2.6. *Consider a discrete-time dynamical system* (2.13) *where* $H : \mathcal{C} \to \mathcal{C}$ *is continuous and* $\mathcal{S} \subset \mathcal{C}$ *is an invariant set. Further, let* $V : \mathcal{S} \to \mathbb{R}$ *be a continuous function such that* $V(T(x)) \leq V(x)$ *for all* $x \in \mathcal{S}$*. Let* $\mathbf{x} = \{x_0, x_1, x_2, \ldots\} \subset \mathcal{S}$ *be a bounded solution. Then, for some* $r \in V(\mathcal{S})$*,* $\mathbf{x}$ *converges to the non-empty set that is the largest invariant subset of* $V^{-1}(r) \cap \mathcal{S} \cap \{x \,|\, V(T(x)) - V(x) = 0\}$*.*

CHAPTER 3

Coordinate-Free Prox-Regularity

3.1 Low-Regularity Riemannian metrics

Throughout this thesis, wherever possible, we adopt a *geometric* viewpoint according to which qualitative features of dynamical systems need to be preserved under nonlinear coordinate transformations. For this purpose, we borrow the idea of a *metric* from Riemannian geometry.

Simply speaking, a metric defines an inner product at every point of a subset of $\mathbb{R}^n$ (or, more generally, at every point on a manifold). Under standard assumptions adopted in differential geometry, a metric allows one to define the length of a curve, the (geodesic) distance between points, different notions of curvature, etc.

In this thesis, we employ the notion of a metric for two purposes:

- In the context of projected dynamical systems (Chapter 4), a metric is a convenient object to model *oblique* projection directions at the boundary of a domain.

The material presented in this chapter is largely taken from [Ha2]. In particular Propositions 3.2, 3.3, and 3.4 constitute the main technical results of this chapter. Further expository material from [Ha4, Ha5] is included and will be required in Chapter 5 and Section 9.3.

- We use metrics to study variations of optimization algorithms under a common model of projected gradient flows.

As a consequence, we do not require any additional notions from Riemannian geometry. Instead, we derive our results from first principles, which also allows us to retain a higher degree of generality and use fewer advanced notions from differential geometry.

For simplicity, we define a *metric* to be a map from some set $\mathcal{C} \subset \mathbb{R}^n$ to the space of symmetric positive definite matrices $\mathbb{S}^n_+$. We consider $\mathbb{S}^n_+$ to be supplied with the $\lambda^{\max}$-norm. Hence, a metric is (Lipschitz) continuous if it is (Lipschitz) continuous as a map $G : \mathcal{C} \to \mathbb{S}^n_+$.

From a more abstract viewpoint, a metric is a map from $\mathcal{C}$ to the space of bilinear forms over the tangent space $T_x\mathbb{R}^n$ at $x \in \mathcal{C}$. This more formal perspective was adopted in [Ha2].

In accordance with our notation for symmetric positive definite matrices, we also define the functions $\lambda^{\max}_G : \mathcal{C} \to \mathbb{R}_+$, $\lambda^{\min}_G : \mathcal{C} \to \mathbb{R}_+$, and $\kappa_G : \mathcal{C} \to \mathbb{R}_+$ to assign to every $x \in \mathcal{C}$ the maximum or minimum eigenvalue, or condition number of $G(x)$, respectively. In particular, since $G(x)$ is positive definite for all $x \in \mathcal{C}$, by definition, it follows that $\lambda^{\max}_{G(x)}, \lambda^{\min}_{G(x)}$ and $\kappa_{G(x)}$ are well-defined for all $x \in \mathcal{C}$. However, $\kappa_{G(x)}$ is not necessarily locally bounded, even if G is bounded as a map. In particular, $\lambda^{\min}_{G(x)}$ might not be bounded below, away from 0. Hence, for metrics we require the following definition of local boundedness.

Definition 3.1. *A metric $G : \mathcal{C} \to \mathbb{S}^n_+$ on $\mathcal{C} \subset \mathbb{R}^n$ is* locally weakly bounded *if for every $x \in \mathcal{C}$ there exist $\ell, \bar{\ell} > 0$ such that $\ell \leq \kappa_G(y) \leq \bar{\ell}$ holds for all $y \in \mathcal{C}$ in a neighborhood of x. The metric is* weakly bounded *if $\ell \leq \kappa_G(x) \leq \bar{\ell}$ holds for all $x \in \mathcal{C}$.*

A metric can be locally weakly bounded even if it is not locally bounded. Furthermore, since maximum and minimum eigenvalues (and hence the condition number) are continuous functions $\mathbb{S}^n_+ \to \mathbb{R}_+$ it follows that a continuous metric is always locally weakly bounded.

Given a metric G on a Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$, we can define the (oblique) normal cones induced by G as follows:

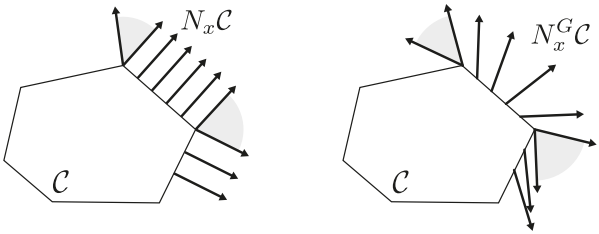

(a) (b)

Figure 3.1: (a) Euclidean normal cones for $\mathcal{C}$; (b) oblique normal cones induced by G.

Definition 3.2. *Let $\mathcal{C} \subset \mathbb{R}^n$ be Clarke regular and let $G : \mathcal{C} \to \mathcal{S}_+^n$ be a metric, then the* normal cone at $x \in \mathcal{C}$ with respect to G *is defined as the polar cone of $T_x^C\mathcal{C}$ with respect to G, i.e.,*

$$N_x^G\mathcal{C} := \left\{ \eta \,\middle|\, \forall v \in T_x^C\mathcal{C} : \langle v, \eta \rangle_{G(x)} \leq 0 \right\} . \tag{3.1}$$

The normal cone with respect to the Euclidean metric $G \equiv \mathbb{I}$ is simply denoted by $N_x\mathcal{C}$. For $x \notin \mathcal{C}$, we adopt the convention $N_x^G\mathcal{C} := \emptyset$.

Note that the Euclidean normal cone and the normal cone with respect to any metric G are isomorphic. Namely, we have

$$\eta \in N_x\mathcal{C} \quad \Longleftrightarrow \quad G^{-1}(x)\eta \in N_x^G\mathcal{C} . \tag{3.2}$$

Figure 3.1 illustrates the Euclidean and oblique normal cones of $\mathcal{C}$. This effect of the metric G will allow us in the forthcoming chapters to define state-dependent oblique projections onto the tangent cones of $\mathcal{C}$.

If the set $\mathcal{C}$ is defined by constraint functions, the oblique normal cone takes an explicit form similar to the expression in Example 2.2:

Example 3.1 (normal cone to constraint-defined sets). Consider $\mathcal{C} := \{x \,|\, h(x) = 0,\, g(x) \leq 0\}$ where $h : \mathbb{R}^n \to \mathbb{R}^m$ and $g : \mathbb{R}^n \to \mathbb{R}^p$ are C^1 and assume that $\mathcal{C}$ satisfies LICQ. Further, let $G : \mathcal{C} \to \mathbb{S}_+^n$ be a metric. Then, the normal cone of $\mathcal{C}$ at x with respect to G is given by

$$N_x\mathcal{C} = \left\{ \eta \,\middle|\, \eta = G^{-1}(x) \left(\nabla h(x)^T \lambda + \nabla g_{\mathbf{I}(x)}(x)^T \mu_{\mathbf{I}} \right),\, \mu_{\mathbf{I}} \in \mathbb{R}_{\geq 0}^{|\mathbf{I}(x)|} \right\} ,$$

where $\mathbf{I}(x)$ denotes the set of active constraints at x. ■

The following result is a generalization of [211, Prop. 6.5] to the case of a continuous metric instead of the standard Euclidean metric:

Lemma 3.1. *Let $\mathcal{C}$ be Clarke regular. If the metric $G : \mathcal{C} \to \mathbb{S}^n_+$ is continuous, then the set-valued map $x \mapsto N^G_x\mathcal{C}$ is outer semicontinuous.*

Proof. Consider any two sequences $x_k \to x$ with $x_k \in \mathcal{C}$ and $\eta_k \to \eta$ with $\eta_k \in N^G_{x_k}\mathcal{C}$. To complete the proof we need to show that $\eta \in N^G_x\mathcal{C}$. By definition of $N^G_{x_k}\mathcal{C}$ we have $\langle v, \eta_k \rangle_{g(x_k)} \leq 0$ for all $v \in T^C_x\mathcal{C}$. Furthermore, by continuity of G we have $\langle v, \eta \rangle_{g(x)} \leq 0$ for all $v \in \limsup_{x_k \to x} T^C_{x_k}\mathcal{C}$. (Namely, we must have $\langle v_k, \eta_k \rangle_{g(x_k)} \leq 0$ for every sequence $v_k \to v$ with $v_k \in T^C_{x_k}\mathcal{C}$, hence the use of lim sup.) By definition of the Clarke tangent cone, we note that $\langle v, \eta \rangle_{g(x)} \leq 0$ holds for all

$$v \in T^C_x\mathcal{C} = \liminf_{x_k \to x} T_{x_k}\mathcal{C} = \liminf_{x_k \to x} T^C_{x_k}\mathcal{C} \subset \limsup_{x_k \to x} T^C_{x_k}\mathcal{C}\,,$$

and therefore $\eta \in N^G_x\mathcal{C}$. □

3.2 Prox-Regular Sets

Throughout the thesis we require the notion of prox-regular sets which are, roughly speaking, closed sets onto which the projection is unique in a neighborhood. Prox-regular sets include closed convex sets (onto which the projection is globally single-valued).

In particular, prox-regular sets turn out to be crucial for (1) the uniqueness of solutions of projected dynamical systems (Section 4.4), (2) anti-windup approximations of projected dynamical systems (Chapter 5), and (3) iterative schemes using constraint linearizations (Chapter 6).

In this subsection we recall the basic definition and properties of prox-regular sets and show that if $\mathcal{C} := \{x \,|\, g(x) \leq 0\}$ satisfies the LICQ at all $x \in \mathcal{C}$ and g is $C^{1,1}$, then $\mathcal{C}$ is prox-regular. The following definition follows [211, Ex. 13.31].

Definition 3.3. *A Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$ is* α-prox-regular at $x \in \mathcal{C}$ *with $\alpha > 0$ if for every $y, z \in \mathcal{C}$ in a neighborhood of x, and every $\eta \in N_z\mathcal{C}$ we have*

$$\langle \eta, y - z \rangle \leq \alpha \|\eta\| \|y - z\|^2 . \tag{3.3}$$

The set $\mathcal{C}$ is prox-regular *if, at every $x \in \mathcal{C}$, it is α-prox-regular for some $\alpha > 0$ (with α possibly depending on x). Further, $\mathcal{C}$ is* α-prox-regular *if it is α-prox-regular at every $x \in \mathcal{C}$ and $\mathcal{C}$ is* uniformly prox-regular *if it is α-prox-regular for some $\alpha > 0$.*

Note that every closed convex set is α-prox-regular for any $\alpha > 0$, since for closed convex sets we have $\langle \eta, y - x \rangle \leq 0$ for any $\eta \in N_x\mathcal{C}$ and all $y \in \mathcal{C}$.

One of the key features of a uniformly prox-regular set $\mathcal{C}$ is that for every point in a neighborhood of $\mathcal{C}$ there exists a unique projection on the set [5, Def. 2.1 & Thm. 2.2].

Proposition 3.1. *If $\mathcal{C} \subset \mathbb{R}^n$ is α-prox-regular, then for any $x \in \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ the set $P_{\mathcal{C}}(x)$ is a singleton and $d_{\mathcal{C}}^2$ is differentiable with $\nabla d_{\mathcal{C}}(x)^T = 2(x - P_{\mathcal{C}}(x))$.*

The following example illustrates how the non-prox-regular sets can fail to have a locally single-valued projection map.

Example 3.2 (Prox-regularity in Euclidean spaces). Consider the parametric set

$$\mathcal{C}_\gamma := \{(x_1, x_2) \,|\, |x_2| \geq \max\{0, x_1\}^\gamma\} , \tag{3.4}$$

where $0 < \gamma < 1$ and which is illustrated in Figure 3.2. For $\gamma \leq 0.5$ the set is (uniformly) prox-regular. In particular for the origin, a ball with non-zero radius can be placed tangentially such that it intersects the set only at 0. For $\gamma > 0.5$ on the other hand the set is not α-prox-regular at the origin for any $\alpha > 0$. In fact, all points on the positive axis have a non-unique projection on $\mathcal{C}_\gamma$ as illustrated in Figure 3.2c. ■

Furthermore, for uniformly prox-regular sets, the pre-image of $P_{\mathcal{C}}$ is spanned by a truncated normal cone, and $P_{\mathcal{C}}$ is locally Lipschitz [5, Thm. 2.2 & Prop. 2.3]:

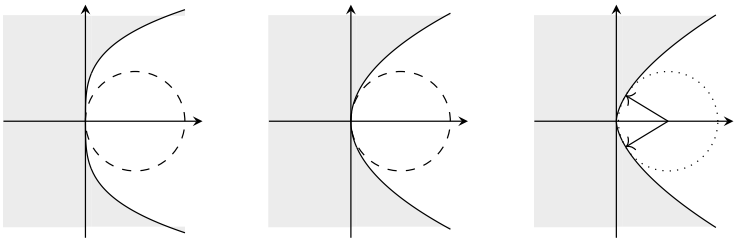

(a) $\gamma = 0.3$ (b) $\gamma = 0.5$ (c) $\gamma = 0.6$

Figure 3.2: $\mathcal{C}_\gamma$ (shaded) for different γ. In (a) and (b) the set $\mathcal{C}_\gamma$ is prox-regular, unlike in (c).

Lemma 3.2. *If $\mathcal{C} \subset \mathbb{R}^n$ is α-prox-regular, then for every $x \in \mathcal{C}$ and all $v \in N_x\mathcal{C} \cap \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ it holds that $P_{\mathcal{C}}(x + v) = x$. In particular, for $y \in \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$, we have $y - P_{\mathcal{C}}(y) \in N_{P_{\mathcal{C}}(y)}\mathcal{C}$.*

Lemma 3.3. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular, then the projection $x \mapsto P_{\mathcal{C}}(x)$ is locally Lipschitz on $\mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$.*

Another crucial property of uniformly prox-regular sets is that the normal cone mapping $x \mapsto N_x\mathcal{C}$ admits a *hypomonotone localization* [211, Ex. 13.38]. We will exploit this property through the following lemma which, in contrast to [211, Ex. 13.38], quantifies the hypomonotonicity in terms of α.

Lemma 3.4. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular. Then, for all $x, x' \in \mathcal{C}$, $\eta \in N_x\mathcal{C} \cap \mathbb{B}$, and $\eta' \in N_{x'}\mathcal{C} \cap \mathbb{B}$, we have*

$$\langle \eta' - \eta, x' - x \rangle \geq -2\alpha \|x' - x\|^2 .$$

Proof. Since $0 \leq \|\eta\| \leq 1$ and $0 \leq \|\eta'\| \leq 1$ it follows from the definition of prox-regularity that

$$\langle \eta, x' - x \rangle \leq \alpha \|x' - x\|^2 \quad \text{and} \quad \langle -\eta', x' - x \rangle \leq \alpha \|x' - x\|^2 .$$

Adding up both inequalities yields the desired result. □

We conclude this section by showing that feasible domains defined by $C^{1,1}$ constraint functions are prox-regular under the usual constraint qualifications (see also Example 2.2). For this, we require the well-known Descent Lemma, e.g. [198, Lem. 1.30].

Lemma 3.5 (Descent Lemma)**.** *Let $f : \mathcal{V} \to \mathbb{R}$ be $C^{1,1}$ and $\mathcal{V} \subset \mathbb{R}^n$ is open. Given $x \in \mathcal{V}$, there exists $\ell > 0$ such that for all $z, y \in \mathcal{V}$ in a neighborhood of x it holds that*

$$|f(z) - f(y) - D_y f(z - y)| \leq \ell \|z - y\|^2 .$$

Proposition 3.2 (Prox-regularity of constraint-defined sets)**.** *Let $g : \mathbb{R}^n \to \mathbb{R}^m$ be $C^{1,1}$ and let $\mathcal{C} := \{x \,|\, g(x) \leq 0\}$ satisfy the LICQ at all $x \in \mathcal{C}$. Then, $\mathcal{C}$ is prox-regular.*

Proof. Given $\hat{x} \in \mathcal{C}$, consider a compact neighborhood $\mathcal{N} \subset \mathcal{C}$ of $\hat{x}$ such that $\mathbf{I}(x) \subset \mathbf{I}(\hat{x})$ for all $x \in \mathcal{N}$, i.e., no other constraint is active on $\mathcal{N}$ other than those active at $\hat{x}$. Clearly, such a neighborhood exists around every $\hat{x} \in \mathcal{C}$. Further, if necessary, reduce the size of the neighborhood to guarantee that $\|\nabla g_i(x)\| > \delta_i$ for some $\delta_i > 0$, all $x \in \mathcal{N}$ and all $i \in \mathbf{I}(\hat{x})$. This is possible since $\nabla g_i(\hat{x}) \neq 0$ for all $i \in \mathbf{I}(\hat{x})$ by LICQ and ∇g_i is continuous.

Using Lemma 3.5, there exists $\ell_i > 0$ such that for all $z, y \in \mathcal{N}$ we have

$$-\ell_i \|z - y\|^2 \leq g_i(z) - g_i(y) - \nabla g_i(y)(z - y) .$$

In particular, for all $i \in \mathbf{I}(\hat{x})$ and all y such that $g_i(y) = 0$ we have

$$\nabla g_i(y)(z - y) \leq g_i(z) + \ell_i \|z - y\|^2 \leq \ell_i \|z - y\|^2 \tag{3.5}$$

since $z \in \mathcal{N} \subset \mathcal{C}$ and thus $g_i(z) \leq 0$.

Recall from Example 3.1 that for $y \in \mathcal{C}$, under LICQ, we have

$$N_y \mathcal{C} = \left\{ \eta \,\middle|\, \eta = \sum\nolimits_{i \in \mathbf{I}(y)} \gamma_i \nabla g_i^T(y),\ \gamma_i \geq 0 \right\} .$$

In particular, since $\mathbf{I}(y) \subset \mathbf{I}(\hat{x})$ and thus $\|\nabla g_i(y)\| \geq \delta_i$, we have that for $\eta \in N_y \mathcal{C} \cap \mathbb{B}$ it holds that $\gamma_i \leq 1/\delta_i$.

Let $\eta \in N_y\mathcal{C} \cap \mathbb{B}$. Using (3.5), we have for all $z \in \mathcal{C}$

$$\begin{aligned}\langle \eta, z-y\rangle &= \left\langle \sum\nolimits_{i\in\mathbf{I}(y)} \gamma_i \nabla g_i(y)^T, z-y \right\rangle \\ &\leq \left(\sum\nolimits_{i\in\mathbf{I}(y)} \gamma_i \ell_i\right) \|z-y\|^2 \leq \underbrace{\sum\nolimits_{i\in\mathbf{I}(\hat{x})} \frac{\ell_i}{\delta_i}}_{=:\alpha(\hat{x})} \|z-y\|^2 .\end{aligned}$$

Thus, $\mathcal{C}$ is $\alpha(\hat{x})$-prox-regular at $\hat{x}$ according to Definition 3.3. □

3.3 Non-Euclidean Prox-Regular Sets

Definition 3.3 cannot be directly generalized to non-Euclidean spaces since it requires the distance $\|y-x\|$ between two points in $\mathcal{C}$. Hence, in [133] prox-regularity is defined on smooth (i.e., C^∞) Riemannian manifolds and $\|y-x\|$ is replaced with the geodesic distance between y and x. For our purposes we can avoid the notational complexity of Riemannian geometry, yet preserve a higher degree of generality. We introduce the following definitions:

Definition 3.4. *Given a Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$ and a metric $G : \mathcal{C} \to \mathbb{S}^n_+$, a normal vector $\eta \in N_x^G\mathcal{C}$ at $x \in \mathcal{C}$ is* α-proximal in G *for $\alpha > 0$ if, for all $y \in \mathcal{C}$ in a neighborhood of x, we have*

$$\langle \eta, y-x\rangle_{G(x)} \leq \alpha \|\eta\|_{G(x)} \|y-x\|^2_{G(x)} . \tag{3.6}$$

The cone of all α-proximal normal vectors at x in G is denoted by $\bar{N}_x^{G,\alpha}\mathcal{C}$.

A crucial detail in (3.6) is the fact that G is evaluated at x and is used as an inner product on $\mathbb{R}^n$. This is a slight abuse of notation because from a conceptual viewpoint G defines a bilinear form between (tangent) vectors at x, i.e., vectors in $T_x\mathbb{R}^n$. However, we have that $y - x \in \mathbb{R}^n$ and, to use $G(x)$ as an inner product on $\mathbb{R}^n$, we exploit the canonical isomorphism between $\mathbb{R}^n$ and $T_x\mathbb{R}^n$.

Definition 3.5. *A Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$ with a metric $G : \mathcal{C} \to \mathbb{S}^n_+$ is* α-prox-regular at $x \in \mathcal{C}$ in G *if $\bar{N}_y^{G,\alpha}\mathcal{C} = N_y^G\mathcal{C}$ for all $y \in \mathcal{C}$ in a*

neighborhood of x. The set $\mathcal{C}$ is α-prox-regular in G *if it is α-prox-regular in G at every $x \in \mathcal{C}$ and $\mathcal{C}$ is* prox-regular in G *if there exists $\alpha > 0$ such that $\mathcal{C}$ is α-prox-regular in G.*

Note that if G is the Euclidean metric, Definition 3.5 reduces to Definition 3.3. The following result shows that prox-regularity is in fact independent of the metric. This is the first step towards a coordinate-free characterization of prox-regularity.

Proposition 3.3. *Let $\mathcal{C} \subset \mathbb{R}^n$ be Clarke regular. If $\mathcal{C}$ is prox-regular in a continuous metric $G : \mathcal{C} \to \mathbb{S}^n_+$, then it is prox-regular in any other continuous metric.*

In particular if $\mathcal{C}$ is prox-regular in the Euclidean metric, i.e., according to Definition 3.3, then it is prox-regular in any other continuous metric on $\mathbb{R}^n$. For the proof of Proposition 3.3 we require the following lemma which establishes the pointwise correspondence of proximal normal cones.

Lemma 3.6. *Let $\mathcal{C} \subset \mathbb{R}^n$ be Clarke regular and consider two metrics $G, G' : \mathcal{C} \to \mathbb{S}^n_+$. Given $x \in \mathcal{C}$, if $\bar{N}_x^{G,\alpha}\mathcal{C} = N_x^G\mathcal{C}$ holds for some $\alpha > 0$ then $\bar{N}_x^{G',\alpha'}\mathcal{C} = N_x^{G'}\mathcal{C}$ holds for $\alpha' \geq \kappa_{G(x)}\kappa_{G'(x)}\alpha$.*

Proof. First, note that for every $x \in \mathcal{C}$ the two metrics G and G' induce a bijection between $N_x^G\mathcal{C}$ and $N_x^{G'}\mathcal{C}$ (see (3.2)). Namely, we define $\psi : T_x\mathbb{R}^n \to T_x\mathbb{R}^n$ as the unique element $\psi(v)$ that satisfies by $\langle v, w\rangle_{G(x)} = \langle \psi(v), w\rangle_{G'(x)}$ for all $w \in T_x\mathbb{R}^n$. To clarify, in matrix notation we can write $v^T G(x) w = \psi(v)^T G'(x) w$ and since $G(x), G'(x)$ are symmetric positive definite we have $\psi(v) := G'(x)^{-1}G(x)v$. It follows that if $\eta \in N_x^G\mathcal{C}$ (hence, by definition $\langle \eta, w\rangle_{G(x)} \leq 0$ for all $w \in T_x\mathcal{C}$), then $\psi(\eta) \in N_x^{G'}\mathcal{C}$. Furthermore, omitting the argument x, we have $\|\psi(\eta)\|_{G'} = \eta^T G G'^{-1} G\eta \geq 1/\lambda_{G'}^{\max}\|G\eta\|$ and $\|\eta\|_G = \eta^T G G^{-1} G\eta \leq 1/\lambda_G^{\min}\|G\eta\|$, and therefore $\|\psi(\eta)\|_{G'(x)} \geq \lambda_{G(x)}^{\min}/\lambda_{G'(x)}^{\max}\|\eta\|_{G(x)}$.

Hence, let $\eta \in N_x^G\mathcal{C} \setminus \{0\}$ be a α-proximal normal vector, then

$$\left\langle \frac{\psi(\eta)}{\|\psi(\eta)\|_{G'(x)}}, y - x \right\rangle_{G'(x)} \leq \frac{\lambda_{G'(x)}^{\max}}{\lambda_{G(x)}^{\min}} \left\langle \frac{\eta}{\|\eta\|_{G(x)}}, y - x \right\rangle_{G(x)} \leq \frac{\lambda_{G'(x)}^{\max}}{\lambda_{G(x)}^{\min}} \alpha \|y - x\|^2_{G(x)}\,.$$

Finally, using the equivalence of norms, we have

$$\frac{\lambda^{\max}_{G'(x)}}{\lambda^{\min}_{G(x)}}\alpha\|y-x\|^2_{G(x)} \leq \frac{\lambda^{\max}_{G'(x)}}{\lambda^{\min}_{G(x)}}\frac{\lambda^{\max}_{G(x)}}{\lambda^{\min}_{G'(x)}}\alpha\|y-x\|^2_{G'(x)} \leq \alpha'\|y-x\|^2_{G'(x)}\,, \quad (3.7)$$

where $\alpha' \geq \kappa_{G(x)}\kappa_{G'(x)}\alpha$. Thus, we have shown that if $v \in \bar{N}^{G,\alpha}_x\mathcal{C} = N^G_x\mathcal{C}$ then $\psi(v) \in \bar{N}^{G',\alpha'}_x = N^{G'}_x\mathcal{C}$ which completes the proof. □

Proof of Proposition 3.3. Since G and G' are continuous it follows that $\kappa_{G(x)}$ and $\kappa_{G'(x)}$ are continuous in x and therefore locally bounded. Given any $x \in \mathcal{C}$ and using the pointwise result in Lemma 3.6, we can choose $\alpha' > 0$ such that (3.7) is satisfied for all $y \in \mathcal{C}$ in a neighborhood of x, thus satisfying Definition 3.5 for α'-prox regularity at x. Since x is arbitrary, we conclude that $\mathcal{C}$ is prox-regular (albeit not uniformly, unless $\kappa_{G(x)}$ and $\kappa_{G'(x)}$ are bounded). □

Thus, we have shown that prox-regularity as in Definition 3.5 is an inherent property of a set $\mathcal{C} \subset \mathbb{R}^n$ independent of the metric (as long as the metric is continuous). Therefore, we will refer to a set simply as being *prox-regular*, rather than being *prox-regular in* G. Also note that this insight extends the scope of Proposition 3.2.

Finally, we show that prox-regularity is preserved under sufficiently smooth coordinate transformations. This will be key in Section 4.5 where we show coordinate invariance of projected dynamical systems.

Proposition 3.4. *Let $\mathcal{V}, \mathcal{W} \subset \mathbb{R}^n$ be open and consider a $C^{1,1}$ diffeomorphism $\phi : \mathcal{V} \to \mathcal{W}$, i.e., ϕ is a $C^{1,1}$ map and has a $C^{1,1}$ inverse $\phi^{-1} : \mathcal{W} \to \mathcal{V}$. If $\mathcal{C} \subset \mathcal{V}$ is prox-regular then the image $\phi(\mathcal{C})$ is prox-regular.*

Proof. By Proposition 3.3 it suffices to show prox-regularity in a single metric on $\mathcal{V}$ and $\mathcal{W}$ respectively. Hence, let $\mathcal{W}$ be endowed with the canonical Euclidean metric which we denote explicitly by E, i.e., $E : \mathcal{W} \to \mathbb{S}^n_+$ with $E \equiv \mathbb{I}$. Next, let E^* denote the so-called *pull-back metric* on $\mathcal{V}$ along ϕ, i.e., E^* is defined such that $\langle v, w\rangle_{E^*(x)} := \langle D_x\phi(v), D_x\phi(w)\rangle$ (or, more explicitly, $E^\star : x \mapsto \nabla\phi(x)\nabla\phi(x)^T$).

Analogously to Lemma 2.2, we show that (proximal) normal cones are preserved under $C^{1,1}$ coordinate transformations, i.e., we prove that

$$\eta \in N_x^{E^*}\mathcal{C} \iff D_x\phi(\eta) \in N_{\phi(x)}\phi(\mathcal{C}) \qquad \forall x \in \mathcal{C} \tag{3.8}$$

$$\eta \in \bar{N}_y^{E^*,\alpha}\mathcal{C} \iff D_y\phi(\eta) \in \bar{N}_{\phi(y)}^{\alpha'}\phi(\mathcal{C}) \qquad \forall y \in \mathcal{N}_x \tag{3.9}$$

for some $\alpha', \alpha > 0$ where $\mathcal{N}_x \subset \mathcal{C}$ is a neighborhood of x. Since ϕ is a diffeomorphism it suffices to show one direction only.

Hence, we first consider a (standard) normal vector $\eta \in N_x^{E^*}\mathcal{C}$. By Definition 3.2 of the normal cone and using Lemma 2.2 we have

$$\begin{aligned} \eta \in N_x^{E^*}\mathcal{C} \quad &\Leftrightarrow \quad \langle \eta, w\rangle_{E^*(x)} = \langle D_x\phi(\eta), D_x\phi(w)\rangle \leq 0 \quad \forall w \in T_x\mathcal{C} \\ &\Leftrightarrow \quad \langle D_x\phi(\eta), \tilde{w}\rangle \leq 0 \quad \forall \tilde{w} \in D_x\phi(T_x\mathcal{C}) = T_{\phi(x)}\phi(\mathcal{C})\,. \end{aligned}$$

We conclude that $D_x\phi(\eta) \in N_{\phi(x)}\phi(\mathcal{C})$ and (3.8) holds.

Next, for (3.9) we consider $y \in \mathcal{C}$ in a neighborhood of x and a α-proximal normal vector $\eta \in \bar{N}_y^{E^*,\alpha}\mathcal{C}$ at y such that, by definition,

$$\langle \eta, z - y\rangle_{E^*(y)} = \langle D_y\phi(\eta), D_y\phi(z-y)\rangle \leq \alpha\|z-y\|^2_{E^*G(y)}$$

holds for all $z \in \mathcal{C}$ in a neighborhood of y. However, we need to show that for some $\alpha' > 0$ we have

$$\langle D_y\phi(\eta), \phi(z) - \phi(y)\rangle \leq \alpha'\|\phi(z) - \phi(y)\|^2\,. \tag{3.10}$$

We define the $C^{1,1}$ function $\psi : \mathcal{V} \to \mathbb{R}$ as $\psi(z) := \langle D_y\phi(\eta), \phi(z)\rangle$ and note that by linearity we have $D_z\psi(v) := \langle D_y\phi(\eta), D_z\phi(v)\rangle$. This enables us to apply the Descent Lemma (Lemma 3.5) and state that for some $M > 0$ it holds that

$$\begin{aligned} |\psi(z) - \psi(y) - D_y\psi(z-y)| &= \underbrace{|\,\langle D_y\phi(\eta), \phi(z) - \phi(y) - D_y\phi(z-y)\rangle\,|}_{=:\gamma(z)} \\ &\leq M\|z-y\|^2\,. \end{aligned}$$

This bound can be used to establish

$$\begin{aligned} \langle D_y\phi(\eta), \phi(z) - \phi(y)\rangle &\leq \langle D_y\phi(\eta), D_y\phi(z-y)\rangle + \gamma(z) \\ &\leq (\alpha + M)\|z-y\|^2\,. \end{aligned}$$

Finally note that $\|z-y\|^2 \leq \ell\|\phi(z)-\phi(y)\|^2$ for some $\ell > 0$ since ϕ^{-1} is locally Lipschitz continuous. Hence, (3.10) and therefore (3.9) holds for $\alpha' = \ell(\alpha + M)$. □

In summary, in this section, we have proposed a new coordinate-free and intrinsic definition of prox-regular sets which generalizes the standard definition of prox-regularity for sets. For this purpose, we have defined prox-regularity with respect to a (Riemannian) metric, but then we have shown that the choice of metric does not affect the prox-regularity. The key takeaway is Proposition 3.4 which states that prox-regularity is preserved by $C^{1,1}$ coordinate transformations.

Note that Propositions 3.3 and 3.4 do not make statements about *uniform prox-regularity* of $\mathcal{C}$. However, it is easy to see that Proposition 3.3 can be adapted to uniformly prox-regular sets by considering only globally weakly bounded metrics, i.e., metrics with bounded conditions numbers. Similarly, Proposition 3.4 can be reformulated by considering a $C^{1,1}$ diffeomorphism ϕ with globally Lipschitz $\nabla\phi$ and $\nabla\phi^{-1}$. In that case uniform prox-regularity is preserved by the coordinate transformation ϕ (although the modulus of prox-regularity may change).

3.4 Notes & Comments

The definition of a metric as a variable inner product is at the center of classical Riemannian geometry. In that context, a Riemannian metric is formally defined as symmetric, positive definite bilinear form defined over the tangent bundle of a manifold [161].

Most work in differential geometry deals with sufficiently smooth objects and most standard textbooks limit themselves to C^∞ manifolds and metrics. The reason for this seems to be that many sophisticated tools are available only under sufficient smoothness (see e.g. [124] for a counter-example where geodesics are not well-defined on a C^2 manifold).

For our purposes, these advanced tools are not required. We primarily use a variable metric to induce oblique projection directions in Part I

and adapt the notion of *steepest descent* for gradient descent algorithms in Part II. Hence, from an engineering perspective, a variable metric is, to a large extent, a design parameter.

Historically, prox-regularity has first been established and studied for functions (in the sense that prox-regular functions are those functions with a prox-regular epigraph) [199]. The notion of prox-regular sets is, however, of great importance. Particularly for the uniqueness of solutions of projected dynamical systems and sweeping processes (i.e., time-varying projected dynamical systems; see also Chapter 12) prox-regularity is a key component which was identified and studied in [68] and others. We will revisit the uniqueness results for projected dynamical systems in Section 4.4 and generalize them to oblique projection directions.

Various definitions of prox-regularity (either via the locally single-valued projections or hypomonotonicity of unit normal vectors) have been proposed and shown to be equivalent [5]. The recent papers [28, 133] have also extended the concept of prox-regularity to C^∞ Riemannian manifolds by imposing that prox-regular subsets on a manifold need to admit a locally single-valued projection (in geodesic terms).

The results presented in this chapter go in a different direction: We have shown that prox-regularity is an inherent property of a subset of $\mathbb{R}^n$ and preserved by sufficiently smooth coordinate transformations. This is important to motivate the modeling of physical processes, which are inherently coordinate-free, by projected dynamical systems.

Finally, we note that Proposition 3.2 provides an important link between the geometry and analysis of prox-regular sets, essentially extending Example 2.2 where we gave explicit descriptions of tangent and normal cones for Clarke regular sets.

Part I

Projected Dynamical Systems and their Feedback Implementation

CHAPTER 4

Coordinate-Free Projected Dynamical Systems

We now rigorously define *projected dynamical system*s (PDSs) and establish basic existence and uniqueness results for their trajectories.

Our goal in this chapter is to establish a "geometric" model of PDSs that is coordinate-free. In other words, the properties of PDSs should be invariant under (nonlinear) coordinate transformations. This feature is highly desirable, if not essential, to motivate the use of PDS to model physical feedback control systems.

As a specific example, the laws of AC power flow that govern the physical flow of power in an electricity grid at steady state can be formulated in different coordinates. This fact will be discussed in more detail in Part III. While the design and implementation of a control scheme for this kind of system is generally based on a specific choice of coordinates, it is judicious to require that any qualitative properties such as well-posedness, stability, and convergence hold irrespective of a

This chapter presents the key results of [Ha2] in a slightly simplified form. Namely, we develop existence and uniqueness results for PDS over *closed* instead of *locally compact* domains. This yields a more concise presentation in Sections 4.1, 4.2, 4.3, and 4.4 at the expense of a less rigorous treatment of PDS on manifolds in Section 4.5. Corollaries 4.1 and 4.3 and Theorems 4.2 and 4.3 constitute the main results of this chapter.

specific numerical representation.

To achieve such a geometric viewpoint we use (Riemannian) metrics, as introduced in Chapter 3, to describe oblique projection directions that are preserved along sufficiently smooth coordinate transformations.

Besides the use of a variable metric, we take a different approach in showing existence compared to the classical works [17, 68, 127] and the popular, albeit slightly unconventional book [186]. Namely, we start by considering a Krasovskii-regularized inclusion, which allows us to generalize PDSs beyond Clarke regular sets.

For our purposes, a PDS is defined as a differential inclusion of the form

$$\dot{x} \in \Pi_{\mathcal{C}}^{G}\left[f(x)\right](x) \qquad x \in \mathcal{C}\,, \tag{4.1}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is a closed set, $G : \mathcal{C} \to \mathbb{S}_+^n$ is a metric, and $f : \mathcal{C} \to \mathbb{R}^n$ is a vector field. Given $x \in \mathcal{C}$ and $w \in \mathbb{R}^n$, the operator $\Pi_{\mathcal{C}}^{G}[w](x)$ projects w onto the tangent cone of $\mathcal{C}$ at x with respect to the metric G, i.e.,

$$\Pi_{\mathcal{C}}^{G}[w](x) := \operatorname*{arg\,min}_{v \in T_x\mathcal{C}} \|v - w\|_{G(x)}\,. \tag{4.2}$$

Since $T_x\mathcal{C}$ is non-empty for any $x \in \mathcal{C}$ (namely, $0 \in T_x\mathcal{C}$) and closed, a minimum norm projection exists, and therefore $\Pi_{\mathcal{C}}^{G}[f](x)$ is non-empty for all $x \in \mathcal{C}$.[1] If $x \notin \mathcal{C}$, we have $\Pi_{\mathcal{C}}^{G}[w](x) = \emptyset$ and therefore $\operatorname{dom} \Pi_{\mathcal{C}}^{G}[w] = \mathcal{C}$ for all $w \in \mathbb{R}^n$. If f is a vector field, we abuse notation and write $\Pi_{\mathcal{C}}^{G}[f](x) := \Pi_{\mathcal{C}}^{G}[f(x)](x)$ for brevity. For simplicity, we call $\Pi_{\mathcal{C}}^{G}[f]$ a *vector field* even though $\Pi_{\mathcal{C}}^{G}[f](x)$ is, in general, not a singleton for every $x \in \mathcal{C}$. However, if $\mathcal{C}$ is Clarke regular and therefore $T_x\mathcal{C}$ is convex for all x, then $\Pi_{\mathcal{C}}^{G}[f](x)$ is a singleton for all $x \in \mathcal{C}$ (note that $\|v - f(x)\|_{G(x)}^2$ is strictly convex as function of v). In this case, we can consider the *ordinary differential equation* (ODE) $\dot{x} = \Pi_{\mathcal{C}}^{G}[f](x)$, $x \in \mathcal{C}$ instead of the inclusion (4.1).

The key properties of $\Pi_{\mathcal{C}}^{G}[\cdot]$ are summarized in the following lemma which is a simple application of Moreau's Decomposition Theorem [131, Thm. 3.2.5].

[1] See, e.g., the first part of the proof of Hilbert's projection theorem [198, Prop. 1.37].

Lemma 4.1. *Given $\mathcal{C} \subset \mathbb{R}^n$, $x \in \mathcal{C}$, a metric $G : \mathcal{C} \to \mathbb{S}^n_+$, and $w \in \mathbb{R}^n$, for any $v \in \Pi^G_{\mathcal{C}}[w]$ one has $\langle w, v \rangle_{G(x)} = \|v\|^2_{G(x)}$. If in addition $\mathcal{C}$ is Clarke regular at x, then $\Pi^G_{\mathcal{C}}[w](x)$ is a singleton and there exists a unique $\hat{\eta} \in N^G_x\mathcal{C}$ such that the following equivalent statements hold:*

(i) $\Pi^G_{\mathcal{C}}[w](x) = w - \hat{\eta}$,

(ii) $\arg\min_{\eta \in N^G_x\mathcal{C}} \|\eta - w\|_{G(x)} = \hat{\eta}$, *and*

(iii) $w - \hat{\eta} \in T_x\mathcal{C}$ *and* $\langle x - \hat{\eta}, \hat{\eta} \rangle_{G(x)} = 0$.

Proof. Let $v \in \Pi^G_{\mathcal{C}}[f](x)$. As $T_x\mathcal{C}$ is a cone we have $\lambda v \in T_x\mathcal{C}$ for all $\lambda \geq 0$. Since v (locally) minimizes $\|v - f(x)\|^2_{G(x)}$ over $T_x\mathcal{C}$, it follows that $\lambda = 1$ minimizes $M(\lambda) := \frac{1}{2}\|\lambda v - f(x)\|^2_{G(x)}$ for v fixed. Hence, for $\lambda = 1$ the optimality condition $\frac{dM}{d\lambda}(\lambda) = \lambda \langle v - f(x), v \rangle_{G(x)} = 0$ holds. This proves the first part. The second part follows from Moreau's Theorem [131, Thm. 3.2.5] since $T_x\mathcal{C}$ is convex by Clarke regularity. □

Recall from Example 2.2 that the tangent cone $T_x\mathcal{C}$ for a given $x \in \mathcal{C}$ can be expressed explicitly as a polyhedral cone whenever the set $\mathcal{C}$ is defined by differentiable constraints and satisfies the LICQ. In this case, the projection onto $T_x\mathcal{C}$ as given by $\Pi^G_{\mathcal{C}}[\cdot]$ takes the form of a quadratic program. This fact that $\Pi^G_{\mathcal{C}}[f](x)$ can be computed numerically will be exploited in Chapter 6 to develop an iterative scheme to implement PDSs as feedback control loops.

Example 4.1 (pointwise evaluation of a projected vector field). As in Example 3.1 let $g : \mathbb{R}^n \to \mathbb{R}^m$ be C^1, let $\mathcal{C} := \{x \,|\, g(x) \leq 0\}$ satisfy the LICQ for all $x \in \mathcal{C}$, and let $G : \mathcal{C} \to \mathbb{S}^+_n$ a metric. Furthermore, consider a vector field $f : \mathcal{C} \to \mathbb{R}^n$. Then, the projected vector field $\Pi^G_{\mathcal{C}}[f](x)$ at $x \in \mathcal{C}$ is given as the solution of the convex quadratic program

$$\underset{v \in \mathbb{R}^n}{\text{minimize}} \quad (f(x) - v)^T G(x)(f(x) - v) \qquad \text{subject to} \quad \nabla h_{\mathbf{I}(x)}(x) v \leq 0\,.$$

Note that x is not an optimization variable. Hence, the properties of f and G as functions of x are irrelevant when doing a pointwise evaluation of $\Pi^G_{\mathcal{C}}[f](x)$. ■

Remark 4.1. Projected dynamical systems can be generalized to the case where the (unprojected) vector field $f : \mathcal{C} \to \mathbb{R}^n$ is replaced with a set-valued map $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$. This avenue has been explored in [18, 17, 68, 127], albeit only for the Euclidean metric $G \equiv \mathbb{I}$ and Clarke regular sets $\mathcal{C}$. For our purposes, we assume that f is single-valued, in order not to overload the forthcoming sections with technicalities. ■

4.1 Existence of Krasovskii Solutions

Existence of solutions to PDSs of the form (4.1) is not covered by Theorem 2.4 (viability theorem) because $x \mapsto \Pi_{\mathcal{C}}^{G}[f](x)$ is in general neither convex nor outer semicontinuous.

In previous work such as [68, 127] as well as [18, Ch. 10.1] or [17, Ch. 6], existence of solutions to (4.1) has been shown under the additional assumption of $\mathcal{C}$ being Clarke regular or even convex. These works have shown, in particular, that *slow*[2] solutions of the so-called differential variational inequalities $\dot{x} \in f(x) - N_x\mathcal{C}$ are equivalent to solutions of (4.1) (assuming the Euclidean metric $G \equiv \mathbb{I}$; see also Section 4.3).

In this thesis we take a different approach by considering a *Krasovskii regularization* of the PDS (4.1) which is well-posed almost by definition and which does not rely on a reformulation as a DVI. As a consequence, we do not require the use of normal cones and can therefore deal with general domains $\mathcal{C}$. Yet, for Clarke regular sets $\mathcal{C}$, we recover results that are analogous to the results in [18, 17, 68, 127], albeit generalized to arbitrary metrics.

Definition 4.1. *Given* $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, *the* Krasovskii regularization of H *is defined as*

$$\mathrm{K}[H] : \mathbb{R}^n \rightrightarrows \mathbb{R}^n \qquad x \mapsto \overline{\mathrm{co}} \limsup_{y \to x} H(y)\,.$$

By definition, the Krasovskii regularization $\mathrm{K}[H]$ of H takes only closed convex values. In addition, it is outer semicontinuous and preserves local

[2] A solution $x : [0, T] \to \mathbb{R}^n$ to an inclusion $\dot{x} \in F(x)$ is *slow* if $\|\dot{x}(t)\| = \min\{\|v\| \,|\, v \in F(x(t))\}$ holds for almost all $t \in [0, T]$

boundedness as the following lemma corresponding to [118, Lem. 5.16] shows:

Lemma 4.2. *Given $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, the Krasovskii regularization of H is outer semicontinuous and, if H is locally bounded, then $\mathrm{K}[H]$ is locally bounded.*

Hence, a *Krasovskii solution* to the differential inclusion $\dot{x} \in H(x)$, $x \in \mathcal{C}$ is simply a solution $x : [0, T] \to \mathcal{C}$ to the Krasovskii-regularized inclusion

$$\dot{x} \in \mathrm{K}[H](x), \qquad x \in \mathcal{C}. \tag{4.3}$$

In other words, a solution of (4.3) is a Krasovskii solution of $\dot{x} \in H(x)$, $x \in \mathcal{C}$.

The Krasovskii regularization being well-posed almost by definition, we can state the following existence result about Krasovskii solutions.

Corollary 4.1 (Existence of Krasovskii solutions)**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be closed, $f : \mathcal{C} \to \mathbb{R}^n$ a locally bounded vector field, and $G : \mathcal{C} \to \mathbb{R}^n$ a locally weakly bounded metric. Then, for any initial condition $x(0) \in \mathcal{C}$, the PDS (4.1) admits a Krasovskii solution $x : [0, T] \to \mathcal{C}$ with either $T = \infty$ or $T < \infty$ and $\limsup_{t \to T^-} \|x(t)\| = \infty$.*

Corollary 4.1 is slightly less general than [Ha12, Thm. 4.2] because $\mathcal{C}$ needs to be closed. Assuming the closedness of $\mathcal{C}$ allows for a more concise presentation of the following proof and the proof of the forthcoming Corollary 4.2. However, Corollary 4.1, is strictly speaking, not general enough to establish existence on abstract manifolds (Section 4.5). In this case, $\mathcal{C}$ is in general only locally compact (which is covered by [Ha12, Thm. 4.2]).

Proof. We apply Theorem 2.4. First, note that $\operatorname{dom} \Pi_{\mathcal{C}}^{G}[f] = \mathcal{C}$ since $T_x\mathcal{C} = \emptyset$ for all $x \notin \mathcal{C}$, by convention. Because $\mathcal{C}$ is closed we also have $\operatorname{dom} \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right] = \mathcal{C}$. Given $x \in \mathcal{C}$, for all $v \in \Pi_{\mathcal{C}}^{G}[f](x) \subset \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$ we have $v \in T_x\mathcal{C}$ (by definition of $\Pi_{\mathcal{C}}^{G}[\cdot]$). Therefore, we have $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x) \cap T_x\mathcal{X} \neq \emptyset$ for all $x \in \mathcal{C}$ and $\mathcal{C}$ is a viability domain for $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$.

To show that the regularized inclusion $\dot{x} \in \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$, $x \in \mathcal{C}$ is well-posed, note that $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$ is non-empty for all $x \in \mathcal{C}$ (since $\emptyset \neq \Pi_{\mathcal{C}}^{G}[f](x) \subset \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$ for all $x \in \mathcal{C}$) and $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$ is convex by definition. Furthermore, by Lemma 4.2, $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$ is osc. It remains to show that $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$ is locally bounded which follows immediately by Lemma 4.2 if $\Pi_{\mathcal{C}}^{G}[f]$ is locally bounded.

To show that $\Pi_{\mathcal{C}}^{G}[f]$ is locally bounded, we first introduce an auxiliary metric $\hat{G}$ defined as $\hat{G}(x) := \frac{1}{\lambda_{G}^{\max}(x)} G(x)$, that is, we scale the metric at every $x \in \mathcal{C}$ by dividing it by its maximum eigenvalue at that point. This implies that $\|f(x)\|_{\hat{G}(x)} \leq \|f(x)\|$ for all $x \in \mathcal{C}$. Note that the projected vector field is unchanged, i.e., $\Pi_{\mathcal{C}}^{\hat{G}}[f] = \Pi_{\mathcal{C}}^{G}[f]$, since in (4.2) only the objective function is scaled. Furthermore, $\kappa_{G}(x) = \kappa_{\hat{G}}(x)$ for all $x \in \mathcal{C}$, and consequently $\hat{G}$ is locally weakly bounded since G is locally weakly bounded.

Given any $x \in \mathcal{C}$, we have $0 \in T_x\mathcal{C}$. It follows that $\|v\|_{\hat{G}(x)} \leq \|f(x) - 0\|_{\hat{G}(x)}$ for every $v \in \Pi_{\mathcal{C}}^{\hat{G}}[f](x)$. Consequently, by local boundedness of f there exists $\ell'' > 0$ such that $\|\Pi_{\mathcal{C}}^{\hat{G}}[f](y)\|_{\hat{G}(y)} \leq \ell''$ for every $y \in \mathcal{C}$ in a neighborhood of x. Furthermore, by weak local boundedness of $\hat{G}$ there exists $\ell' > 0$ such that $\kappa_{\hat{G}}(x) \leq \ell'$ in a neighborhood of x. Since $\lambda_{\hat{G}}^{\max}(x) = 1$, it follows that $\lambda_{G(x)}^{\max} \geq 1/\ell'$ and therefore $\|v\| \leq \ell' \|v\|_{G(y)}$ for all $v \in T_y\mathbb{R}^n$ and all $y \in \mathcal{C}$ in a neighborhood of x. Combining these arguments, there exist $\ell', \ell'' > 0$ such that for every $y \in \mathcal{C}$ in a neighborhood of x it holds that

$$\tfrac{1}{\ell'} \|\Pi_{\mathcal{C}}^{\hat{G}}[f](y)\| \leq \|\Pi_{\mathcal{C}}^{\hat{G}}[f](y)\|_{\hat{G}(y)} \leq \|f(y)\|_{\hat{G}(y)} \leq \|f(y)\| \leq \ell''. \quad (4.4)$$

Hence, since $\Pi_{\mathcal{C}}^{\hat{G}}[f] = \Pi_{\mathcal{C}}^{G}[f]$, it follows that $\Pi_{\mathcal{C}}^{G}[f]$ is locally bounded.

Consequently, the regularized inclusion $\dot{x} \in \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$, $x \in \mathcal{C}$ is well-posed and $\mathcal{C}$ is a viability domain for $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$. Theorem 2.4 can thus be applied to guarantee the existence of Krasovskii solutions to (4.1) which completes the proof. □

The next result establishes conditions for the existence of complete solutions.

Corollary 4.2. *Consider the same setup as in Corollary 4.1. If either*

(i) f is bounded, and G is weakly bounded, or

(ii) $\mathcal{C}$ is compact, f and G are continuous, or

(iii) f is globally Lipschitz and G is weakly bounded,

then for every $x(0) \in \mathcal{C}$ there exists a complete Krasovskii solution to (4.1).

Proof. According to Theorem 2.4, the linear growth condition

$$\left\|\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)\right\| \leq c(\|x\|+1)$$

for some $c > 0$ guarantees the existence of complete solutions.

(i) If f is bounded and G is weakly bounded, then the local boundedness argument of the proof of Corollary 4.1 can be applied globally, i.e., (4.4) holds for all $y \in \mathcal{C}$ for the same ℓ', ℓ'' and hence $\Pi_{\mathcal{C}}^{G}[f]$ is bounded and (trivially) has linear growth.

(ii) Since f is continuous it takes only bounded values on a compact set. Furthermore, continuity of G implies local weak boundedness, i.e., for every $x \in \mathcal{C}$ there exist $\ell_x, \ell_x > 0$ such that $\ell_x < \kappa_G(y) < \ell_x$ for all $y \in \mathcal{C}$ in a neighborhood of x. Since $\mathcal{C}$ is compact, there exist $\ell := \min_{x\in\mathcal{C}} \ell_x$ and $\ell := \max_{x\in\mathcal{C}} \ell_x$ and (4.4) holds for all $y \in \mathcal{C}$. Hence, G is weakly bounded. Then, the same arguments as for (i) apply.

(iii) Assume without loss of generality that $0 \in \mathcal{C}$. Lipschitz continuity of f on $\mathcal{C}$ implies the existence of $\ell'' > 0$ such that $\|f(x)\| \leq \ell''(\|x\|+1)$ for all $x \in \mathcal{C}$. To see this, recall that by the reverse triangle inequality and the definition of Lipschitz continuity there exists $\ell' >$ such that $|\|f(x)\| - \|f(0)\|| \leq \|f(x) - f(0)\| \leq \ell'\|x\|$ for all $x, y \in \mathcal{C}$. It follows that $\|f(x)\| \leq \ell'\|x\|+\|f(0)\|$ and hence ℓ'' can be chosen as the maximum of ℓ' and $\|f(0)\|$ to yield the linear growth property.

Since G is weakly bounded, the same arguments as for (4.4) can be used to establish the existence of $\ell''' > 0$ such that for all $x \in \mathcal{C}$ we have

$$\begin{aligned}\ell'''\|\Pi_{\mathcal{C}}^{G}[f](x)\| < \|\Pi_{\mathcal{C}}^{G}[f](x)\|_{G(x)} &\leq \|f(x)\|_{G(x)} \\ &\leq \|f(x)\| < \ell''(\|x\|+1)\,.\end{aligned}$$

It follows by the same arguments as in the proof of Corollary 4.1 that $\| \mathrm{K} \left[\Pi_{\mathcal{C}}^{G} [f] \right] (x) \| \leq \ell(\|x\| + 1)$ where $\ell = \ell''/\ell'''$, i.e., the linear growth condition applies to $\mathrm{K} \left[\Pi_{\mathcal{C}}^{G} [f] \right]$. □

4.2 Existence of Carathéodory Solutions

We now ask whether the PDS (4.1) admits a (Carathéodory) solution without being regularized. We show that under additional assumptions on the problem parameters, all Krasovskii solutions are Carathódory solutions, consequently proving the existence of the latter.

Consider a set $\mathcal{C} \subset \mathbb{R}^n$, a metric $G : \mathcal{C} \to \mathbb{S}_+^n$ and a vector field $f : \mathcal{C} \to \mathbb{R}^n$. The *sets of Carathéodory* and *Krasovskii solutions* of (4.1) with initial condition $x_0 \in \mathcal{C}$ are denoted by $\mathfrak{S}_C(x_0)$ and $\mathfrak{S}_K(x_0)$, respectively.

Since $\Pi_{\mathcal{C}}^{G} [f] (x) \subset \mathrm{K} \left[\Pi_{\mathcal{C}}^{G} [f] \right] (x)$, every Carathéodory solution of (4.1) is a Krasovskii solution, i.e., $\mathfrak{S}_C(x_0) \subset \mathfrak{S}_K(x_0)$ for all $x_0 \in \mathcal{C}$. To show the opposite inclusion we need the following lemma about the Krasovskii regularization of a projected vector field.

Lemma 4.3. *Consider $\mathcal{C} \subset \mathbb{R}^n$ closed, let $G : \mathcal{C} \to \mathbb{S}_+^n$ and $f : \mathcal{C} \to \mathbb{R}^n$ be continuous. Then, for every $v \in \mathrm{K} \left[\Pi_{\mathcal{C}}^{G} [f] \right] (x)$, one has $\langle f(x), v \rangle_{G(x)} \geq \|v\|_{G(x)}^2$. If, in addition, $\mathcal{C}$ is Clarke regular, then it holds that $f(x) - v \in N_x^G \mathcal{C}$.*

Proof. Let $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ be the *osc-regularization* of $\Pi_{\mathcal{C}}^{G} [f]$, i.e., for all $x \in \mathbb{R}^n$ we have $F(x) := \limsup_{y \to x} \Pi_{\mathcal{C}}^{G} [f] (y)$. By definition of the outer limit, there exist sequences $x_k \to x$ with $x_k \in \mathcal{C}$ and $v_k \to v$ with $v_k \in \Pi_{\mathcal{C}}^{G} [f] (x_k)$ for every $v \in F(x)$ and every $x \in \mathcal{C}$. In particular, $\langle f(x_k), v_k \rangle_{G(x_k)} = \|v_k\|_{G(x_k)}^2$ holds for every k by Lemma 4.1. Since f and G are continuous the equality holds in the limit, i.e., $\langle f(x), v \rangle_{G(x)} = \|v\|_{G(x)}^2$ for every $v \in F(x)$. Taking any convex combination $v = \sum_i \gamma_i v_i$ with $v_i \in F(x)$ and $\gamma_i \geq 0$ and $\sum_i \gamma_i = 1$, we have

$$\sum\nolimits_i \langle f(x), \gamma_i v_i \rangle_{G(x)} = \sum\nolimits_i \gamma_i \|v_i\|_{G(x)}^2 \geq \left\| \sum\nolimits_i \gamma_i v_i \right\|_{G(x)}^2 = \|v\|_{G(x)}^2 \ ,$$

and therefore we get $\langle f(x), v\rangle_{G(x)} \geq \|v\|^2_{G(x)}$ for every vector $v \in \overline{\text{co}}\, F(x) = \text{K}\left[\Pi^G_{\mathcal{C}}[f]\right](x)$.

According to Lemma 4.1, if $\mathcal{C}$ is Clarke regular, given a sequence $x_k \to x$, the sequences $v_k = \Pi^G_{\mathcal{C}}[f](x_k)$ and $\hat{\eta}_k \in N^G_{x_k}\mathcal{C}$ for which $\hat{\eta}_k = f(x_k) - \Pi^G_{\mathcal{C}}[f](x_k)$ are uniquely defined. Since G is continuous, the mapping $x \mapsto N^G_x\mathcal{C}$ is outer semicontinuous (Lemma 3.1) and therefore $\lim_{k\to\infty} \hat{\eta}_k \in N^G_x\mathcal{C}$. In other words, for every $v \in F(x)$ it holds that $f(x) - v \in N^G_x\mathcal{C}$. Since by Clarke regularity $N^G_x\mathcal{C}$ is convex, it follows that, for any convex combination $\eta = \sum_i \gamma_i(f(x) - v_i)$ with $v_i \in F(x)$ and $\gamma_i \geq 0$ and $\sum_i \gamma_i = 1$, it must hold that $\eta \in N^G_x\mathcal{C}$, which completes the proof. □

The proof of the next result follows ideas from [68]. The requirement that G and f need to be continuous deserves particular attention.

Theorem 4.1 (equivalence of solution sets)**.** *Consider the PDS (4.1). If $\mathcal{C}$ is Clarke regular, $G : \mathcal{C} \to \mathbb{S}^n_+$ is continuous, and $f : \mathcal{C} \to \mathbb{R}^n$ is continuous, then $\mathfrak{S}_C(x_0) = \mathfrak{S}_K(x_0)$ for all $x_0 \in \mathcal{C}$.*

Proof. Since $\mathfrak{S}_C(x_0) \subset \mathfrak{S}_K(x_0)$, we need to show only that $x \in \mathfrak{S}_K(x_0)$ implies $x \in \mathfrak{S}_C(x_0)$. For this, we first show that, for all $x \in \mathcal{C}$, $\text{K}\left[\Pi^G_{\mathcal{C}}[f]\right](x) \cap T_x\mathcal{C} = \Pi^G_{\mathcal{C}}[f](x)$.

By definition of $\Pi^G_{\mathcal{C}}[f](x)$ we have $\Pi^G_{\mathcal{C}}[f](x) \subset \text{K}\left[\Pi^G_{\mathcal{C}}[f]\right](x) \cap T_x\mathcal{C}$. For the converse, let $v \in \text{K}\left[\Pi^G_{\mathcal{C}}[f]\right](x) \cap T_x\mathcal{C}$. By Lemma 4.3, $v = f(x) - \hat{\eta}$ for some $\hat{\eta} \in N^G_x\mathcal{C}$ and $\|v\|^2_{G(x)} \leq \langle v, f(x)\rangle_{G(x)}$. Since $\langle v, \eta\rangle_{G(x)} \leq 0$ for all $\eta \in N^G_x\mathcal{C}$ we have

$$\|v\|^2_{G(x)} \leq \langle v, f(x)\rangle_{G(x)} - \langle v, \eta\rangle_{G(x)} \leq \|v\|_{G(x)}\|f(x) - \eta\|_{G(x)}$$

holds for all $\eta \in N^G_x\mathcal{C}$, where the second inequality is due to Cauchy-Schwarz, and therefore $\|v - \hat{\eta}\|_{G(x)} \leq \|f(x) - \eta\|_{G(x)}$ holds for all $\eta \in N^G_x\mathcal{C}$. However, according to Lemma 4.1, $\hat{\eta} = \arg\min_{\eta \in N^G_x\mathcal{C}} \|f(x) - \eta\|_{G(x)}$ is equivalent to $v \in \Pi^G_{\mathcal{C}}[f](x)$.

Consequently, we have $\text{K}\left[\Pi^G_{\mathcal{C}}[f]\right](x) \cap T_x\mathcal{C} = \Pi^G_{\mathcal{C}}[f](x)$ holds for all $x \in \mathcal{C}$ and by Lemma 2.3, $\dot{x}(t) \in T_{x(t)}\mathcal{C}$ holds for almost all t.

Therefore, we have that $\dot{x}(t) \in \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x(t)) \cap T_{x(t)}\mathcal{C}$ which implies that $\dot{x}(t) = \Pi_{\mathcal{C}}^{G}[f](x(t))$ for almost all t and therefore $x \in \mathfrak{S}_C$. □

Theorem 4.1, combined with Corollary 4.1, guarantees the existence Carathéodory solutions. Namely, the following corollary recovers the conditions for existence derived in [68], but generalizes them to a variable metric.

Corollary 4.3 (Existence of Carathéodory solutions)**.** *If $\mathcal{C}$ is Clarke regular, and $G : \mathcal{C} \to \mathbb{S}_{+}^{n}$ and $f : \mathcal{C} \to \mathbb{R}^n$ are continuous, then there exists a (Carathéodory) solution $x : [0, T] \to \mathcal{C}$ for (4.1) for any initial condition $x(0) \in \mathcal{C}$ with either $T = \infty$ or $T < \infty$ and $\limsup_{t \to T^-} \|x(t)\| = \infty$.*

4.3 Differential Variational Inequalities

We now discuss an equivalent formulation of a PDS which is known as a *differential variational inequality* (DVI)[3]. Namely, under Clarke regularity of the feasible set $\mathcal{C}$, we may define the differential inclusion

$$\dot{x} \in f(x) - N_x^G \mathcal{C}\,, \qquad x \in \mathcal{C}\,. \tag{4.5}$$

The set of (Carathéodory) solutions of (4.5) starting at $x_0 \in \mathcal{C}$ is denoted by $\mathfrak{S}_N(x_0)$.

For the Euclidean metric $G \equiv \mathbb{I}$, inclusions of the form (4.5) have been studied extensively, particularly for convex sets $\mathcal{C}$ [52, 17, 68] and the connection to PDSs is well-established.

The inclusion (4.5) is also the starting point for the definition of time-varying PDSs which are a special case of so-called *sweeping processes* [28, 248, 154] in which the set $\mathcal{C}$ is time-varying.

To show existence of solutions for (4.5), Theorem 2.4 cannot be applied directly, because $x \mapsto f(x) - N_x^G \mathcal{C}$ is not well-posed. In particular, this set-valued map is not locally bounded. However, we can establish

[3]The name "differential variational inequality" refers to the fact that, for a solution $x : [0, T] \to \mathcal{C}$ where $\mathcal{C}$ is convex, the derivative $\dot{x}(t)$, whenever it exists, satisfies the variational inequality $\langle \dot{x}(t), y - x \rangle \leq 0$ for all $y \in \mathcal{C}$.

equivalence with PDSs, i.e., any solution of (4.5) is also a solution of (4.1) and vice versa. Hence, the following result is an generalization of [68, Thm. 2.3] to arbitrary metrics.

Corollary 4.4. *Consider a Clarke regular set $\mathcal{C} \subset \mathbb{R}^n$, a continuous vector field f, and a continuous metric G, both defined on $\mathcal{C}$. Then, $\mathfrak{S}_N(x_0) = \mathfrak{S}_C(x_0)$ holds for systems of the form (4.1) and (4.5), and for all $x_0 \in \mathcal{C}$.*

Proof. We first note that $\mathfrak{S}_C(x_0) \subset \mathfrak{S}_N(x_0)$ since $\Pi_{\mathcal{C}}^G[f](x) \subset f(x) - N_x^G\mathcal{C}$ for all $x \in \mathcal{C}$ by virtue of Lemma 4.3 and since $\mathcal{C}$ is Clarke regular. Conversely, let $x \in \mathfrak{S}_N(x_0)$ be defined for $t \in [0, T)$ for $T > 0$. Then for almost all t, we have $\dot{x}(t) \in f(x(t)) - N_{x(t)}^G\mathcal{C}$ and $\dot{x}(t) \in T_{x(t)}\mathcal{C} \cap -T_{x(t)}\mathcal{C}$ by Lemma 2.3. Thus, for $\dot{x}(t) = f(x(t)) - \eta(x(t))$ with $\eta(x(t)) \in N_{x(t)}^G\mathcal{C}$ it must hold that

$$\begin{aligned}\langle f(x(t)) - \eta(x(t)), \eta(x(t))\rangle_{G(x(t))} &\leq 0\\ \langle f(x(t)) - \eta(x(t)), -\eta(x(t))\rangle_{G(x(t))} &\leq 0\,.\end{aligned}$$

Thus, $\langle f(x(t)) - \eta(x(t)), \eta(x(t))\rangle_{G(x(t))} = 0$, and using Lemma 4.1 it follows that $\dot{x}(t) = \Pi_{\mathcal{C}}^G[f](x(t))$. □

We immediately conclude the existence of solutions for (4.5) from Corollary 4.3:

Corollary 4.5 (Existence of solutions for DVI)**.** *If $\mathcal{C}$ is Clarke regular, and $G : \mathcal{C} \to \mathbb{S}_+^n$ and $f : \mathcal{C} \to \mathbb{R}^n$ are continuous, then there exists a solution $x : [0, T] \to \mathcal{C}$ for (4.5) for any initial condition $x(0) \in \mathcal{C}$ with either $T = \infty$ or $T < \infty$ and $\limsup_{t \to T^-} \|x(t)\| = \infty$.*

Remark 4.2. Defining inclusions of the form (4.5) for a set $\mathcal{C}$ that is not Clarke regular is possible but technical since one needs to distinguish between different types of normal cones. Furthermore, depending on the choice of normal cone the resulting set of solutions can be overly relaxed or too restrictive. In the context of sweeping processes, this idea has been explored in [248]. ■

Remark 4.3. Using (ii) in Lemma 4.1 it follows that whenever $\dot{x}$ exists, we have

$$\dot{x} = \arg \min_{v \in f(x) - N_x^G \mathcal{C}} \|v\|_{G(x)} .$$

When G is the Euclidean metric, this *minimum norm* property gives rise to so-called *slow* solutions of (4.5) [18, Ch. 10.1]. For a general metric, the definition of a slow solution generalizes accordingly. However, the property of being "slow" depends on the metric. ■

4.4 Uniqueness of Solutions

To guarantee the uniqueness of solutions for (4.1), we require additional assumptions on $\mathcal{C}$, f and G. From the classical theory for continuous ODEs, it is well-known that local Lipschitz continuity of the vector field f guarantees uniqueness. Since PDSs include unconstrained ODEs by setting $\mathcal{C} = \mathbb{R}^n$, it is unsurprising that Lipschitz continuity of f is generally required for uniqueness.

Before formulating our main uniqueness result, we present an example that illustrates the impact of prox-regularity on the uniqueness of solutions.

Example 4.2 (prox-regularity and uniqueness of solutions). We consider the set

$$\mathcal{C}_\gamma := \{(x_1, x_2) \,|\, |x_2| \geq \max\{0, x_1\}^\gamma\}$$

for $0 < \gamma < 1$, as in Example 3.2. We study how the value of γ affects the uniqueness of solutions of the PDS defined by the uniform "horizontal" vector field $f(x) = (1, 0)$ for all $x \in \mathcal{C}$ and the initial condition $x(0) = 0$ as illustrated in Figure 4.1.

Since $\mathcal{C}_\gamma$ is Clarke regular and closed, since the vector field is uniform, and since we use the Euclidean metric, the existence of Krasovskii solutions and the equivalence of Carathéodory solutions is guaranteed by Corollary 4.2 and Theorem 4.1, respectively. The prox-regularity of $\mathcal{C}_\gamma$ at the origin is, however, guaranteed only for $0 < \gamma \leq \frac{1}{2}$ (Example 3.2).

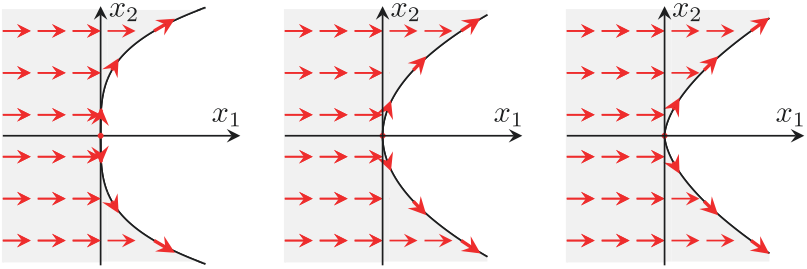

(a) $\gamma = 0.3$ (b) $\gamma = 0.5$ (c) $\gamma = 0.6$

Figure 4.1: Projected vector field on $\mathcal{C}_\gamma$ (shaded) for different values of γ in Example 4.2. The origin is a strong equilibrium in (a) and (b), and it is a weak equilibrium in (c).

A rigorous analysis reveals that for $0 < \gamma \leq \frac{1}{2}$ the origin is a strong equilibrium, i.e., the constant solution $x(t) = 0$ is the unique solution to the PDS. For $\frac{1}{2} < \gamma < 1$, however, the origin is only a weak equilibrium point. Namely, a solution may remain at 0 for an arbitrary amount of time before leaving 0 on either upper or lower halfplane, and thus uniqueness is not guaranteed. ■

For the proof of uniqueness under prox-regularity, we require the following lemma for which local Lipschitz continuity of the metric G is key.

Lemma 4.4. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular at x with respect to a $C^{0,1}$ metric $G : \mathcal{C} \to \mathbb{S}^n_+$. Then, there exist $\bar{\alpha} > 0$ such that for all $y \in \mathcal{C}$ in a neighborhood of x and all $\eta \in N_y^{G,\alpha}$ with $\|\eta\|_{G(y)} = 1$ we have $\langle \eta, z - y \rangle_{G(x)} \leq \bar{\alpha} \|z - y\|^2_{G(x)}$.*

Proof. By the definition of the proximal normal cone, we know that $\eta \in N_y^{G,\alpha}\mathcal{C}$ satisfies $\langle \eta, y - x \rangle_{G(y)} \leq \alpha \|y - x\|^2_{G(y)}$ for y close enough to x. Furthermore, by the equivalence of norms there exists $\alpha' > 0$ such that $\langle \eta, y - x \rangle_{G(y)} \leq \alpha' \|y - x\|^2_{G(x)}$.

Next, we show that $|\langle \eta, x - y \rangle_{G(y)} - \langle \eta, x - y \rangle_{G(x)}| \leq \ell \|y - x\|^2_{G(x)}$ for some $\ell > 0$. Since the space of square symmetric $n \times n$-matrices $\mathbb{S}^n$

is a vector space, we may write

$$\langle \eta, x-y\rangle_{G(y)} - \langle \eta, x-y\rangle_{G(x)} = \langle \eta, x-y\rangle_{G(y)-G(x)}$$

which is a slight abuse of notation since $\langle\cdot,\cdot\rangle_{G(y)-G(x)}$ is not necessarily positive definite and therefore does not induce a norm. Nevertheless, any map of the form $(u,v,G)\mapsto \langle u,w\rangle_G$ where $G\in\alpha_2^n$ is linear in u,v and in G (e.g., $(u,v,G)\mapsto \langle u,w\rangle_{\lambda G} = \lambda\langle u,w\rangle_G$ for any $\lambda\in\mathbb{R}$). Therefore, using Cauchy-Schwarz on $\langle G(y)\eta - G(x)\eta, x-y\rangle$ there exists $\ell>0$ such that

$$\left|\langle \eta, x-y\rangle_{G(y)-G(x)}\right| \le \|\eta\|\|G(y)-G(x)\|_{\mathbb{S}^n}\|x-y\|_{G(x)} \le \ell\|x-y\|^2_{G(x)}\,,$$

where $\|\cdot\|_{\mathbb{S}^n}$ denotes the induced norm on the vector space $\mathbb{S}^n$, and the second inequality follows directly from the local Lipschitz continuity of G. Hence, we can conclude that

$$\begin{aligned}\langle \eta, x-y\rangle_{G(x)} &\le \langle \eta, x-y\rangle_{G(y)} + \left|\langle \eta, x-y\rangle_{G(y)-G(x)}\right| \\ &\le (\alpha'+\ell)\|y-x\|^2_{G(x)}\,,\end{aligned}$$

which completes the proof. □

Lemma 4.4 allows us to establish the following one-sided Lipschitz property of projected vector fields.

Proposition 4.1. *Let $\mathcal{C}\subset\mathbb{R}^n$ and let $f:\mathcal{C}\to\mathbb{R}^n$ and $G:\mathcal{C}\to\mathbb{S}^n_+$ be $C^{0,1}$ maps. Then, for every $x\in\mathcal{C}$ there exists $\ell>0$ such that for all $y\in\mathcal{C}$ in a neighborhood of x we have*

$$\left\langle \Pi^G_{\mathcal{C}}[f](y) - \Pi^G_{\mathcal{C}}[f](x), y-x\right\rangle_{G(x)} \le \ell\|y-x\|^2_{G(x)}\,.$$

Proof. Using Lemma 4.1, we can decompose $\Pi^G_{\mathcal{C}}[f]$ into f and a normal vector, and write

$$\begin{aligned}&\left\langle \Pi^G_{\mathcal{C}}[f](y) - \Pi^G_{\mathcal{C}}[f](x), y-x\right\rangle_{G(x)} \\ &= \langle f(y)-f(x), y-x\rangle_{G(x)} + \langle \eta_y, x-y\rangle_{G(x)} + \langle \eta_x, y-x\rangle_{G(x)}\,, \quad (4.6)\end{aligned}$$

where $\eta_y \in N_y^G\mathcal{C} = \bar{N}_y^{G,\ell}\mathcal{C}$ and $\eta_x \in N_x^G\mathcal{C} = \bar{N}_x^{G,\ell}$ for some $\ell > 0$.

For the first term, we get

$$\langle f(y) - f(x), y - x\rangle_{G(x)} \leq \|f(y) - f(x)\|_{G(x)}\|y - x\|_{G(x)}$$

by applying Cauchy-Schwarz. Since f is locally Lipschitz and using the equivalence of norms there exists $\ell_a > 0$ such that $\|f(y) - f(x)\|_{G(x)} \leq \ell_a\|y - x\|_{G(x)}$ for all $y \in \mathcal{C}$ in a neighborhood of x. Thus, we have $\langle f(y) - f(x), y - x\rangle_{G(x)} \leq \ell_a\|y - x\|_{G(x)}^2$.

For the second and third term in (4.6) we have

$$\langle \eta_y, x - y\rangle_{G(x)} \leq \ell'\|y - x\|_{G(x)}^2\|\eta_y\|_{G(y)}$$
$$\langle \eta_x, y - x\rangle_{G(x)} \leq \ell\|y - x\|_{G(x)}^2\|\eta_x\|_{G(x)}$$

by Lemma 4.4 and the definition of a ℓ-proximal normal vector, respectively.

By Lemma 4.1 we know that $\|\eta_y\|_{G(y)} \leq \|f(y)\|_{G(y)}$ and $\|\eta_x\|_{G(x)} \leq \|f(x)\|_{G(x)}$. Since G and f are continuous (hence locally bounded) we can choose $M > 0$ such that $\|f(z)\|_{G(z)} \leq M$ for all $z \in \mathcal{C}$ in a neighborhood of x. Thus, (4.6) can be bounded by

$$\left\langle \Pi_{\mathcal{C}}^G[f](y) - \Pi_{\mathcal{C}}^G[f](x), y - x\right\rangle_{G(x)} \leq (\ell_a + \ell' M + \ell M)\|y - x\|_{G(x)}^2$$

which completes the proof. □

Using Proposition 4.1 we can now prove our main uniqueness results for PDSs using standard contraction ideas, e.g. [100]. To reiterate, the following statement extends the uniqueness result in [68] which had identified prox-regularity as a key component, but not considered the implications of oblique projection directions induced by a variable metric.

Theorem 4.2 (uniqueness of solutions)**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be prox-regular and let $f : \mathcal{C} \to \mathbb{R}^n$ and $G : \mathcal{C} \to \mathbb{S}_+^n$ be $C^{0,1}$ maps. Then (4.1) admits a unique Carathéodory solution for every initial condition $x_0 \in \mathcal{C}$ (which is also the unique Krasovskii solution).*

Proof. Because our definition of prox-regularity implies Clarke regularity, and local Lipschitz continuity implies continuity, every Krasovskii

solution is a Carathódory solution according to Theorem 4.1 under the given assumptions on $\mathcal{C}$, f, and G.

Let $x, y : [0, T] \to \mathcal{C}$ be two Carathódory solutions of (4.1) with $y(0) = x(0) = x_0 \in \mathcal{C}$. Using Proposition 4.1, there exists $\ell > 0$ and a neighborhood $\mathcal{V}$ of x_0 such that

$$\begin{aligned} &\tfrac{d}{dt}\left(\tfrac{1}{2}\|y(t) - x(t)\|^2_{G(x_0)}\right) \\ &\quad = \left\langle \Pi^G_{\mathcal{C}}[f](y(t)) - \Pi^G_{\mathcal{C}}[f](x(t)), y(t) - x(t)\right\rangle_{G(x_0)} \\ &\quad \leq \ell ||y(t) - x(t)||^2_{G(x_0)} \end{aligned} \tag{4.7}$$

for all t in some non-empty subinterval $[0, T') \subset [0, T]$ for which $x(t)$ and $y(t)$ remain in $\mathcal{V}$. Next, consider the non-negative, absolutely continuous function $q : [0, T') \to \mathbb{R}$ defined as $q(t) := \frac{1}{2}\|y(t) - x(t)\|^2_{G(x_0)} e^{-2\ell t}$. Note that $q(0) = 0$. Furthermore, using (4.7) and applying the product rule we have

$$\begin{aligned} \tfrac{d}{dt} q(t) = \Big(&\left\langle \Pi^G_{\mathcal{C}}[f](y(t)) - \Pi^G_{\mathcal{C}}[f](x(t)), y(t) - x(t)\right\rangle_{G(x_0)} \\ &- \ell ||y(t) - x(t)||^2_{G(x_0)}\Big) e^{-2\ell t} \end{aligned}$$

and since $y(0) = x(0)$ it follows that $\frac{d}{dt} q(t) \leq 0$ for $t \geq 0$. However, since q is non-negative and absolutely continuous, we conclude that $x(t) = y(t)$ for all $t \in [0, T')$ thus finishing the proof of uniqueness. □

4.5 Coordinate Invariance

We can now phrase the main result of this chapter which states that PDSs are invariant under sufficiently smooth, nonlinear coordinate transformations.

For this purpose, we take inspiration from basic differential geometry and use so-called *pullback metrics* and *pushforward vector fields* along a coordinate transformation ϕ. However, the set-valued and discontinuous nature of the project vector field $\Pi^G_{\mathcal{C}}[f]$ and its Krasovskii regularization introduce several technicalities.

Also recall Lemma 2.2 on Page 53 which states that Clarke regularity is preserved by a C^1 coordinate transformation and Proposition 3.4 on Page 70 which guarantees that prox-regularity is preserved by a $C^{1,1}$ coordinate transformation.

Theorem 4.3. *Let $\mathcal{V}, \mathcal{W} \subset \mathbb{R}^n$ be open and consider a C^1 diffeomorphism $\phi : \mathcal{V} \to \mathcal{W}$. Let $\mathcal{C} \subset \mathcal{V}$ be closed (in $\mathbb{R}^n$). Further, let $G : \phi(\mathcal{V}) \to \mathbb{S}^n_+$ be locally weakly bounded and let $\phi^* G : \mathcal{V} \to \mathbb{S}^n_+$ denote the* pullback metric along ϕ, *i.e.,*

$$\langle v, w \rangle_{\phi^* G(x)} := \langle D_x \phi(x) v, D_x \phi(w) \rangle_{G(\phi(x))} \tag{4.8}$$

for all $x \in \mathcal{V}$ and $v, w \in T_x \mathbb{R}^n$. Further, let $f : \mathcal{C} \to \mathbb{R}^n$ be a locally bounded vector field and define the pushforward vector field of f along ϕ^{-1} *as $f_* : \phi(\mathcal{C}) \to \mathbb{R}^n$ such that $f_*(y) := D_x \phi(f(x))$ for all $y \in \phi(\mathcal{C})$ and $x : \phi^{-1}(y)$.*

Then, $x : [0, T] \to \mathcal{C}$ for some $T > 0$ is a Krasovskii (Carathéodory) solution of

$$\dot{x} \in \Pi_{\mathcal{C}}^{\phi^* G}[f](x), \qquad x \in \mathcal{C}, \tag{4.9}$$

if and only if $y := \phi \circ x : [0, T] \to \phi(\mathcal{C})$ is a Krasovskii (Carathéodory) solution of

$$\dot{y} \in \Pi_{\phi(\mathcal{C})}^{G}[f_*](y), \qquad y \in \phi(\mathcal{C}). \tag{4.10}$$

For the proof of Theorem 4.3 we require [211, Thm. 4.26] which reads as follows:

Lemma 4.5. *For a sequence of sets $\{C_k\}$ in $\mathcal{V} \subset \mathbb{R}^n$ and a continuous map $f : \mathcal{V} \to \mathbb{R}^m$, one has*

$$f\left(\liminf_{k\to\infty} C_k\right) \subset \liminf_{k\to\infty} f(C_k), \qquad f\left(\limsup_{k\to\infty} C_k\right) \subset \limsup_{k\to\infty} f(C_k).$$

Namely, Lemma 4.5 makes a statement about how inner and outer limits of sequences of sets commute with continuous functions.

Proof of Theorem 4.3. First, note that since x is absolutely continuous and ϕ is differentiable, $y := \phi \circ x$ is absolutely continuous [212, Ex. 6.44].

Second, it holds that $y(t) \in \phi(\mathcal{C})$ for all $t \in [0, T)$. Third, using (2.10) which states that tangent cones of $\mathcal{C}$ are mapped to tangent cones of $\phi(\mathcal{C})$ through $D\phi$, we can express $\Pi^G_{\phi(\mathcal{C})}[f_*]$ for every $x \in \mathcal{C}$ and $y := \phi(x)$ as

$$\begin{aligned}\Pi^G_{\phi(\mathcal{C})}[f_*](y) &= \operatorname*{arg\,min}_{w \in T_y\phi(\mathcal{C})} \|w - D_x\phi(f(x))\|_G \\ &= \operatorname*{arg\,min}_{w \in D_x\phi(T_x\mathcal{C})} \|w - D_x\phi(f(x))\|_G \\ &= D_x\phi\left(\operatorname*{arg\,min}_{v \in T_x\mathcal{C}} \|D_x\phi(v) - D_x\phi(f(x))\|_G\right),\end{aligned}$$

where for the last equality we introduce the transformation $w := D_x\phi(v)$ for $v \in T_x\mathcal{C}$. Hence, using the definition of the pullback metric (4.8) we continue with

$$\Pi^G_{\phi(\mathcal{C})}[f_*](y) = D_x\phi\left(\operatorname*{arg\,min}_{v \in T_x\mathcal{C}} \|v - f(x)\|_{\phi^*G}\right) = D_x\phi\left(\Pi^{\phi^*G}_{\mathcal{C}}[f](x)\right).$$

Consequently, if x is a Carathéodory solution of (4.9) and hence $\dot{x}(t) \in \Pi^{\phi^*G}_{\mathcal{C}}[f](x(t))$ holds almost everywhere, then $\phi \circ x$ satisfies

$$\tfrac{d}{dt}y(t) = \tfrac{d}{dt}(\phi \circ x) \in D_x\phi\left(\Pi^{\phi^*G}_{\mathcal{C}}[f](x)\right) = \Pi^G_{\phi(\mathcal{C})}[f_*](\phi \circ x(t))$$

almost everywhere and hence y is a Carathéodory solution to (4.10).

It remains to prove that the statement is also true for Krasovskii solutions. For this, we need to show that $D_x\phi(\mathrm{K}\left[\Pi^{\phi^*G}_{\mathcal{C}}[f]\right](y)) \subset \mathrm{K}\left[\Pi^G_{\phi(\mathcal{C})}[f_*]\right](y)$ for all $x \in \mathcal{C}$ and $y = \phi(x)$. Expanding the definition of the Krasovskii regularization we get

$$\begin{aligned}\mathrm{K}\left[\Pi^G_{\phi(\mathcal{C})}[f_*]\right](y) &= \overline{\mathrm{co}} \limsup_{\tilde{y} \to y} \Pi^G_{\phi(\mathcal{C})}[f_*](\tilde{y}) \\ &= \overline{\mathrm{co}} \limsup_{\tilde{x} \to x} D_{\tilde{x}}\phi\left(\Pi^{\phi^*G}_{\mathcal{C}}[f](\tilde{x})\right) \\ &= \overline{\mathrm{co}} \limsup_{\tilde{x} \to x} D_x\phi\left(\Pi^{\phi^*G}_{\mathcal{C}}[f](\tilde{x})\right),\end{aligned}$$

where the last equation is due to the fact that $D_x\phi$ is continuous in x.

Next, since $D_x\phi$ is a linear map we have

$$\begin{aligned}
\mathrm{K}\left[\Pi^G_{\phi(\mathcal{C})}\left[f_*\right]\right](y) &= \overline{\mathrm{co}}\,\limsup_{\tilde{x}\to x} D_x\phi\left(\Pi^{\phi^*G}_{\mathcal{C}}\left[f\right](\tilde{x})\right) \\
&\supset \overline{\mathrm{co}}\; D_x\phi\left(\limsup_{\tilde{x}\to x}\Pi^{\phi^*G}_{\mathcal{C}}\left[f\right](\tilde{x})\right) \\
&= D_x\phi\left(\overline{\mathrm{co}}\,\limsup_{\tilde{x}\to x}\Pi^{\phi^*G}_{\mathcal{C}}\left[f\right](x_k)\right) \\
&= D_x\phi\left(\mathrm{K}\left[\Pi^{\phi^*G}_{\mathcal{C}}\left[f\right]\right](x)\right)\,.
\end{aligned}$$

Namely, $\supset$-inclusion follows since $D_x\phi$ continuous (as a linear map ${}_x\phi : \mathbb{R}^n \to \mathbb{R}^n$) and therefore we can apply Lemma 4.5. The subsequent equality follows because $D_x\phi$, as a linear map, commutes with taking the convex closure.

To conclude we can proceed similar to the case of Carathéodory solutions. Let x be a Krasovskii solution to (4.9) and $y := \phi \circ x$. Then, $\dot{y}(t) = \frac{d}{dt}(\phi \circ x)(t) = D_{x(t)}\phi(\dot{x}(t))$ holds for almost all $t \in [0,T)$ and we have that

$$\dot{y}(t) \in D_{x(t)}\left(K\left[\Pi^{\phi^*G}_{\mathcal{C}}\left[f\right]\right](x(t))\right) \subset K\left[\Pi^G_{\phi(\mathcal{C})}\left[f_*\right]\right](y(t))$$

for almost all $t \in [0,T)$, and thus y is a Krasovskii solution of (4.10).

Because ϕ is diffeomorphism, it follows that the pullback metric ϕ^*G is locally weakly bounded and the pushforward vector field f_* is locally bounded. Therefore, the entire setup can be reverted to prove the converse statement and arrive at the if-and-only-if statement to complete the proof. □

Applying Corollary 4.4, we can also make the analogous statement about solutions of DVIs:

Corollary 4.6. *Consider the same setup as in Theorem 4.3 and, in addition, assume that $\mathcal{C}$ is Clarke regular and that f and G are continuous. Then, $x : [0,T] \to \mathcal{C}$ for some $T > 0$ is a solution of $\dot{x} \in f(x) - N^{\phi^*G}_{\mathcal{C}}$, $x \in \mathcal{C}$ if and only if $y := \phi \circ x : [0,T] \to \phi(\mathcal{C})$ is a solution of $\dot{y} \in f_*(y) - N^G_{\phi(\mathcal{C})}$, $y \in \phi(\mathcal{C})$.*

Concerning the uniqueness, we make the following important observation: Because Theorem 4.3 works both ways, if solutions of (4.9) are unique, then so are the solutions of (4.10) and vice versa.

Theorem 4.2 guarantees the uniqueness of solutions of (4.9) if $\mathcal{C}$ is prox-regular, f is locally Lipschitz, and $\phi^* G$ is locally Lipschitz. If we assume that ϕ is a $C^{1,1}$ diffeomorphism, then the metric G and f_* of (4.10) are locally Lipschitz, and Proposition 3.4 implies that $\phi(\mathcal{C})$ is prox-regular. Consequently, if the requirements for uniqueness are met for (4.9) and ϕ is a $C^{1,1}$ diffeomorphism, then the requirements for uniqueness are met for (4.10).

4.6 Existence & Uniqueness on Manifolds

As a final, consolidating result of this chapter we define PDSs on abstract manifolds and show existence and uniqueness of solutions. We give only a sketch of this construction though, since Corollaries 4.1 and 4.3 and Theorem 4.2 derived in the previous sections are slightly simplified and not entirely suitable for this purpose. For the full set general statements the reader is referred to [Ha2].

The key enabling results for this section are Lemma 2.2 and Proposition 3.4, which state that Clarke regularity and prox-regularity are preserved by C^1 and $C^{1,1}$ coordinate transformations, as well as Theorem 4.3 which proves that PDSs and their solutions can subjected to a change of coordinates.

Recall that a (topological) manifold is defined as a locally Euclidean, second-countable Hausdorff space[4] [162]. In particular, for every point x on a manifold $\mathcal{V}$, there exists an open set $\mathcal{V} \subset \mathcal{M}$ with $x \in \mathcal{V}$ and a continuous map $\phi : \mathcal{U} \to \mathbb{R}^n$. In particular, ϕ is a homeomorphism

[4] A second-countable space is defined as a space with a countable base. That is, there exists a countable family of open subsets such that any open subset can be written as a union of elements of the family. A Hausdorff space is such that for any two distinct points there exist neighborhoods of each point which are disjoint from each other. As a consequence, in a Hausdorff space, any sequence converges to at most one point [162].

onto its image, i.e., ϕ is continuous and has a continuous inverse $\phi^{-1} : \phi(\mathcal{V}) \to \mathcal{V}$. The tuple $(\mathcal{V}, \phi)$ is called a *coordinate chart*.

A *C^k-manifold* is, additionally, equipped with a *C^k-structure*[5] which essentially means that any two coordinate charts $(\mathcal{V}, \phi)$ and $(\mathcal{W}, \psi)$ with overlapping domains define a change of coordinates $\phi \circ \psi^{-1} : \psi(\mathcal{V} \cap \mathcal{W}) \to \phi(\mathcal{V} \cap \mathcal{W})$ that is a C^k map. In addition, we the require the less common notion of a $C^{1,1}$ *manifold*, which is defined analogously to C^k manifolds but with coordinate transformations that are $C^{1,1}$.

Any C^k manifold with $k \geq 1$ admits a tangent space $T_x\mathcal{M}$ at every point $x \in \mathcal{M}$ which is a linear space isomorphic to $\mathbb{R}^n$. A *C^k (Riemannian) metric* G is a map that assigns to every point $x \in \mathcal{M}$ an inner product on the tangent space $T_x\mathcal{M}$. These objects the standard definitions given in [162]. In particular, we note that the notion of a Riemannian metric on $\mathcal{M}$ expressed in local coordinates is equivalent to our definition of a metric (in Section 3.1) as a variable inner product on a subset of $\mathbb{R}^n$ (represented by a map $\mathbb{R}^n \to \mathbb{S}^n_+$). Also, if $\mathcal{M}$ is embedded in $\mathbb{R}^n$, then the tangent space $T_x\mathcal{M}$ of $\mathcal{M}$ as a manifold is equivalent to the tangent cone of $\mathcal{M}$ as a set [211, Ex. 6.8].

In general, to define an object or property on a manifold it needs to be invariant under coordinate transformations. More precisely, one can usually carry definitions from $\mathbb{R}^n$ to manifolds, if they are independent of a particular coordinate representation.

For instance, a curve $\gamma : [0, T) \to \mathcal{M}$ on a C^1 manifold is *absolutely continuous* if its restriction to any chart domain is absolutely continuous in local coordinates [54, Def. 3.3]. This follows, since absolute continuity is preserved by C^1 coordinate transformations [212, Ex 6.44]. Similarly, a vector field on a subset $\mathcal{C} \subset \mathcal{M}$ is *locally bounded at $x \in \mathcal{C}$* if it is locally bounded in any choice of local coordinates, and a metric on $\mathcal{M}$ is *locally weakly bounded at x* if its locally weakly bounded in local coordinates.

Further, we define (Clarke) tangent cones on a manifold as follows:

Definition 4.2. *Given a C^1 manifold $\mathcal{M}$, let $\mathcal{C} \subset \mathcal{M}$ be locally compact. The (Clarke) tangent cone $T_x\mathcal{C}$ ($T_x^C\mathcal{C}$) is the subset of $T_x\mathcal{M}$ such that*

[5] Strictly speaking, a C^kstructure on a manifold is an equivalence class of C^k atlases.

$D_x\phi(T_x\mathcal{C})$ ($D_x\phi(T_x^C\mathcal{C})$) is the (Clarke) tangent cone of $\phi(\mathcal{C}\cap\mathcal{V})$ at $\phi(x)$ for any coordinate chart $(\mathcal{V},\phi)$ defined at x.

The set $\mathcal{C}$ is Clarke regular *at $x\in\mathcal{C}$ if it is Clarke regular in any local coordinate domain defined at x. The set $\mathcal{C}$ is* Clarke regular *if it is Clarke regular at all $x\in\mathcal{C}$.*

In fact, the notion of (Clarke) tangent cones on C^1 manifolds is well-defined because of Lemma 2.2 which states that tangent cones are preserved by C^1 coordinate transformations. To see this, consider the following setup: Let $(\mathcal{V},\phi)$ and $(\mathcal{W},\psi)$ denote two coordinate charts at x and let $\mathcal{N}\subset\mathcal{V}\cap\mathcal{W}$ be open and $x\in\mathcal{N}$. Consider the tangent cone $T_{\phi(x)}\phi(\mathcal{C}\cap\mathcal{V})$ of the set $\phi(\mathcal{C}\cap\mathcal{V})$ at the point $\phi(x)$. Since $\mathcal{V}$ and $\mathcal{N}$ are open, $x\in\mathcal{N}\subset\mathcal{V}$ and ϕ is continuous, we have $T_{\phi(x)}\phi(\mathcal{C}\cap\mathcal{V}) = T_{\phi(x)}\phi(\mathcal{C}\cap\mathcal{N})$. Similarly, we have in the chart domain $\mathcal{W}$ that $T_{\psi(x)}\psi(\mathcal{C}\cap\mathcal{W}) = T_{\psi(x)}\psi(\mathcal{C}\cap\mathcal{N})$.

Next, we may apply Lemma 2.2 by considering the C^1 coordinate transformation $\psi\circ\phi^{-1}:\phi(\mathcal{V}\cap\mathcal{W})\to\psi(\mathcal{V}\cap\mathcal{W})$ and conclude that

$$T_{\psi(x)}\psi(\mathcal{C}\cap\mathcal{N}) = D_{\phi(x)}(\psi\circ\phi^{-1})\left(T_{\phi(x)}\phi(\mathcal{C}\cap\mathcal{N})\right).$$

The same argument holds for Clarke tangent cones. That definition of $T_x\mathcal{C}$ in Definition 4.2 does not depend on the choice of chart $(\mathcal{V},\phi)$.

Analogously, we can apply Lemma 3.5 that states that prox-regularity is preserved by $C^{1,1}$ coordinate transformations. Hence, we define a prox-regular set on a $C^{1,1}$ manifold in terms of local coordinates:

Definition 4.3. *Let $\mathcal{M}$ be a $C^{1,1}$ manifold and consider a locally compact subset $\mathcal{C}\subset\mathcal{M}$ is* prox-regular *at $x\in\mathcal{C}$ if it is prox-regular in any local coordinate domain defined at x and $\mathcal{C}$ is* prox-regular *if it is prox-regular at all $x\in\mathcal{C}$.*

With these definitions it follows that for a subset $\mathcal{C}\subset\mathcal{M}$ of a C^1 manifold equipped with a metric G, the projection operator $\Pi_{\mathcal{C}}^G[w](x)$ of the vector $w\in T_x\mathcal{M}$ onto the tangent cone $T_x\mathcal{C}\subset T_x\mathcal{M}$ is well-defined (i.e., independent of the specific coordinate representation), and can be used to define the dynamical system

$$\dot{x}\in\Pi_{\mathcal{C}}^G[f](x)\quad x\in\mathcal{C}\,, \tag{4.11}$$

where f is a vector field on $\mathcal{C}$.

Assume for now that $\mathcal{C}$ is closed and can be covered by single coordinate chart$(\mathcal{V}, \phi)$, i.e., $\mathcal{C} \subset \mathcal{V}$. Then, we say that $x : [0, T] \to \mathcal{V}$ is a (Carathéodory; Krasovskii) solution of (4.11) if $\phi \circ x : [0, T] \to \phi(\mathcal{V})$ is a solution of (4.11) in the local coordinate representation. For this case Corollaries 4.1 and 4.3 guarantee the existence of Krasovskii and Carathéodory solutions, respectively. Most importantly, however, Theorem 4.3 states that if x a solution in the local coordinated defined by $(\mathcal{V}, \phi)$, then it is a solution in any other coordinate chart $(\mathcal{W}, \psi)$ for which $\mathcal{C} \subset \mathcal{W}$. Hence, this definition of a solution is, again, *coordinate-free* and hence well-defined on manifolds.

The assumption that $\mathcal{C}$ can be covered by a single coordinate domain $\mathcal{V}$ is clearly too restrictive. To relax it, one may consider a PDS (in local coordinates) on the set $\phi(\mathcal{C} \cap \mathcal{V})$ and "stitch together" local solutions on different coordinate domains. However, because $\mathcal{V}$ is open, the set $\phi(\mathcal{C} \cap \mathcal{V})$ is, in general, locally compact, but not closed. Hence, Corollaries 4.1 and 4.3 and Theorem 4.2 derived previously do not apply. Instead, we need to consider the corresponding results in [Ha2] which work for locally compact domains. These results lead to the following statements:

Theorem 4.4 (existence on manifolds)**.** *Let $\mathcal{M}$ be C^1 manifold, G a locally weakly bounded Riemannian metric, $\mathcal{C} \subset \mathcal{M}$ locally compact, and f a locally bounded vector field on $\mathcal{C}$. Then for every $x_0 \in \mathcal{C}$ there exists a Krasovskii solution $x : [0, T) \to \mathcal{C}$ for some $T > 0$ that solves $\dot{x}(t) \in \Pi_{\mathcal{C}}^{G}[f](x(t))$ with $x(0) = x_0$. Furthermore, if $\mathcal{C}$ is Clarke regular, and if f and G are continuous, then every Krasovskii solution is a Carathéodory solution and vice versa.*

For uniqueness, we require stronger conditions. Apart from the fact that prox-regularity is preserved by $C^{1,1}$ coordinate transformations (Lemma 3.5), we also note that local Lipschitz continuity of a metric and of vector fields is preserved under $C^{1,1}$ coordinate transformations.

Theorem 4.5 (uniqueness on manifolds)**.** *Let $\mathcal{M}$ be $C^{1,1}$ manifold, G a $C^{0,1}$ Riemannian metric, $\mathcal{C} \subset \mathcal{M}$ is prox-regular, and f a $C^{0,1}$ vector field on $\mathcal{C}$. Then, for every $x_0 \in \mathcal{C}$ there exists a unique Carathéodory*

	f	G	$\mathcal{C}$	$\mathcal{M}$	
Existence of Krasovskii solutions	LB	LWB	closed[6]	C^1	Corollary 4.1 Theorem 4.4
Existence of Carathéodory solutions	C^0	C^0	C. regular	C^1	Theorem 4.1 Theorem 4.4
Uniqueness of solutions	$C^{0,1}$	$C^{0,1}$	prox-regular	$C^{1,1}$	Theorem 4.2 Theorem 4.5

Table 4.1: Summary of results: regularity requirements for PDSs for a vector field f, metric G, feasible domain $\mathcal{C}$ and regularity of the manifold $\mathcal{M}$. (LB: locally bounded; LWB: locally weakly bounded)

solution $x : [0, T) \to \mathcal{C}$ *for some* $T > 0$ *that solves* $\dot{x}(t) \in \Pi_{\mathcal{C}}^{G}[f](x(t))$ *with* $x(0) = x_0$.

4.7 Summary

The regularity assumption required for the main existence and uniqueness results of this chapter are summarized in Table 4.1. To reiterate, the requirements on the vector field f and the set $\mathcal{C}$ to guarantee existence and uniqueness of Carathéodory solutions have essentially been known since [68]. The results in this chapter provide a substantial generalization by considering the use of a variable metric and consequently defining PDS on abstract manifolds.

For future reference, we provide the following statement as a condensation of most important results in this chapter:

Theorem 4.6. *Consider* (4.1) *and let* $\mathcal{C}$ *be Clarke regular, and* f *and* G *be continuous. Then,* (4.1) *admits a (Carathéodory) solution for every initial condition* $x(0) \in \mathcal{C}$.

If there exists $\kappa > 0$ *such that* $\sup_{x \in \mathcal{C}} \lambda_{G(x)}^{\max} / \lambda_{G(x)}^{\min} \leq \kappa$ *and* f *is Lipschitz on* $\mathcal{C}$, *then* (5.2) *admits a complete solution for every* $x(0) \in \mathcal{C}$.

If $\mathcal{C}$ *is prox-regular, and if* f *and* G *are locally Lipschitz, then* (4.1) *admits a unique solution for every initial condition* $x(0) \in \mathcal{C}$.

CHAPTER 5

Anti-Windup Approximation of PDS

In the context of feedback-based optimization one aims at reducing the need for model information in the control design as much as possible. When it comes to constraints on the control inputs, one would like to exploit physical saturation to enforce them. However, the cascade of an (integral) controller with a saturation element is well-known to be prone to integrator windup. In this scenario, anti-windup schemes are a well-established and effective remedy.

In this chapter we rigorously show how high-gain anti-windup controllers can be used to approximate projected dynamical systems. More precisely, we show uniform convergence and semiglobal practical asymptotic stability. The forthcoming Chapter 9 will revisit this type of *anti-windup approximation* (AWA) of PDS and provide additional results for specific optimization dynamics.

The primary advantage of this type of constraint enforcement is the fact that the set of feasible inputs does not need to be known. Instead, one relies on the physical system to saturate inputs, in other words, to project them onto the set of feasible inputs. This feature is highly desirable in a setting with variable, intermittent input capabilities such as the time-varying availability of renewable energy in power systems.

The material presented corresponds to the first part of [Ha4]. Theorems 5.1 and 5.2 and their respective corollaries constitute the main technical results of this chapter. Simplified results can also be found in [Ha6].

From a technical viewpoint, we focus on the (unconstrained) inclusion

$$\dot{x} \in F_K(x) := f(x, P_{\mathcal{C}}(x)) - \tfrac{1}{K} G^{-1}(P_{\mathcal{C}}(x))\big(x - P_{\mathcal{C}}(x)\big)\,, \tag{5.1}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is a closed set, $f : \mathbb{R}^n \times \mathcal{C} \to \mathbb{R}^n$ is a continuous vector field, $G : \mathcal{C} \to \mathbb{S}^n_+$ is a continuous metric, and $K > 0$ is a constant parameter. Because $P_{\mathcal{C}}$ is in general not single-valued (unless $\mathcal{C}$ is convex), (5.1) has to be treated as a differential inclusion.

We study the behavior of solutions of (5.1) as $K \to 0^+$ and show that under appropriate assumption on $\mathcal{C}, f$, G, and for an initial condition $x(0) \in \mathcal{C}$, these solutions *converge uniformly* to solutions of the projected dynamical system

$$\dot{x} = \Pi^G_{\mathcal{C}}\left[\hat{f}\right](x)\,, \qquad x \in \mathcal{C}\,, \tag{5.2}$$

where we use $\hat{f}(x) := f(x, P_{\mathcal{C}}(x))$. Further, we show that a compact gas set of (5.2) is spas for (5.1) in K.

The key idea for studying (5.1) is to exploit α-prox-regularity of $\mathcal{C}$ which, according to Proposition 3.1, guarantees that $P_{\mathcal{C}}(x)$ is single-valued for all

$$x \in \mathcal{C}^\circ_\alpha := \mathcal{C} + \tfrac{1}{2\alpha} \operatorname{int} \mathbb{B}\,.$$

Hence, on $\mathcal{C}^\circ_\alpha$, (5.1) reduces to an ODE. Furthermore, under appropriate conditions on the problem parameters and for small enough K, trajectories starting in $\mathcal{C}$ remain in $\mathcal{C}^\circ_\alpha$. These insights will be rigorously established in Section 5.1. In Section 5.2 we then show that (5.1) corresponds to a σ-perturbation of (5.2) as a function of K. Standard results from [118] can then be applied to establish uniform convergence and semiglobal practical asymptotic stability in Sections 5.3 and 5.4, respectively.

The following example shows how systems of the form (5.1) arise in the context of anti-windup control for feedback loops with integral controllers and thus motivates the name "anti-windup approximation".

Example 5.1. We show how (5.1) models physical systems and how anti-windup implementations can be used in the context of feedback-based

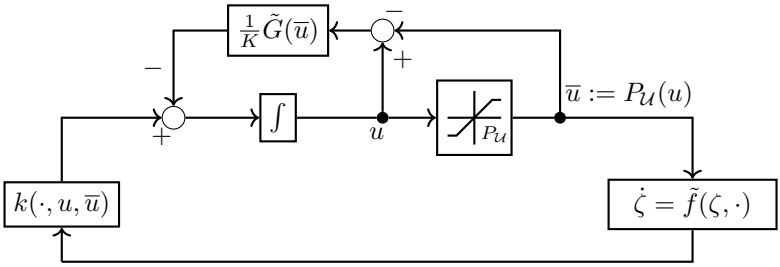

Figure 5.1: Feedback loop with anti-windup (dependence of k and $\tilde{G}$ on $u, \overline{u}$ is not drawn)

optimization to approximate closed-loop optimization dynamics which are formulated as projected dynamical systems.

First, consider the feedback control loop illustrated in Figure 5.1. Namely, we study a plant controlled by an integral feedback controller that is subject to input saturation modeled as an Euclidean projection. An anti-windup scheme is in place to avoid integrator windup. Formally, we consider a dynamical system of the form

$$\dot{\zeta} = \tilde{f}(\zeta, P_{\mathcal{U}}(u)) \qquad x \in \mathbb{R}^m \quad (5.3a)$$

$$\dot{u} = k(\zeta, u, P_{\mathcal{U}}(u)) - \tfrac{1}{K}\tilde{G}^{-1}(P_{\mathcal{U}}(u))(u - P_{\mathcal{U}}(u)) \qquad u \in \mathbb{R}^p\,, \quad (5.3b)$$

where $\mathcal{U} \subset \mathbb{R}^p$ is uniformly prox-regular, $\tilde{f} : \mathbb{R}^m \times \mathcal{U} \to \mathbb{R}^m$ and $k : \mathbb{R}^m \times \mathbb{R}^p \times \mathcal{U} \to \mathbb{R}^p$ are continuous vector fields, $\tilde{G} : \mathcal{U} \to \mathbb{S}_p^+$ is a continuous metric, and $K > 0$ a tuning parameter.

The system (5.3) can be brought into the form (5.1) with $n = m + p$ by defining

$$x := \begin{bmatrix} \zeta \\ u \end{bmatrix} \qquad \mathcal{C} := \mathbb{R}^m \times \mathcal{U} \quad \text{and} \quad G(x) := \begin{bmatrix} \mathbb{I} & 0 \\ 0 & \tilde{G}(u) \end{bmatrix}.$$

Thus, we further have

$$P_{\mathcal{C}}(x) = \begin{bmatrix} \zeta \\ P_{\mathcal{U}}(u) \end{bmatrix} \qquad \text{and} \qquad f(x, P_{\mathcal{C}}(x)) := \begin{bmatrix} \tilde{f}(\zeta, P_{\mathcal{U}}(u)) \\ k(\zeta, u, P_{\mathcal{U}}(u)) \end{bmatrix}.$$

With these definitions, the projected dynamical system (5.2) takes the form

$$\begin{aligned} \dot{x} &= \tilde{f}(\zeta, u) && x \in \mathbb{R}^m \\ \dot{u} &= \Pi_{\mathcal{U}}^{\tilde{G}}\left[k(\zeta, u, u)\right](u) && u \in \mathcal{U}\,, \end{aligned}$$

where we can ignore the projection onto $\mathcal{U}$, because any solution needs to be viable (i.e., remain in $\mathcal{U}$). ■

Remark 5.1. Figure 5.1 shows one limitation of our problem setup: Compared to existing work on anti-windup control [264, 245], we do not model any proportional controller subject to input saturation. This is motivated, on one hand, by theoretical necessity. On the other hand, for our application scenario of feedback-based optimization, stability of the physical plant is usually considered a prerequisite (see Chapter 11 on timescale separation requirements in feedback-based optimization). ■

5.1 Existence, Boundedness, and Equicontinuity

As a first step in studying (5.1), we prove the following lemma for future reference:

Lemma 5.1. *Let $\mathcal{C} \subset \mathbb{R}^n$ be closed and $f : \mathbb{R}^n \times \mathcal{C}$ be continuous. Then, $x \mapsto \hat{f}(x) := f(x, P_{\mathcal{C}}(x))$ is locally bounded and osc. Furthermore, if $\mathcal{C}$ is α-prox-regular for $\alpha > 0$, then $\hat{f}$ is single-valued and continuous for all $x \in \mathcal{C}_\alpha^\circ$.*

Proof. The projection $P_{\mathcal{C}} : \mathbb{R}^n \rightrightarrows \mathcal{C}$ is osc and locally bounded [211, Ex. 5.23], and $P_{\mathcal{C}}(x)$ is non-empty and closed for all $x \in \mathbb{R}^n$ (since $\mathcal{C}$ is closed). By continuity of f it follows that $\hat{f}$ is osc and locally bounded, since both properties are preserved under addition and composition [211, Prop. 5.51 & 5.52]. Using Proposition 3.1 it follows that $\hat{f}$ is single-valued (hence continuous) for $x \in \mathcal{C}_\alpha^\circ$. □

Lemma 5.1 and Proposition 3.1 imply that, on $\mathcal{C}_\alpha^\circ$, F_K is single-valued and continuous. Consequently, standard results for ODEs guarantee

that (5.1) admits a (local) solution for every initial condition $x(0) \in \mathcal{C}_\alpha^\circ$. However, outside of $\mathcal{C}_\alpha^\circ$, (5.1) is a differential inclusion for which the existence of solutions is not immediately guaranteed. Nevertheless, one can establish the existence of Krasovskii solutions (Section 4.1 or [118, Ch. 5]).

For the main result of this section we consider the following (local) setup:

Assumption 5.1. *Consider (5.1) and $x_0 \in \mathcal{C}$. Let $M, \nu, \mu, \alpha, \epsilon > 0$ be such that*

$$\|f(x, P_{\mathcal{C}}(x))\| \leq M \qquad \text{and} \qquad \mu\mathbb{I} \preceq G^{-1}(P_{\mathcal{C}}(x)) \preceq \nu\mathbb{I} \tag{5.5}$$

hold for all $x \in (x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$, and $\mathcal{C}$ is α-prox-regular at every $x \in (x_0 + \epsilon\mathbb{B}) \cap Z$.

Parameters M, μ, ν, ϵ that satisfy (5.5) can always be found for any $x_0 \in \mathcal{C}$ since $x \mapsto f(x, P_{\mathcal{C}}(x))$ is locally bounded by Lemma 5.1, G is continuous, and $P_{\mathcal{C}}$ is single-valued on $(x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$.

Assumption 5.1 allows us to formulate the following proposition which combines the existence of truncated solutions, the invariance of a neighborhood of $\mathcal{C}$, and equicontinuity (i.e. uniform Lipschitz continuity):

Proposition 5.1. *Let Assumption 5.1 be satisfied for $x_0 \in \mathcal{C}$. Given any $T > 0$ and $K < \frac{\mu}{2\alpha M}$, there exists a (T, ϵ)-truncated solution x for (5.1) with $x(0) = x_0$ (where ϵ stems from Assumption 5.1). Furthermore, x satisfies, for almost all $t \in \operatorname{dom} x$,*

$$x(t) \in \mathcal{C} + \tfrac{KM}{\mu}\mathbb{B} \qquad \text{and} \qquad \|\dot{x}(t)\| \leq \left(1 + \tfrac{\nu}{\mu}\right) M\,.$$

Proof. First, we consider the existence of solutions: As mentioned, Lemma 5.1 and Proposition 3.1 imply that, on $(x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$, (5.1) reduces to a continuous ODE which is a well-posed inclusion (trivially). Hence, Theorem 2.4 guarantees the existence of a *maximal* solution which is either complete or given by $x : [0, T'] \to (x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$ starting at x_0 and with $x(T')$ on the boundary of $(x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$.

Next, by Proposition 3.1, we have $\nabla d_{\mathcal{C}}^2(x)^T = 2(x - P_{\mathcal{C}}(x))$ for all $x \in \mathcal{C}_\alpha^\circ$. Hence, the Lie derivative of $d_{\mathcal{C}}^2$ along F_K for all $x \in (x_0+\epsilon\mathbb{B})\cap\mathcal{C}_\alpha^\circ$ is well-defined and satisfies

$$\begin{aligned}\mathcal{L}_{F_K}\left(\tfrac{1}{2}d_{\mathcal{C}}^2(x)\right) &= (x - P_{\mathcal{C}}(x))^T \left(f(x, P_{\mathcal{C}}(x)) - \tfrac{1}{K}G^{-1}(P_{\mathcal{C}}(x))(x - P_{\mathcal{C}}(x)\right)\\ &\leq d_{\mathcal{C}}(x)\|f(x, P_{\mathcal{C}}(x))\|\\ &\qquad - \tfrac{1}{K}(x - P_{\mathcal{C}}(x))^T G^{-1}(P_{\mathcal{C}}(x))\,(x - P_{\mathcal{C}}(x))\\ &\leq d_{\mathcal{C}}(x)\|f(x, P_{\mathcal{C}}(x))\| - \tfrac{\mu}{K}d_{\mathcal{C}}^2(x)\\ &= d_{\mathcal{C}}(x)(M - \tfrac{\mu}{K}d_{\mathcal{C}}(x))\,.\end{aligned}$$

It follows that $\mathcal{L}_{F_K}\left(\frac{1}{2}d_{\mathcal{C}}^2(x)\right) < 0$ whenever $d_{\mathcal{C}}(x) > \frac{KM}{\mu}$. Since $K < \frac{\mu}{2\alpha M}$ and using an invariance argument, it follows that $x(t) \in \mathcal{C}+\frac{KM}{\mu}\mathbb{B} \subset \mathcal{C}_\alpha^\circ$ for all $t \in [0, T']$.

In other words, for small enough K, any solution of (5.1) starting at x_0 remains within a neighborhood of $\mathcal{C}$ on which the projection $P_{\mathcal{C}}$ is single-valued.

Since $x(T')$ lies on the boundary of $(x_0 + \epsilon\mathbb{B}) \cap \mathcal{C}_\alpha^\circ$, but at the same time $x(T') \in \mathcal{C} + \frac{KM}{\mu}\mathbb{B}$, it follows that $\|x(T') - x_0\| = \epsilon$. In other words, $x(T')$ lies on the boundary of $x_0 + \epsilon\mathbb{B}$ (rather than the boundary of $\mathcal{C}_\alpha^\circ$). Hence, (after restricting x to $[0, T]$ if $T' > T$) it can be concluded that x is a (T, ϵ)-truncated solution of (5.1).

Finally, we have that for all $x \in (\mathcal{C} + \frac{KM}{\mu}\mathbb{B}) \cap (x_0 + \epsilon\mathbb{B})$ it holds that

$$\left\|\tfrac{1}{K}G(x)^{-1}(x - P_{\mathcal{C}}(x))\right\| \leq \tfrac{1}{K}\nu\tfrac{KM}{\mu} \leq M\tfrac{\nu}{\mu}\,.$$

It then follows from the definition of M and the triangle inequality that $\|F_K(x)\| \leq M + M\frac{\nu}{\mu}$, thus establishing the bound on $\|\dot{x}(t)\|$. □

The proof of Proposition 5.1 suggests that the prox-regularity assumption on $\mathcal{C}$ is primarily required for $d_{\mathcal{C}}^2(x)$ to have a single-valued derivative in a neighborhood of $\mathcal{C}$. The following example shows, however, that prox-regularity is a more fundamental requirement which, in general, cannot be avoided.

Example 5.2. Consider the set $\mathcal{C} := \{(x_1, x_2) \in \mathbb{R}^2 \mid |x_2| \geq \max\{0, x_1\}^\kappa\}$ for any $\frac{1}{2} < \kappa < 1$. Further assume that $G(x) = \mathbb{I}$ and $f(x) = (1, 0)$

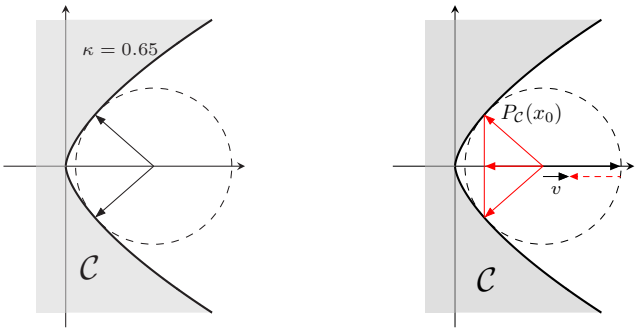

(a) (b)

Figure 5.2: Non-prox-regular set for Example 5.2: (a) non-uniqueness of projection for every $(x_1, 0)$ (b) construction of Krasovskii regularization, namely $v \in \overline{\text{co}}\, F_K(x_0)$.

for all $x \in \mathbb{R}^n$. Hence, we can choose $M = \nu = \mu = 1$ and any $\epsilon > 0$ to satisfy Assumption 5.1. Note, however, that $\mathcal{C}$ is not prox-regular at $(0, 0)$. Namely, every point on the positive x_1-axis has a non-unique projection onto $\mathcal{C}$ as illustrated in Figure 5.2a.

We claim that for every $K > 0$ there exists a Krasovskii[1] solution (i.e., a solution of the inclusion $x \in \overline{\text{co}}\, F_K(x)$) starting on the x_1-axis that leaves the set $\mathcal{C} + \frac{KM}{\mu}\mathbb{B}$ established in Proposition 5.1. This can be deduced graphically from Figure 5.2b. Namely, let $x_0 = (x_{01}, 0)$ be such that $d_{\mathcal{C}}(x_0) = K$. Then, there exists $v = (v_1, 0)$ with $v_1 > 0$ in the Krasovskii regularization of $F_K(x_0)$, i.e., $v \in \overline{\text{co}}\, F_K(x_0)$. In other words, on the boundary of $\mathcal{C} + K\mathbb{B}$, the vector v points out of the (supposedly) invariant set. This, in turn, can be used to rigorously establish that the set $\mathcal{C} + K\mathbb{B}$ is not invariant, illustrating that the conclusion of Proposition 5.1 does not hold without prox-regularity of $\mathcal{C}$, even when considering more general Krasovskii solutions. ■

[1]We cannot rely on the existence of Carathéodory solutions because $\mathcal{C}$ is not prox-regular and Proposition 5.1 does not apply, but every Carathéodory solution (if it exists) is a Krasovskii solution.

5.2 Anti-Windup Approximations as Perturbed PDS

As a key technical result, we establish that solutions of the AWA (5.1) are also solutions of a σ-perturbation of the PDS (5.2). To prove this fact, we need to revisit the equivalence between PDS and differential variational inequalities from Section 4.3.

Namely, similarly to Corollary 4.4, we can show that solutions of (5.2) are equivalent to solutions of the *truncated* differential variational inequality

$$\dot{x} \in F(x) := \hat{f}(x) - N_x^G\mathcal{C} \cap \gamma\mathbb{B} \qquad x \in \mathcal{C}\,, \tag{5.6}$$

where $\gamma \geq \sup_{x\in\mathcal{C}} \|f(x)\|_{G(x)}$ (assuming $\sup_{x\in\mathcal{C}} \|f(x)\|_{G(x)} < \infty$). The advantage of this latter inclusion is that the mapping F is bounded.

Proposition 5.2. *If $\mathcal{C}$ is Clarke regular and $x \mapsto \|f(x)\|_{G(x)}$ is bounded, then, $x : [0, T] \to \mathcal{C}$ with $T > 0$ is a solution of (5.2) if and only if it is a solution of (5.6).*

The proof of Proposition 5.2 is the same as for Corollary 4.4, except for the additional boundedness assumption. For an explicit proof the reader is referred to [Ha4].

Lemma 5.2. *If f and G are continuous and $\mathcal{C}$ Clarke regular, (5.6) is well-posed.*

Proof. Non-emptiness and convexity of $F(x)$ are immediate because $N_x^G\mathcal{C} \cap \gamma\mathbb{B}$ is non-empty (in particular, $0 \in N_x^G\mathcal{C}$) and convex for all $x \in \mathcal{C}$ (and f is single-valued). For outer semicontinuity recall that for a Clarke regular $\mathcal{C}$ and continuous G the mapping $x \mapsto N_x^G\mathcal{C}$ is osc [Ha2, Lem. A.6]. It then follows that the truncation $N_x^G\mathcal{C}\cap\gamma\mathbb{B}$ is osc and locally bounded [211, p.161]. Finally, since f is continuous and single-valued, $x \mapsto f(x) - N_x^G\mathcal{C} \cap \gamma\mathbb{B}$ is osc and locally bounded. □

Next, consider $x_0 \in \mathcal{C}$, and let $M, \mu, \nu, \alpha, \epsilon > 0$ be such that Assumption 5.1 is satisfied. From Proposition 5.2 it follows that, for some $T > 0$,

every (T,ϵ)-truncated solution $x:[0,T']\to(x_0+\epsilon\mathbb{B})$ of the PDS (5.2) with $x(0)=x_0$ is also a (T,ϵ)-truncated solution of the inclusion

$$\dot{x}\in\hat{F}(x):=f(x,P_{\mathcal{C}}(x))-N_x^G\mathcal{C}\cap\gamma\mathbb{B}\quad\text{where}\quad\gamma:=\max\left\{\tfrac{M}{\sqrt{\mu}},\tfrac{\nu}{\mu}M\right\} \tag{5.7}$$

and vice versa. This choice of γ will be convenient in the proof of Proposition 5.3 below. For now, note that using Cauchy-Schwarz, it holds that

$$\sup_{x\in x_0+\epsilon\mathbb{B}}\|f(x,P_{\mathcal{C}}(x))\|_{G(x)}\leq\sup_{x\in x_0+\epsilon\mathbb{B}}\underbrace{\sqrt{\|G(x)\|}}_{\leq 1/\sqrt{\mu}}\underbrace{\|f(x,P_{\mathcal{C}}(x))\|}_{\leq M}\leq\gamma\,,$$

thus satisfying the condition on γ in (5.6) and Proposition 5.2.

Furthermore, given $x_0\in\mathcal{C}$, let Assumption 5.1 hold with some $\epsilon>0$. By Lemma 5.1 we have that $x\mapsto f(x,P_{\mathcal{C}}(x))$ is continuous on $x\in\mathcal{C}_\alpha^\circ$ and hence uniformly continuous on the bounded set $\mathcal{C}_\alpha^\circ\cap(x_0+\epsilon\mathbb{B})$. As a consequence of uniform continuity there exists $\omega\in\mathcal{K}_\infty$ such that, for all $x,x'\in\mathcal{C}_\alpha^\circ\cap(x_0+\epsilon\mathbb{B})$, we have

$$\|f(x,P_{\mathcal{C}}(x))-f(x',P_{\mathcal{C}}(x'))\|\leq\omega(\|x-x'\|)\,. \tag{5.8}$$

Proposition 5.3. *Consider $x_0\in\mathcal{C}$ and let Assumption 5.1 hold with M,ν,μ,α and ϵ. Further, let $K<\frac{\mu}{2\alpha M}$. Then, for some $T>0$, every (T,ϵ)-truncated solution $x:[0,T']\to(x_0+\epsilon\mathbb{B})$ of (5.1) is a solution of the σ-perturbation of (5.7) with $\sigma:=\max\left\{\frac{KM}{\mu},\omega\left(\frac{KM}{\mu}\right)\right\}$, where $\omega\in\mathcal{K}_\infty$ satisfies (5.8).*

Proof. We need to show that the (T,ϵ)-truncated solution x satisfies

$$\dot{x}(t)\in\hat{F}_\sigma(x(t))\,,\qquad x(t)\in\mathcal{C}_\sigma \tag{5.9}$$

for almost all $t\in[0,T']$, where $\mathcal{Z}_\sigma:=\mathcal{C}+\sigma\mathbb{B}$ and $\hat{F}_\sigma(x):=\operatorname{co}\hat{F}((x+\sigma\mathbb{B})\cap\mathcal{C})+\sigma\mathbb{B}$ for all $x\in\mathcal{C}_\sigma$ and with $\hat{F}$ defined in (5.7). Note that for $x\in\mathcal{C}_\sigma$ we have that

$$P_{\mathcal{C}}(x)\subset(x+\sigma\mathbb{B})\cap\mathcal{C}\,. \tag{5.10}$$

Proposition 5.1 guarantees that $x(t) \in \mathcal{C} + \frac{KM}{\mu}\mathbb{B}$, and since $\sigma \geq \frac{KM}{\mu}$ it follows that $x(t) \in \mathcal{C}_\sigma$ for all $t \in [0, T']$. For the remainder of the proof we omit the argument of $x(t)$ to simplify notation. All statements hold for almost all $t \in [0, T']$.

Since $x - P_{\mathcal{C}}(x) \subset N_{P_{\mathcal{C}}(x)}\mathcal{C}$ for all $x \in \mathbb{R}^n$ [211, Ex. 6.16] and using (3.2) we have

$$\tfrac{1}{K} G(P_{\mathcal{C}}(x))^{-1} \left(x - P_{\mathcal{C}}(x)\right) \in N^G_{P_{\mathcal{C}}(x)}\mathcal{C}\,. \tag{5.11}$$

Furthermore, since $x \in \mathcal{C} + \frac{KM}{\mu}$ and using γ as defined in (5.7) we have that

$$\left\| \tfrac{1}{K} G(P_{\mathcal{C}}(x))^{-1} \left(x - P_{\mathcal{C}}(x)\right) \right\| \leq \tfrac{1}{K}\nu\tfrac{KM}{\mu} = \tfrac{\nu M}{\mu} \leq \gamma. \tag{5.12}$$

Combining (5.11) and (5.12) we have

$$\dot{x} \in f(x, P_{\mathcal{C}}(x)) - N^G_{P_{\mathcal{C}}(x)}\mathcal{C} \cap \gamma\mathbb{B}\,. \tag{5.13}$$

Note that, in contrast to (5.7), the normal cone is evaluated at $P_{\mathcal{C}}(x)$.

Next, using the fact that ω, as defined in (5.8), is strictly increasing, and exploiting the definition of σ, we have

$$\|f(x, P_{\mathcal{C}}(x)) - f(P_{\mathcal{C}}(x), P_{\mathcal{C}}(x))\| \leq \omega\left(\|x - P_{\mathcal{C}}(x)\|\right) \leq \omega\left(\tfrac{KM}{\mu}\right) \leq \sigma\,. \tag{5.14}$$

Therefore, in summary, using (5.14) on (5.13), as well as (5.10), we have that

$$\begin{aligned}
\dot{x} &\in f(x, P_{\mathcal{C}}(x)) - N^G_{P_{\mathcal{C}}(x)}\mathcal{C} \cap \gamma\mathbb{B} \\
&\subset f(P_{\mathcal{C}}(x), P_{\mathcal{C}}(x)) + \sigma\mathbb{B} - N^G_{P_{\mathcal{C}}(x)}\mathcal{C} \cap \gamma\mathbb{B} \\
&= \hat{F}(P_{\mathcal{C}}(x)) + \sigma\mathbb{B} \subset \hat{F}((x + \sigma\mathbb{B}) \cap \mathcal{C}) + \sigma\mathbb{B} \subset \hat{F}_\sigma(x)\,.
\end{aligned}$$

Hence, $x(\cdot)$ satisfies (5.9) which completes the proof. □

5.3 Uniform Convergence

We establish the graphical/uniform convergence of solutions of the anti-windup approximation (5.1) to solutions of the projected dynamics (5.2).

This proof requires two arguments: On the one hand, we need to show that a graphically convergent sequence of solutions of (5.1) converges to a solution of (5.2). On the other hand, we need that such a graphically convergent sequence exists.

Starting with the latter requirement, we first recall that from a bounded sequence of sets, we can always extract a graphically convergent subsequence [118, Thm. 5.7]. This applies in particular to a sequence of (uniformly) truncated solutions:

Lemma 5.3. *Consider a sequence $K_n \to 0^+$ and $x_0 \in \mathcal{C}$. Given $T, \epsilon > 0$, any sequence $\{x_n\}$ of (T,ϵ)-truncated solution of (5.1) with $K = K_n$ and $x_n(0) = x_0$ has a graphically convergent subsequence.*

Lemma 5.3 is purely set-theoretic and does not imply that the limit $\operatorname{gph} \lim_{n\to\infty} x_n$ is a single-valued map. Hence, we need the following simplification[2] of [118, Thm. 5.29]:

Lemma 5.4. *Let the inclusion (2.12) be well-posed and $x_0 \in \mathcal{C}$. Further, given any $T, \epsilon, \rho > 0$ and $\delta_i \to 0^+$, let $x_i : [0, T_i] \to \mathcal{X}_i$ denote a (T,ϵ)-truncated solution of the $\delta_i\sigma$-perturbation of (2.12). If the sequence $\{x_i\}$ converges graphically, then convergence is to a solution $x : [0, T] \to \mathcal{X}$ of (2.12), where $T = \lim_{i\to\infty} T_i$.*

Remark 5.2. In the context of Lemma 5.4, graphical convergence implies uniform convergence to x on every subinterval of $[0, T)$ [118, Lem. 5.28]. Furthermore, if $T_i \geq T$ for all i, then convergence is uniform on $[0, T]$. ■

Next, we require the following lemma about $\mathcal{K}_\infty$-functions [229, Cor. 10]:

Lemma 5.5. *Given $\omega \in \mathcal{K}_\infty$, there exist $\sigma_1, \sigma_2 \in \mathcal{K}_\infty$ with $\omega(rs) \leq \sigma_1(r)\sigma_2(s)$ for all $r, s \geq 0$.*

Since, by Proposition 5.3, solutions of (5.1) are solutions of a σ-perturbation of an alternate form PDS (5.7) we can use Lemma 5.4 to establish the following result:

[2] We require only the first of the two statements of the original theorem. Further, we consider the case where ρ is constant. Finally, we work with truncated solutions which have, by definition, a compact domain (and thus are trivially *locally eventually bounded* [118, Def. 5.24]).

Proposition 5.4. *Given $x_0 \in \mathcal{C}$, let Assumption 5.1 be satisfied. Consider $T > 0$ and a sequence $K_i \to 0^+$, and assume that a sequence of (T, ϵ)-truncated solutions x_i of the AWA (5.1) with $K = K_i$ and $x_i(0) = x_0$ for all i converges graphically. Then, the limit is a (T, ϵ)-truncated solution of the PDS (5.2).*

Proof. Let $M, \mu, \nu > 0$ and $\omega \in \mathcal{K}_\infty$ be defined as in Assumption 5.1 and (5.8), respectively. Using Lemma 5.5, there exist $\sigma_1, \sigma_2 \in \mathcal{K}_\infty$ such that $\omega(rs) \leq \sigma_1(r)\sigma_2(r)$ for all $r, s \geq 0$. Hence, we define $\delta_i := \max\{K_i, \sigma_1(K_i)\}$ and $\rho := \max\left\{\frac{M}{\mu}, \sigma_2\left(\frac{M}{\mu}\right)\right\}$.

Proposition 5.3 states that for every K_i, the solution x_i of (5.1) is also a solution of the σ-perturbation of (5.7) with $\sigma := \max\left\{K_i\frac{M}{\mu}, \omega\left(K_i\frac{M}{\mu}\right)\right\}$. It follows that x_i is also a solution of every σ'-perturbation of (5.7) with $\sigma' \geq \sigma$. In particular, we can set

$$\sigma' := \delta_i\rho = \max\{K_i, \sigma_1(K_i)\} \max\left\{\tfrac{M}{\mu}, \sigma_2(\tfrac{M}{\mu})\right\} \geq \sigma\,,$$

and thus we have that x_i is a solution of the $\delta_i\rho$-perturbation of (5.7).

Since, by assumption, $\{x_i\}$ converges graphically to x it follows from Lemma 5.4 that x is a solution of (5.7), and, by Proposition 5.2, x is a solution of (5.2).

Finally, we need to show that $x : [0, T'] \to (x_0 + \epsilon\mathbb{B})$ is a (T, ϵ)-truncated solution. Namely, we need to show that either $T = T'$ or $\|x(T) - x_0\| = \epsilon$. This requirement is equivalent to $(T', x(T'))$ lying on the boundary of the cylinder $\mathcal{X} := [0, T] \times (x_0 + \epsilon\mathbb{B})$. Since, by definition, for every i, x_i is a (T, ϵ)-truncated solution of (5.1) we have that $(T_i, x_i(T_i)) \in \partial\mathcal{X}$ for all i. Since $\partial\mathcal{X}$ is closed, it follows that the limit also lies in $\partial\mathcal{X}$. □

Now, we can immediately combine Lemma 5.3 and Proposition 5.4 to arrive at our first main result about the graphical convergence of truncated solutions (i.e., local) solutions of anti-windup approximations to a projected dynamical system:

Theorem 5.1. *Let Assumption 5.1 be satisfied for some $x_0 \in \mathcal{C}$. Given any $T > 0$ (and $\epsilon > 0$ from Assumption 5.1), consider a sequence*

$K_n \to 0^+$ and let $\{x_n\}$ denote a sequence of (T,ϵ)-truncated solutions of the AWA (5.1) with $K = K_n$ and $x_n(0) = x_0$. Then, there exists a subsequence of $\{x_n\}$ that converges graphically to a (T,ϵ)-truncated solution of the PDS (5.2).

Under certain circumstances, it can be useful to know that, rather than a subsequence of gains $\{K_n\}$, any sequence $K_n \to 0^+$ will lead to a converging sequence of solutions. This is guaranteed if it is known that the PDS (5.2) has a unique solution:

Corollary 5.1. *Let Assumption 5.1 be satisfied for some $x_0 \in \mathcal{C}$. Given any $T > 0$ (and $\epsilon > 0$ from Assumption 5.1), assume that the PDS (5.2) admits a unique (T,ϵ)-truncated solution x with $x(0) = x_0$. Then, any sequence $\{x_n\}$ of (T,ϵ)-truncated solutions of the AWA (5.1) with $x_n(0) = x_0$ and $K = K_n$ with $K_n \to 0^+$ converges graphically to the (unique) (T,ϵ)-truncated solution of the PDS (5.2).*

Proof. Assume, for the sake of contradiction, that $\{x_n\}$ does not converge to the unique solution x of (5.2). This implies that there exists $\nu > 0$ and a subsequence $\{x_m\}$ of $\{x_n\}$ such that $d_\infty(\operatorname{gph} x_m, \operatorname{gph} x) \geq \nu$ for all m where d_∞ denotes the Hausdorff distance between two sets. (In particular, since x is a truncated solution $\operatorname{gph} x$ is compact and thus graphical convergence is equivalent to convergence with respect to d_∞ [211, Ex. 4.13].) However, by Lemma 5.3, the sequence $\{x_m\}$ has a convergent subsequence that converges to some limit $\tilde{x}$. By Proposition 5.4, $\tilde{x}$ is a solution of (5.2), but we also have $\|\tilde{x} - x\|_\infty \geq \varepsilon$ which contradicts the uniqueness of x. □

Finally, we can state the following *ready-to-use* result about uniform convergence in the case when the existence of unique complete solutions is guaranteed:

Corollary 5.2. *Consider the AWA (5.1), let $\mathcal{C}$ be uniformly prox-regular, f globally Lipschitz, and there exist $\mu, \nu > 0$ such that $\mu\mathbb{I} \preceq G^{-1}(x) \preceq \nu\mathbb{I}$ for all $x \in \mathbb{R}^n$. Given $x_0 \in \mathcal{C}$ and a sequence $K_n \to 0^+$, every sequence of complete solutions x_n of the AWA (5.1) with initial condition x_0 and $K = K_n$ converges uniformly to the unique complete solution of the PDS (5.2) on every compact interval $[0,T]$.*

Proof. Note that the assumptions on $\mathcal{C}$, G, and f guarantee that for every initial condition (5.2) admits a unique complete (Carathéodory) solution (Theorem 4.6).

Hence, given any $T > 0$, let $x : [0, T] \to \mathcal{C}$ denote the unique solution of the PDS (5.2) and define $\epsilon > \sup_{t\in[0,T]} \|x(t) - x_0\|$. Since f is continuous and hence bounded over a compact set, Assumption 5.1 is satisfied with ν, μ, α and by choosing $M := \max_{x\in x_0+\epsilon\mathbb{B}} \|f(x, P_{\mathcal{C}}(x))\|$. Theorem 5.1 guarantees convergence of a subsequence to the (T, ϵ)-truncated solution $x : [0, T'] \to \mathcal{X}$ of (5.2). Moreover, for the same reason as in Corollary 5.1 the sequence itself converges.

Finally, by definition of ϵ, we have that x is defined on $[0, T']$ with $T' = T$ and $\|x(T) - x_0\| < \epsilon$ and, in this case, graphical convergence of (T, ϵ)-truncated solutions implies their uniform convergence on $[0, T]$ (see Remark 5.2). □

Remark 5.3. Theorem 5.1 and its corollaries can be slightly generalized, albeit at the expense of additional technicalities. For instance, instead of considering a single initial condition $x_0 \in \mathcal{C}$, it is in general possible to consider a sequence of initial conditions (under some additional restrictions) that converges to x_0. ■

5.4 Semiglobal Practical Robust Stability

Since anti-windup approximations can be seen as perturbations of projected dynamical systems, we can establish semiglobal practical asymptotic stability in K with the following simplification[3] of [118, Lem. 7.20]:

Lemma 5.6. *Let the inclusion (2.12) be well-posed and let $\mathcal{A} \subset \mathcal{X}$ be a compact and asymptotically stable set for (2.12), i.e., $d_{\mathcal{A}}(x(t)) \leq \beta(d_{\mathcal{A}}(x(0)), t)$ for all $t \geq 0$ holds for some $\beta \in \mathcal{KL}$ and any (complete)*

[3] We consider only global asymptotic stability, which allows us to use the distance function instead of more general indicator functions. Further, we limit ourselves to ρ being a positive constant instead of a function. As noted in Remark 2.2, compactness and stability of $\mathcal{A}$ guarantee the existence of complete solutions since finite-time escape is not possible.

solution x of (2.12). Then, for every $\rho > 0$, every compact $\mathcal{B} \subset \mathbb{R}^n$, and every $\zeta > 0$ there exists $\delta \in (0,1)$ such that every solution x_δ of the $\delta\rho$-perturbation of (2.12) starting in $\mathcal{B} \cap C_{\delta\rho}$ satisfies $d_{\mathcal{A}}(x_\delta(t)) \leq \beta(d_{\mathcal{A}}(x_\delta(0)), t) + \zeta$ for all $t \geq 0$.

Hence, using Proposition 5.3, we arrive at the following second main result:

Theorem 5.2. *Consider a PDS (5.2) where $\mathcal{C}$ is Clarke regular, f and G are continuous, and for which the compact set $\mathcal{A} \subset \mathcal{C}$ is gas, i.e., there is $\beta \in \mathcal{KL}$ such that for every solution x it holds that*

$$d_{\mathcal{A}}(x(t)) \leq \beta(d_{\mathcal{A}}(x(0)), t) \quad \forall t \geq 0\,.$$

Then, for every $\zeta > 0$ and every compact $\mathcal{B} \subset \mathcal{C}$ there exists $K^\star > 0$ such that for all $K \in (0, K^\star)$ every solution x_K of the AWA (5.1) with $x_K(0) \in \mathcal{B}$ satisfies

$$d_{\mathcal{A}}(x_K(t)) \leq \beta(d_{\mathcal{A}}(x_K(0)), t) + \zeta \qquad \forall t \geq 0\,.$$

Proof. First, we establish that Assumption 5.1 holds for every $x_0 \in \mathcal{B}$. Since $\mathcal{B}$ is compact, let $\overline{\beta} := \max_{x\in\mathcal{B}} \beta(d_{\mathcal{A}}(x), 0)$. Since β is strictly increasing and unbounded, and, since $\mathcal{A}$ is compact, the set $\mathcal{V} := \{x \,|\, \beta(d_{\mathcal{A}}(x), 0) \leq \overline{\beta}\}$ is compact. Hence, we can choose $\epsilon > 0$ such that $\mathcal{V} \subset \mathcal{B} + \epsilon\mathbb{B}$. It follows that any solution of (5.2) starting in $\mathcal{B}$ remains in $\mathcal{B} + \epsilon\mathbb{B}$. By continuity over the compact set $\mathcal{B} + \epsilon\mathbb{B}$, we can further choose and $M, \mu, \nu > 0$ such that $\|f(x, P_{\mathcal{C}}(x))\| \leq M$ and $\mu\mathbb{I} \preceq G^{-1}(x) \preceq \nu\mathbb{I}$ holds for all $x \in \mathcal{B} + \epsilon\mathbb{B}$. Thus, Assumption 5.1 is satisfied for all $x_0 \in \mathcal{B}$. Further, every (complete) solution of the PDS (5.2) starting in $\mathcal{B}$ remains in $\mathcal{B} + \epsilon\mathbb{B}$ and hence can be written in its alternate form (5.7). Next, fix any $\rho > 0$. Lemma 5.6 implies that for every $\zeta > 0$ and every compact $\mathcal{B} \subset \mathcal{C}$ there exists $\delta \in (0,1)$ such that the $\delta\rho$-perturbation is ζ-practically pre-asymptotically stable. Given such a δ, we conclude that there exists $K^\star > 0$ that, for all $K' < K^\star$, $\max\{K'\frac{M}{\mu}, \omega(K'\frac{M}{\mu})\} \leq \delta\rho$ since ω is strictly increasing and $\omega(0) = 0$. Thus, Proposition 5.3 states that the solution of (5.1) with $K = K'$ is a solution of the σ-perturbation of (5.7) with $\sigma = \max\{K'\frac{M}{\mu}, \omega(K'\frac{M}{\mu})\}$.

Moreover, it is also solution to any σ'-perturbation with $\sigma' \geq \sigma$ and, in particular, for $\sigma' = \delta\rho$. □

Since the asymptotic stability of $\mathcal{A}$ can often be established with a smooth Lyapunov function (see [118, Thm. 3.18]), we can also state the following corollary:

Corollary 5.3. *Consider the PDS (5.2) where $\mathcal{C}$ is Clarke regular, f and G are continuous. Further, consider a compact set $\mathcal{A} \subset \mathcal{C}$ for which there exists a* Lyapunov function[4]*. Then, for every $\zeta > 0$ and every compact set $\mathcal{B} \subset \mathbb{R}^n$, there exists $K^\star$ such that for all $K \in (0, K^\star)$ every solution of (5.1) converges to a subset of $\mathcal{A} + \zeta\mathbb{B}$.*

5.5 Notes & Comments

Although the design of AWA for PDS is inspired by anti-windup controllers which are ubiquitous in practical control applications, the approach presented in this chapter is almost entirely unrelated to other works in this field, which, in the last decades, has mostly focused on LMI-based approaches [264, 245].

On the one hand, our modeling remains limited because, with respect to the standard use of anti-windup control, we have considered only a very simple type of *static* anti-windup controller in contrast to more general *dynamic* compensators. On top of that, as stated in Remark 5.1, we consider only pure integral controllers without any (saturated) proportional control component.

On the other hand, we deal with general nonlinear systems (as opposed to the anti-windup literature which seems to mostly focus on LTI plants) and we have used a generalized notion of saturation as projection onto a closed set (instead of componentwise saturation onto closed intervals).

[4]Namely, $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is a Lyapunov function for $\mathcal{A}$ if it is differentiable everywhere on $\mathcal{C}$, there exist $\underline{\alpha}, \overline{\alpha} \in \mathcal{K}_\infty$ such that $\underline{\alpha}(d_{\mathcal{A}}(x)) \leq V(x) \leq \overline{\alpha}(d_{\mathcal{A}}(x))$ for all $x \in \mathcal{C}$, and $\nabla V(x)\Pi_{\mathcal{C}}^{G}[f](x) \leq -\alpha(x)$ for all $x \in \mathcal{C}$ where $\alpha : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is continuous and positive definite with respect to $\mathcal{A}$, i.e., $\alpha(x) > 0$ for all $x \notin \mathcal{A}$ and $\alpha(x) = 0$ for all $x \in \mathcal{A}$.

CHAPTER 6

A Feedback-Enabled Discretization of PDS

Besides driving a plant to a cost-minimizing steady state, a key task that can be addressed with feedback-based optimization is the satisfaction of constraints on the plant outputs. Unfortunately, satisfying these output constraints is more challenging than handling input constraints which can be enforced directly, e.g., using the anti-windup schemes presented in the previous chapter.

In this chapter we therefore introduce a new discretization of projected dynamical systems that can be implemented as a closed-loop system for the purpose of enforcing output constraints asymptotically.

The structure of this chapter is analogous to the previous one. Namely, we show that (interpolated) trajectories of our iterative scheme converge uniformly to solutions of a projected dynamical system as we let the step-size go to zero. Furthermore, we establish semiglobal practical asymptotic stability and other robustness properties of our new discretization.

This chapter slightly generalizes the algorithm in [Ha1] and presents new theoretical results that were unpublished by the time of writing. Preliminary results for this chapter resulted from a collaboration with V. Häberle.

Unless noted otherwise, we consider a feasible set $\mathcal{C} \subset \mathbb{R}^n$ of the form

$$\mathcal{C} := \{x \,|\, g(x) \leq 0,\, k(x) \leq 0\}\,, \tag{6.1}$$

where $g : \mathbb{R}^n \to \mathbb{R}^p$ and $k : \mathbb{R}^n \to \mathbb{R}^q$ are continuously differentiable and we define $\mathcal{U} := \{x \,|\, g(x) \leq 0\}$.

In addition, we consider a metric $G : \mathcal{U} \to \mathbb{S}^n_+$ and a vector field $f : \mathcal{U} \to \mathbb{R}^n$. Together with some $\alpha > 0$, these objects allow us to define the map $\Sigma^G_{\mathcal{C}}[f] : \mathcal{U} \times \mathbb{R} \to \mathbb{R}^n$ as the solution of the parametrized optimization problem

$$\begin{aligned} \Sigma^G_{\mathcal{C}}[f](x,\alpha) := \arg\min_{w \in \mathbb{R}^n} \quad & \|f(x) - w\|^2_{G(x)} && \text{(6.2a)} \\ \text{subject to} \quad & x + \alpha w \in \mathcal{U} && \text{(6.2b)} \\ & k(x) + \alpha \nabla k(x) w \leq 0\,. && \text{(6.2c)} \end{aligned}$$

We claim that, for small α and under various assumptions, the discrete-time system

$$x^+ = x + \alpha\, \Sigma^G_{\mathcal{C}}[f](x,\alpha) \tag{6.3}$$

approximates the projected dynamical system

$$\dot{x} = \Pi^G_{\mathcal{C}}\left[f\right](x)\,, \qquad x \in \mathcal{C}\,. \tag{6.4}$$

More precisely, we will show in the following that interpolated solutions of (6.3) converge uniformly to a solution of (6.4) as $\alpha \to 0^+$ and we show that a gas set of (6.4) is spas with respect to the discrete approximation (6.3).

The key idea behind the constraint structure of $\mathcal{C}$ is that, in a feedback setup, $\mathcal{U}$ denotes constraints on plant inputs that can be enforced directly. On the other hand, constraints on the outputs (or possibly joint constraints on inputs and outputs) can be enforced only approximately due to disturbances and model uncertainty, hence, it is reasonable to linearly approximate these output constraints as a best guess. We therefore refer to (6.3) as a *linearized output projection* (LOP) *discretization* of (6.4). The following example illustrates how constraints on a physical system can be brought into the form (6.1).

Example 6.1. Consider an algebraic plant with a steady-state input-to-output map $h : \mathcal{U} \to \mathbb{R}^m$ and let $\mathcal{U} \subset \mathbb{R}^n$ denote the set of admissible inputs. Further, assume that outputs y (and inputs u) need to satisfy the constraints $\tilde{g}(y, u) \leq 0$, where $\tilde{g} : \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}^q$ is continuously differentiable. Define $k : \mathbb{R}^n \to \mathbb{R}^q$ as $k(u) := \tilde{g}(h(u), u)$. Thus, we arrive at a set $\mathcal{C}$ of the form (6.1) that combines constraints on the inputs and joint constraints on the inputs and outputs.

Next, let f be of the form $f(u) = \tilde{f}(h(u), u)$ where $\tilde{f} : \mathcal{U} \times \mathbb{R}^m \to \mathbb{R}^n$. Together with the metric $G : \mathcal{U} \to \mathbb{S}^n_+$, the vector field $\tilde{f}$ defines a desired (unconstrained) control law. For example, f might be the negative gradient of an objective function to be minimized (Chapter 10).

With these definitions, we can write (6.3) equivalently as an integral control law

$$u^+ = u + \alpha \, \tilde{\Sigma}^G_{\mathcal{C}}(y, u, \alpha) \quad y = h(u) \tag{6.5}$$

with $\tilde{\Sigma}^G_{\mathcal{C}}(y, u, \alpha)$ as the *measurement-based* evaluation of (6.2), i.e.,

$$\begin{aligned} \tilde{\Sigma}^G_{\mathcal{C}}(y, u, \alpha) := \arg\min_{w \in \mathbb{R}^n} \; & \|w - \tilde{f}(y, u)\|^2_{G(u)} \\ \text{subject to} \quad & u + \alpha w \in \mathcal{U} \\ & \tilde{g}(y, u) + \alpha \underbrace{\nabla \tilde{g}(y, u) \begin{bmatrix} \nabla h(u) \\ \mathbb{I}_n \end{bmatrix}}_{\nabla k(u)} w \leq 0 \,, \end{aligned}$$

where $y = h(u)$ is the measured system output. In particular, the evaluation of $\tilde{\Sigma}^G_{\mathcal{C}}(y, u, \alpha)$ does not directly rely on h. Instead, it is enough to know ∇h.

Furthermore, note that $\tilde{\Sigma}^G_{\mathcal{C}}(y, u, \alpha) = 0$ implies that $u \in \mathcal{U}$ and $\tilde{g}(y, u) \leq 0$. In other words, any equilibrium point of the closed-loop system has to satisfy both the constraints on the inputs as well as the joint constraints on the inputs and outputs. ■

Remark 6.1 (Related Work on Discretizations of PDS). The numerical integration of ODEs is a classical, well-studied topic with wide-ranging results [236]. Similarly, numerical methods for unconstrained inclusions have been studied extensively [84]. However, for constrained inclusions and projected dynamical systems in particular, there seem to be much

fewer results. A few general methods are proposed in [18]. Besides that, [186] seems to deliver the most comprehensive results for PDS.

Namely, one well-established approach to discretize projected dynamical system such as (6.4) (but only for $G \equiv \mathbb{I}$ and $\mathcal{C}$ convex) is the *projected Euler integration* of the form

$$x^+ = P_{\mathcal{C}}\left(x + \alpha f(x)\right) , \tag{6.6}$$

where an unconstrained forward Euler step is followed by a projection onto the feasible set.

Similarly, one can also generalize Runge-Kutta methods and other numerical integration schemes to approximate projected dynamical systems. For these methods, uniform convergence of (interpolated) solutions to the trajectories of (6.4) is well-known has been established in [186]. Finite-difference methods for constrained inclusions have also been sketched in [18]. However, none of these schemes can be implemented as feedback control loops. For instance, in the context of Example 6.1, the projected Euler method takes the form

$$u^+ = P_{\mathcal{C}}\left(u + \alpha \hat{f}(y,u)^T\right) \qquad y = \hat{h}(u). \tag{6.7}$$

This approach, however, is not practical for feedback optimization, because it requires the controller to perform a projection on $\mathcal{C}$ which might not tractable unless $\mathcal{C}$ is convex. In addition, full knowledge of k is required to evaluate $P_{\mathcal{C}}$ (as opposed to the evaluation of $\hat{\Sigma}_{\mathcal{C}}^{G}(u,y,\alpha)$ which requires only ∇h) and generally results in a lack of robustness and entails performance degradation. ■

6.1 Continuity of Approximate Projections

To establish the properties of the LOP discretization, we first need to establish properties of the approximate projection operator $\Sigma_{\mathcal{C}}^{G}[f]$. For this purpose we need to make several assumptions on the problem setup.

Assumption 6.1. *Given the discrete-time system (6.3) with $\mathcal{C}$ of the form (6.1), the set $\mathcal{C}$ is non-empty and satisfies LICQ. Further, ∇k, ∇g, f and G are locally Lipschitz.*

In particular, under Assumption 6.1 and by Proposition 3.2, the set $\mathcal{C}$ is prox-regular and Theorem 4.2 guarantees the existence and uniqueness of solutions of (6.4).

Furthermore, we make the following assumption:

Assumption 6.2. *Given the discrete-time system (6.3) with $\mathcal{C}$ of the form (6.1), the set $\mathcal{U} = \{x \,|\, g(x) \leq 0\}$ is compact.*

Assuming that $\mathcal{U}$ is compact is both practically relevant and technically convenient: On one hand, since no realistic physical system can handle unbounded inputs, it is reasonable to assume that the set of admissible inputs is always bounded. On the other hand, compactness of $\mathcal{U}$ simplifies our analysis because continuous maps on $\mathcal{U}$ are automatically uniformly continuous, have upper and lower bounds, etc. and local Lipschitz continuity implies (global) Lipschitz continuity on $\mathcal{U}$. Nevertheless, the compactness assumption can presumably be relaxed using similar ideas as in Chapter 5 by passing to a local setup (see Assumption 5.1).

Hence, by applying Corollary 4.2, we note that under Assumptions 6.1 and 6.2 the PDS (6.4) admits unique complete solutions.

Corollary 6.1. *Under Assumptions 6.1 and 6.2 the continuous-time PDS (6.4) admits a unique complete (Caratheódory) solution for every initial condition $x(0) \in \mathcal{C}$.*

Moreover, for (6.3) to be well-defined, we make the following assumption:

Assumption 6.3. *The set $\mathcal{U} = \{x \,|\, g(x) \leq 0\}$ is convex and, for all $x \in \mathcal{U}$ and all $\alpha > 0$, the feasible set of (6.2) defined as*

$$\tilde{\mathcal{C}}(x, \alpha) := \{w \,|\, g(x + \alpha w) \leq 0,\, k(x) + \alpha \nabla k(x) w \leq 0\}$$

is non-empty and satisfies LICQ.

Remark 6.2. Although the LICQ requirement in Assumption 6.3 is in general not obvious to verify, it is commonly made in the study of *sequential quadratic programming* (SQP) algorithms [250, 191, 31]. Furthermore, presumably, the genericity of LICQ for parametrized feasible sets can be exploited to weaken Assumption 6.3. ■

Invariance of $\mathcal{U}$ under (6.3) is immediate if Assumption 6.3 holds:

Lemma 6.1. *Under Assumption 6.3 we have $x^+ \in \mathcal{U}$ for all $x \in \mathcal{U}$, all $\alpha > 0$ x^+ as in (6.3).*

Furthermore, under Assumptions 6.1 and 6.3, $\Sigma_{\mathcal{C}}^G[f](x,\alpha)$ is single-valued for all $x \in \mathcal{U}$ and all $\alpha > 0$ since $\mathcal{U}$ (and hence $\tilde{\mathcal{C}}(x,\alpha)$) is convex and non-empty, and the objective is strongly convex. Moreover, $\Sigma_{\mathcal{C}}^G[f]$ is continuous:

Lemma 6.2. *Under Assumptions 6.1 and 6.3, $\Sigma_{\mathcal{C}}^G[f](x,\alpha)$ is continuous in (x,α) and the map of Lagrange multipliers of (6.2) is continuous. Further, for all $x \in \mathcal{C}$ it holds that*

$$\lim_{\alpha \to 0^+} \Sigma_{\mathcal{C}}^G[f](x,\alpha) = \Pi_{\mathcal{C}}^G[f](x)\,. \tag{6.8}$$

Thus, (6.8) offers a first glimpse at why the LOP discretization (6.3) approximates the PDS (6.4).

Proof. We can apply the sensitivity result Corollary 2.2 for parametrized convex problems. Namely, (6.2) is parametrized in (x,α), the problem is feasible and $\tilde{\mathcal{C}}(x,\alpha)$ satisfies LICQ for all (x,α) by Assumption 6.3. The objective of (6.2) is strongly convex (in w) for all (x,α) and $\tilde{\mathcal{C}}(x,\alpha)$ is convex as a set (by linearization of k and convexity of $\mathcal{U}$). Continuity of $\Sigma_{\mathcal{C}}^G[f]$ and the associated Lagrange multipliers follows directly from Corollary 2.2.

To show (6.8) holds, we reformulate (6.2) as

$$\begin{aligned} \Sigma_{\mathcal{C}}^G[f](x,\alpha) := \arg\min_{w \in \mathbb{R}^n} \quad & \|f(x) - w\|_{G(x)}^2 && \text{(6.9a)} \\ \text{subject to} \quad & w \in (\mathcal{U} - x)/\alpha && \text{(6.9b)} \\ & \nabla k(x) w \leq -\tfrac{1}{\alpha} k(x)\,. && \text{(6.9c)} \end{aligned}$$

For (6.9b) we recall that the definition of the tangent cone of the Clarke regular set $\mathcal{U}$ as $T_x\mathcal{U} := \lim_{\alpha \to 0^+} \frac{1}{\alpha}(\mathcal{U} - x)$ (Definition 2.4) and therefore conclude that $w \in T_x\mathcal{U}$ as $\alpha \to 0^+$.

Next, for α small enough, we can split (6.9c) into active and inactive constraints. For instance, if $k(x) = 0$, then we have, as $\alpha \to 0^+$, that

$\nabla k(x)w = 0$. On the other hand, if $k(x) < 0$ (recall that $x \in \mathcal{C}$, by assumption), the constraint is relaxed to infinity as $\alpha \to 0^+$ and can be omitted in the limit. Consequently, we have $\nabla k_{\mathbf{I}(x)}(x)w \leq 0$ where $\mathbf{I}(x) := \{i \,|\, k_i(x) = 0\}$ denotes the indices of the active constraints k_i at x. Since LICQ holds for $\mathcal{Y} := \{x \,|\, k(x) \leq 0\} \subset \mathcal{C}$ it follows from the explicit expression for the tangent cone in Example 2.2 that $w \in T_x\mathcal{Y}$ for $\alpha \to 0^+$.

Hence, in the limit we have $w \in T_x\mathcal{U}$ and $w \in T_x\mathcal{Y}$. Using the fact that under LICQ we have $T_x\mathcal{C} = T_x(\mathcal{U} \cap \mathcal{Y}) = T_x\mathcal{U} \cap T_x\mathcal{Y}$ [211, Thm 6.42] it follows that, $\lim_{\alpha\to 0^+} \Sigma_{\mathcal{C}}^{G}[f](x, \alpha)$ is a projection of $f(x)$ onto $T_x\mathcal{C}$ with respect to G. □

6.2 Linearized Output Projection as Perturbed PDS

Similarly to Section 5.2 in the previous chapter, we show that the LOP system (6.3) can be interpreted a perturbed PDS. The details of this process are different from Section 5.2 because we have to pass from a discrete-time to a continuous-time system. But the key arguments leading to uniform convergence and semiglobal practical asymptotic stability are the same.

We start this proof with the key technical result that states that the distance to a set defined by inequality constraints can be locally upper bounded by the constraint violation:

Proposition 6.1. *Consider a set $\mathcal{X} := \{x \,|\, b(x) \leq 0\}$ where $b : \mathbb{R}^n \to \mathbb{R}^r$ is continuously differentiable. Let $\mathcal{X}$ be compact and satisfy LICQ. Then, there exist $\alpha, \gamma > 0$ such that for all $x \in \mathcal{X} + \alpha \operatorname{int} \mathbb{B}$ it holds that $d_{\mathcal{X}}(x) \leq \gamma \|\max\{b(x), 0\}\|$.*

The proof of Proposition 6.1 is conceptually simple but technically tedious: The main idea is to locally "straighten" the set $\mathcal{X}$ by applying local coordinate transformations that map subsets of $\mathcal{X}$ to the non-negative orthant. The distance of any point to the non-negative orthant can be expressed explicitly, and, because the coordinate transformation

has a Lipschitz inverse, the distance $d_{\mathcal{X}}(x)$ in the original domain can be bounded.

Proof sketch. Recall that equality-constrained sets of the form $\mathcal{X}_{eq} = \{x \,|\, b(x) = 0\}$ with $b : \mathbb{R}^n \to \mathbb{R}^r$ are smooth $n-r$-dimensional submanifolds of $\mathbb{R}^n$ if b has constant rank on $\mathcal{X}$ [48, Thm. 5.8]. Namely, for every $y \in \mathcal{X}_{eq}$ the exists a chart $(\mathcal{U}, \phi)$ where $\mathcal{U} \subset \mathbb{R}^n$ is open, $y \in \mathcal{U}$, and $\phi : \mathcal{U} \to \mathbb{R}^n$ is such that $\phi(y) = 0$ and $\phi(\mathcal{U} \cap \mathcal{X}_{eq}) \subset \{x \,|\, x_{r+1} = \ldots = x_n = 0\}$. This fundamental result is a consequence of the rank theorem [48, Thm 7.1].

Similarly, a set of the form $\mathcal{X} := \{x \,|\, b(x) \leq 0\}$ for which LICQ holds, defines a *submanifold with corners* [162, Ch. 16]. In particular, around every point $\hat{x} \in \mathcal{X}$ there exists a neighborhood $\mathcal{U}$ of $\hat{x}$ such that for all $x \in \mathcal{U}$ we have $b_i(x) < 0$ for all $i \notin \hat{\mathbf{I}} := \mathbf{I}(\hat{x})$ and $\nabla b_{\hat{\mathbf{I}}}(x)$ has full rank. In particular, on $\mathcal{U}$, no constraint is active other than those active at $\hat{x}$ and the constraints active at $\hat{x}$ have full rank on the entire neighborhood. Furthermore, there exists a $C^{1,1}$ chart $\phi : \mathcal{U} \to \mathbb{R}^n$ such that $\phi(\hat{x}) = 0$ and, for all $x \in \mathcal{U}$, we have

$$x \in \mathcal{U} \cap \mathcal{X} \quad \Leftrightarrow \quad \phi(x) \in \mathbb{R}^{|\hat{\mathbf{I}}|}_{\geq 0} \times \mathbb{R}^{n-|\hat{\mathbf{I}}|}.$$

In other words, ϕ maps the feasible set $\mathcal{U} \cap \mathcal{X}$ to the non-negative orthant, as illustrated in Figure 6.1a. Furthermore, without loss of generality, we may assume that for all $x \in \mathcal{U}$ we have

$$[\phi(x)]_{1:|\hat{\mathbf{I}}|} = -b_{\hat{\mathbf{I}}}(x),$$

i.e., the first $|\hat{\mathbf{I}}|$ components of ϕ corresponds to the negative constraint values $-b_{\hat{\mathbf{I}}}(x)$.

Next, consider an open ball $\mathcal{V} := \epsilon \operatorname{int} \mathbb{B}$ with radius $\epsilon > 0$ in the co-domain of ϕ such that $\operatorname{cl} \mathcal{V} = \epsilon \mathbb{B} \subset \phi(\mathcal{U})$. Any point inside of $\mathcal{V}$ can be easily projected onto the (transformed) feasible set $\mathcal{W} := \phi(\mathcal{X} \cap \mathcal{U}) \cap \mathcal{V}$.

Namely, consider $x \in \phi^{-1}(\mathcal{V})$ and define $y := \phi(x) \subset \mathcal{V}$. The projection of y onto $\mathcal{W}$ is given by $P_{\mathcal{W}}(y) = \Big(\max\{y_{1:|\hat{\mathbf{I}}|}, 0\}, y_{|\hat{\mathbf{I}}|:n}\Big)$. Namely, the first $|\hat{\mathbf{I}}|$ components of y are projected onto the non-negative orthant

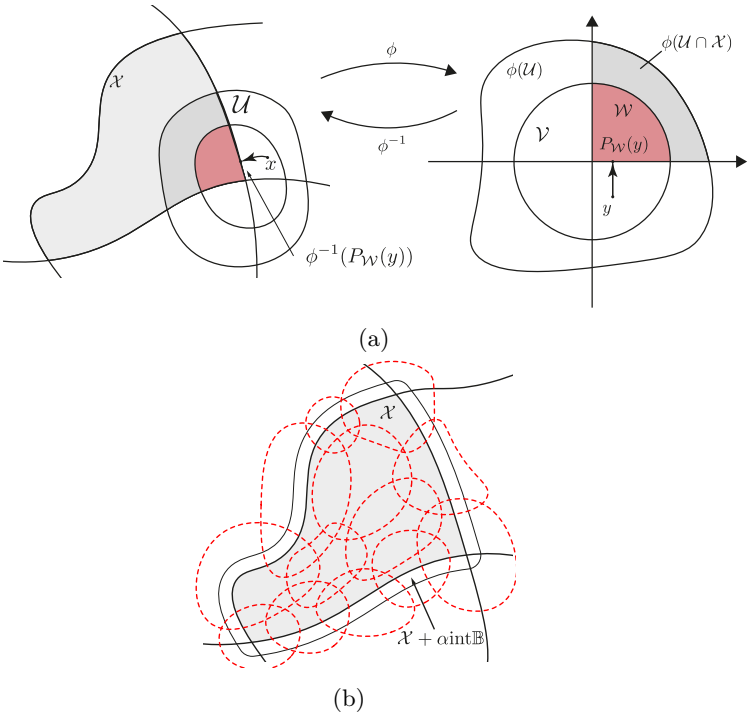

Figure 6.1: Illustrations for proof sketch of Proposition 6.1: (a) coordinate neighborhoods and transformations (b) finite subcover and α-neighborhood of $\mathcal{X}$

and the remaining components are unaffected. This insight, allows us to write

$$\begin{aligned}\|y - P_{\mathcal{W}}(y)\| &= \left\| y_{1:|\hat{\mathbf{i}}|} - \max\left\{ y_{1:|\hat{\mathbf{i}}|}, 0 \right\} \right\| \\ &= \left\| \min\left\{ y_{1:|\hat{\mathbf{i}}|}, 0 \right\} \right\| = \left\| \max\left\{ b_{\hat{\mathbf{i}}}(x), 0 \right\} \right\| .\end{aligned}$$

Next, recall that ϕ is a C^1 diffeomorphism on $\mathcal{U}$. Consequently, ϕ^{-1} is ℓ-Lipschitz on the compact set $\epsilon\mathbb{B}$ (and hence on the set $\epsilon \,\text{int}\,\mathbb{B}$) for

some $\ell > 0$. Therefore, we have

$$\begin{aligned}\|\phi^{-1}(y) - \phi^{-1}(P_{\mathcal{W}}(y))\| &\leq \ell \|y - P_{\mathcal{W}}(y)\| \\ &= \ell \left\|\max\left\{b_{\hat{\mathbf{I}}}(x), 0\right\}\right\| = \ell \|\max\{b(x), 0\}\|\,,\end{aligned}$$

where the last equality holds because on $\mathcal{U}$, by definition, $b_i(x) < 0$ for all $i \notin \hat{\mathbf{I}}$.

Since, $\phi^{-1}(y) = x$ and $\phi^{-1}(P_{\mathcal{W}}(y)) \in \mathcal{X}$, we conclude that for all $x \in \phi^{-1}(\mathcal{V})$ we have $d_{\mathcal{X}}(x) \leq \ell \|\max\{b(x), 0\}\|$.

This construction of a (restricted) chart $(\phi^{-1}(\mathcal{V}), \phi)$ is possible around every point $\hat{x} \in \mathcal{X}$. Hence, let $(\phi_{\hat{x}}^{-1}(\mathcal{V}_{\hat{x}}), \phi_{\hat{x}})$ denote such a chart centered at $\hat{x}$ and note that the union of all $\phi_{\hat{x}}^{-1}(\mathcal{V}_{\hat{x}})$ forms an *open cover* of $\mathcal{X}$, i.e., $\mathcal{X} \subset \bigcup_{\hat{x}\in\mathcal{X}} \phi_{\hat{x}}^{-1}(\mathcal{V}_{\hat{x}})$.

Since $\mathcal{X}$ is assumed to be compact and by the definition of compactness, there exists a *finite subcover* $\{\phi_{\hat{x}_i}^{-1}(\mathcal{V}_{\hat{x}_i})\}_{i=1}^{M}$ of chart domains that covers $\mathcal{X}$ (Figure 6.1b). Moreover, because $\bigcup_{i=1}^{M} \phi_{\hat{x}_i}^{-1}(\mathcal{V}_{\hat{x}_i})$ is open, there exists $\alpha > 0$ such that $\mathcal{X} + \alpha \operatorname{int}\mathbb{B} \subset \bigcup_{i=1}^{M} \phi_{\hat{x}_i}^{-1}(\mathcal{V}_{\hat{x}_i})$. Finally, we can choose $\gamma = \max_{i=1,\dots,M} \ell_{\hat{x}_i}$ where $\ell_{\hat{x}_i}$ is the Lipschitz constant of $\phi_{\hat{x}_i}^{-1}$ on $\mathcal{V}_{\hat{x}_i}$.

To summarize, every $x \in \mathcal{X} + \alpha \operatorname{int}\mathbb{B}$ lies in at least one chart domain $\phi_{\hat{x}_i}^{-1}(\mathcal{V}_{\hat{x}_i})$. On this chart domain it holds that

$$d_{\mathcal{X}}(x) \leq \ell_{\hat{x}_i} \|\max\{b(x), 0\}\| \leq \gamma \|\max\{b(x), 0\}\|\,. \qquad \square$$

For our problem setup we need to modify Proposition 6.1 as follows:

Corollary 6.2. *Under Assumptions 6.1 and 6.2 there exist $\delta, \gamma > 0$ such that for all $x \in \mathcal{U}$ for which $\|\max\{k(x), 0\}\| < \delta$ it holds that $d_{\mathcal{C}}(x) \leq \gamma \|\max\{k(x), 0\}\|$.*

Proof. Starting from Proposition 6.1, define

$$\delta := \min_{x\in\mathcal{U}\setminus(\mathcal{X}+\alpha \operatorname{int}\mathbb{B})} \|\max\{b(x), 0\}\| > 0\,,$$

i.e., δ is the smallest constrained violation outside of $\mathcal{X} + \alpha \operatorname{int}\mathbb{B}$. This minimum exists by compactness of $\mathcal{U} \setminus (\mathcal{X} + \alpha \operatorname{int}\mathbb{B})$ and is positive

because it is over a set that does not contain any feasible point. Consequently, for all $x \in \mathcal{U}$ for which $\|\max\{b(x), 0\}\| < \delta$ it follows that $x \in \mathcal{U} \cap (\mathcal{C} + \alpha \operatorname{int} \mathbb{B}$ and $d_{\mathcal{X}}(x) \leq \gamma \|\max\{b(x), 0\}\|$ applies. □

The following result is a generalization of Lemma 3.5 to vector-valued functions. The proof in [198] generalizes directly and is provided for completeness only.

Lemma 6.3. *Let $\mathcal{V} \subset \mathbb{R}^n$ be open and let $g : \mathcal{V} \to \mathbb{R}^m$ be $C^{1,1}$ and let ∇g be ℓ-Lipschitz. Then,*

$$\|g(x') - g(x) - \nabla g(x)(x' - x)\| \leq \tfrac{1}{2}\ell \|x' - x\|^2 .$$

Proof. By defining $v := x' - x$ and the curve $\gamma : [0, 1] \to \mathbb{R}$ as $\gamma(t) := g(x + tv)$, we may write $\int_0^1 \langle \nabla g(x + \tau v), v \rangle \, d\tau = \int_0^1 \gamma(\tau) d\tau = \gamma(1) - \gamma(0) = g(x') - g(x)$. Hence, we have

$$g(x') - g(x) - \int_0^1 \langle \nabla g(x), v \rangle \, d\tau = \int_0^1 \langle \nabla g(x + \tau v) - \nabla g(x), v \rangle \, d\tau ,$$

and by taking the norm on both sides and using the ℓ-Lipschitz continuity of ∇g we get

$$\begin{aligned} &\Big\| g(x') - g(x) - \underbrace{\int_0^1 \langle \nabla g(x), v \rangle}_{= \langle g(x), x' - x \rangle} \Big\| \\ &\qquad = \Big\| \int_0^1 \underbrace{\langle \nabla g(x + \tau v) - \nabla g(x), v \rangle}_{\leq \|\nabla g(x+\tau v) - \nabla g(x)\| \|v\| \leq \ell \|v\|^2 \tau} \, d\tau \Big\| \leq \tfrac{1}{2}\ell \|x' - x\|^2 \end{aligned}$$

which completes the proof. □

Lemma 6.3 allows us to bound the constraint violation of every iterate of (6.3):

Lemma 6.4. *Under Assumptions 6.1 and 6.3 and $x \in \mathcal{U}$, the iteration (6.3) satisfies*

$$\big\|\max\{k(x^+), 0\}\big\| \leq \tfrac{1}{2}\ell\alpha^2 \left\|\Sigma_{\mathcal{C}}^{G}[f](x, \alpha)\right\|^2 .$$

Proof. Applying the vector-valued Descent Lemma (Lemma 6.3) we have

$$\begin{aligned}&\left\|k(x^+) - (k(x) + \nabla k(x)(x^+ - x))\right\| \\ &\qquad\qquad \leq \tfrac{1}{2}\ell\|x^+ - x\|^2 = \tfrac{1}{2}\ell\left\|\alpha\Sigma_{\mathcal{C}}^G[f](x,\alpha)\right\|^2.\end{aligned}$$

Using the fact that $k(x) + \nabla k(x)(x^+ - x) \leq 0$ (by definition of x^+ and $\Sigma_{\mathcal{C}}^G[f]$ in (6.3) and (6.2), respectively) we have

$$\begin{aligned}&\left\|k(x^+) - (k(x) + \nabla k(x)(x^+ - x))\right\| \\ &\quad\geq \left\| \max\left\{k(x^+) - \underbrace{(k(x) + \nabla k(x)(x^+ - x))}_{\leq 0}, 0\right\}\right\| \\ &\quad\geq \left\|\max\left\{k(x^+), 0\right\}\right\|.\end{aligned}$$

To see that the second inequality holds, consider $a, b \in \mathbb{R}$ with $b \leq 0$. Then, we have

(i) if $a \geq 0$, then $a - b \geq a \geq 0$ and therefore $|\max\{a - b, 0\}| \geq |\max\{a, 0\}|$,

(ii) if $a \leq 0$ and $a - b \geq 0$, we have $\max\{a, 0\} = 0$ and thus $\max\{a - b, 0\} \geq \max\{a, 0\}$, and

(iii) if $a \leq 0$ and $a - b < 0$, then trivially, we have $\max\{a - b, 0\} = \max\{a, 0\} = 0$.

This argument applies componentwise to the inequality above with $a = k_i(x^+)$, $b = k_i(x) + \nabla k_i(x)(x^+ - x)$ and finishes the proof. □

Since $x \in \mathcal{U}$ and $\mathcal{U}$ is compact by Assumption 6.2, $\left\|\Sigma_{\mathcal{C}}^G[f](x,\alpha)\right\|$ is bounded and therefore the constraint violation in Lemma 6.4 can be upper bounded by a function of α.

Next, we combine Lemma 6.4 and Corollary 6.2:

Lemma 6.5. *Consider the discrete-time system (6.3) and let Assumptions 6.1, 6.2, and 6.3 hold. Then, there exist $\sigma \in \mathcal{K}_\infty$ and $\alpha^\star > 0$ such that for all $\alpha \in (0, \alpha^\star)$ and for every iterate $x^+ = x + \alpha\Sigma_{\mathcal{C}}^G[f](x,\alpha)$, we have that $d_{\mathcal{C}}(x^+) \leq \sigma(\alpha)$.*

Proof. By compactness of $\mathcal{U}$ and continuity of $\Sigma_{\mathcal{C}}^{G}[f]$ in x (Lemma 6.2) there exists, for every $\alpha > 0$, an upper bound $\zeta_\alpha > 0$ such that $\|\Sigma_{\mathcal{C}}^{G}[f](x,\alpha)\|^2 \leq \zeta_\alpha$ for $x \in \mathcal{U}$. From Lemma 6.4 it thus follows that $\|\max\{k(x^+),0\}\| \leq \ell\zeta_\alpha\alpha^2$.

By Corollary 6.2, $d_{\mathcal{C}}(x^+) \leq \gamma\|\max\{k(x^+),0\}\| \leq \gamma\zeta_\alpha\ell\alpha^2$ holds for some $\gamma > 0$ and whenever $\ell\zeta_\alpha\alpha^2 \leq \delta$ for some $\delta > 0$. Hence, we define σ such that $\sigma(\alpha) \geq \zeta_\alpha\ell\gamma\alpha^2$ for all $\alpha < \alpha^\star := \sqrt{\zeta_\alpha\ell\gamma/\delta}$ to finish the proof. □

Finally, we can prove the key result that $\Sigma_{\mathcal{C}}^{G}[f]$ is perturbation of $\Pi_{\mathcal{C}}^{G}[f]$:

Proposition 6.2. *Consider the discrete-time system (6.3) and let Assumptions 6.1, 6.2, and 6.3 hold. Then, there exist $\rho \in \mathcal{K}_\infty$ and $\alpha^\star > 0$ such that for all $\alpha \in (0,\alpha^\star)$ we have*

$$\Sigma_{\mathcal{C}}^{G}[f](x,\alpha) \subset \operatorname{co}\Pi_{\mathcal{C}}^{G}[f]\left((x+\rho(\alpha)\mathbb{B})\cap\mathcal{C}\right)+\rho(\alpha)\mathbb{B}\,. \tag{6.10}$$

Furthermore, the set $\mathcal{C}_\alpha^\circ := (\mathcal{C}+\rho(\alpha)\mathbb{B})\cap\mathcal{U}$ is invariant for (6.3).

Proof. The invariance of $\mathcal{C}_\alpha^\circ$ is immediate from Lemma 6.5. In particular, independently of the initial condition, the solution of (6.3) enters $\mathcal{C}_\alpha^\circ$ after at most one iteration.

To show (6.10) we first recall that $\Sigma_{\mathcal{C}}^{G}[f]$ is continuous in (x,α) (Lemma 6.2). In particular, since $\mathcal{U}$ is compact, $\Sigma_{\mathcal{C}}^{G}[f]$ is uniformly continuous on any compact set $\mathcal{U}\times[0,\hat{\alpha}]$ with $\hat{\alpha} > \alpha^\star$ large enough. Namely, there exists a *modulus of continuity* $\omega \in \mathcal{K}_\infty$ such that

$$\left\|\Sigma_{\mathcal{C}}^{G}[f](x',\alpha')-\Sigma_{\mathcal{C}}^{G}[f](x,\alpha)\right\| \leq \omega\left(\left\|\begin{bmatrix}x'-x\\ \alpha'-\alpha\end{bmatrix}\right\|\right)$$

for all $x', x \in \mathcal{U}$ and all $\alpha', \alpha \in [0,\hat{\alpha}]$. Let $\sigma \in \mathcal{K}_\infty$ as in Lemma 6.5. Then, we have

$$\begin{aligned}
\Sigma_{\mathcal{C}}^{G}[f](x,\alpha) &\in \Sigma_{\mathcal{C}}^{G}[f](P_{\mathcal{C}}(x),0)+\omega\left(\sqrt{d_{\mathcal{C}}^2(x)+\alpha^2}\right)\mathbb{B} && (6.11)\\
&\subset \Pi_{\mathcal{C}}^{G}[f]\left((x+\sigma(\alpha)\mathbb{B})\cap\mathcal{C}\right)+\omega\left(\sqrt{\sigma(\alpha)^2+\alpha^2}\right)\mathbb{B} && (6.12)\\
&\subset \operatorname{co}\Pi_{\mathcal{C}}^{G}[f]\left((x+\rho(\alpha)\mathbb{B})\cap\mathcal{C}\right)+\rho(\alpha)\mathbb{B}\,, && (6.13)
\end{aligned}$$

where $\rho(\alpha) := \max\left\{\sigma(\alpha), \omega\left(\sqrt{\sigma(\alpha)^2 + \alpha^2}\right)\right\}$. The relation (6.11) follows directly from the definition of ω. Note that $P_{\mathcal{C}}(x)$ is a set, however for every $\hat{x} \in P_{\mathcal{C}}(x)$ it holds that $\|\hat{x} - x\| = d_{\mathcal{C}}(x)$. The inclusion (6.12) follows from Lemma 6.2, the fact that $P_{\mathcal{C}}(x) \subset (x + d_{\mathcal{C}}(x)\mathbb{B}) \cap \mathcal{C}$, the fact that $d_{\mathcal{C}}(x) \leq \sigma(\alpha)$ according to Lemma 6.5, and the fact that $b \mapsto \omega(\sqrt{(b^2 + \alpha^2)})$ is monotonically increasing. The relation (6.13) follows by definition of ρ and the fact that the convex hull of a set always contains the set itself.

Finally, we have $\rho \in \mathcal{K}_\infty$ because of the following facts [229]: If $\sigma \in \mathcal{K}_\infty$, then $\sigma^2 \in \mathcal{K}_\infty$. In addition the sum of two $\mathcal{K}_\infty$-functions is $\mathcal{K}_\infty$ and therefore $\alpha \mapsto \sigma(\alpha)^2 + \alpha^2$ is $\mathcal{K}_\infty$. Next, the square root is a $\mathcal{K}_\infty$ and therefore the composition $\alpha \mapsto \sqrt{\sigma(\alpha^2) + \alpha^2}$ is $\mathcal{K}_\infty$. Finally, the maximum of two $\mathcal{K}_\infty$-functions is a $\mathcal{K}_\infty$-function. □

6.3 Uniform Convergence

We can now show that solutions of the LOP discretization (6.3) approximate the continuous-time PDS (6.4). Given a solution $x_\alpha : \mathbb{N} \to \mathcal{U}$ of (6.3) for some $\alpha > 0$, let the interpolation of x_α be defined as $\hat{x}_\alpha : \mathbb{R}_{\geq 0} \to \mathcal{U}$ with

$$\hat{x}_\alpha(t) := x_\alpha[k] + \tfrac{t-\alpha k}{\alpha}(x_\alpha[k+1] - x_\alpha[k]) \qquad \forall t \in [\alpha k, \alpha(k+1))\,.$$

For this purpose, we need the following result which we derive from Lemma 5.4 on Page 111, but which can also be obtained from simplifying[1] [216, Thm. 5.2].

Lemma 6.6. *Consider $\mathcal{C} \subset \mathbb{R}^n$ and $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and let the inclusion*

$$\dot{x} \in H(x)\,, \quad x \in \mathcal{C} \tag{6.14}$$

be well-posed and let $\mathcal{C}$ be a viability domain for H. Further, consider a class of discrete-time systems, parametrized in $\delta > 0$, and given by

$$x^+ = H_\delta(x) \quad x \in \mathcal{C}_\delta \tag{6.15}$$

[1] The original result applies to hybrid inclusions and possibly incomplete, non-unique (hybrid) trajectories.

with $\mathcal{C}_\delta \subset \mathbb{R}^n$ and $H_\delta : \mathcal{C}_\delta \to \mathcal{C}_\delta$. Assume that there exists $\rho \in \mathcal{K}_\infty$ and $\delta^\star > 0$ such that for all $\delta \in (0, \delta^\star)$ it holds that $\mathcal{C}_\delta \subset \mathcal{C} + \rho(\delta)\mathbb{B}$ and for every $x \in \mathcal{C}_\delta$ we have

$$H_\delta(x) \subset x + \delta \operatorname{co} H((x + \rho(\delta)\mathbb{B}) \cap \mathcal{C}) + \delta\rho(\delta)\mathbb{B}\,. \tag{6.16}$$

Given a sequence $\delta_i \to 0^+$, let $\hat{x}_{\delta_i}$ denote the interpolated solution of (6.15) starting at $x_0 \in \mathcal{C}$. Then, for any $T, \epsilon > 0$, if the sequence $\{\operatorname{gph} \hat{x}_{\delta_i} \cap (x_0 + \epsilon\mathbb{B}) \times [0, T]\}$ converges graphically, then its convergence is to a (T, ϵ)-truncated solution $x : [0, T] \to \mathcal{C}$ of (6.14).

Proof. First, we need to show that for a given δ with $0 < \delta < \delta^\star$, the interpolated solution $\hat{x}_\delta$ of (6.15) is a solution of a $\delta\rho$-perturbation of (6.14). For this purpose, let $t \in (\delta k, \delta(k+1))$ for some $k \geq 0$. Then x_δ satisfies $\dot{\hat{x}}_\delta(t) = (x_\delta[k+1] - x_\delta[k])/\delta$ and therefore, by (6.16), we have

$$\dot{\hat{x}}_\delta(t) = \tfrac{1}{\delta}\left(H_\delta(x_\delta[k]) - x_\delta[k]\right) \subset \operatorname{co} H(x_\delta[k] + \rho(\delta)) \cap \mathcal{C}) + \rho(\delta)\mathbb{B}\,.$$

Next, we limit ourselves to the (T, ϵ)-truncation of $\hat{x}_\delta$, i.e., we consider the map $\hat{x}'_\delta : T' \to x_0 + \epsilon\mathbb{B}$ such that $\operatorname{gph} \hat{x}'_\delta = \operatorname{gph} \hat{x}_\delta$ and either $T' = T$ or $\|\hat{x}'_\delta(T') - x_0\| = \epsilon$. Define

$$\sigma'(\delta) := \sup_{x \in x_0 + \epsilon\mathbb{B}} \|\operatorname{co} H((x + \rho(\delta)\mathbb{B}) \cap \mathcal{C}) + \rho(\delta)\|\,.$$

This allows us establish the bound $\|x_\delta[k] - \hat{x}'_\delta(t)\| \leq \|x_\delta[k] - x_\delta[k+1]\| \leq \delta\sigma'(\delta)$ for all $t \in [0, T']$ (and k such that $t \in [\delta k, \delta(k+1)]$). Therefore, for all $t \in [0, T']$ and all $\delta < \delta^\star$, we have $\hat{x}'_\delta(t) \in \{x \,|\, x + \delta\sigma'(\delta)\mathbb{B} \cap \mathcal{C}_\delta \neq \emptyset\}$ as well as

$$\begin{aligned}\dot{\hat{x}}'_\delta(t) \in \overline{\operatorname{co}}\, H(\hat{x}'_\delta(t) + (\delta\sigma'(\delta) + \rho(\delta))\mathbb{B}) + \rho(\delta)\mathbb{B} \\ \subset \overline{\operatorname{co}}\, H(\hat{x}'_\delta(t) + \delta\sigma\mathbb{B}) + \delta\sigma\mathbb{B}\,,\end{aligned}$$

where $\sigma > 0$ is such that $\delta\sigma'(\delta) + \rho(\delta) \leq \sigma\delta$ for all $\delta \leq \delta^\star$. Hence, for all $\delta \leq \delta^\star$, every (T, ϵ)-truncated interpolated solution of (6.15) is a (T, ϵ)-truncated solution of

$$\dot{x} \in H(x + \delta\sigma\mathbb{B}) + \delta\sigma\mathbb{B}\,, \quad x \in \mathcal{C} + \delta\sigma\mathbb{B}$$

which is the $\delta\sigma$-perturbation of (6.14). Lemma 5.4 thus yields the desired graphical convergence as $\delta_i \to 0^+$. □

Now, we can combine Lemma 6.6 and Proposition 6.2 to arrive at our first main result about the graphical convergence of truncated solutions (i.e., local solutions) of anti-windup approximations to a projected dynamical system:

Theorem 6.1. *Let Assumptions 6.1, 6.2, and 6.3 be satisfied for the LOP discretization (6.3). Consider a sequence $\alpha_n \to 0^+$ and let $\{x_n\}$ denote the sequence of complete interpolated solutions of (6.3) with $\alpha = \alpha_n$ and $x_n(0) = x_0$. Then, then $\{x_n\}$ converges uniformly to the unique solution of (6.4) on every compact subinterval of $[0, \infty)$.*

Proof. First, note that the PDS (6.4) is not a well-posed inclusion. Instead, for Lemma 6.6 to apply, we need to consider the Krasovskii regularization $\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]$ of $\Pi_{\mathcal{C}}^{G}[f]$. However, since $\Pi_{\mathcal{C}}^{G}[f](x) \subset \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right](x)$ for all $x \in \mathcal{C}$, we can adapt (6.10) to have

$$\Sigma_{\mathcal{C}}^{G}[f](x,\alpha) \subset \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]((x+\rho(\alpha)\mathbb{B}) \cap \mathcal{C}) + \rho(\alpha)\mathbb{B}\,. \tag{6.17}$$

It then follows, that the map $H_\alpha(x) := x + \alpha\Sigma_{\mathcal{C}}^{G}[f](x,\alpha)$ satisfies (6.16), i.e.,

$$H_\alpha(x) \subset x + \alpha\left(\mathrm{K}\left[\Pi_{\mathcal{C}}^{G}[f]\right]((x+\rho(\alpha)\mathbb{B}) \cap \mathcal{C})\right) + \alpha\rho(\alpha)\mathbb{B}\,.$$

and using Lemma 6.6 we conclude uniform convergence to Krasovskii solutions of (6.4).

More precisely, every sequence of bounded solutions of (6.3) has a graphically convergent subsequence. To apply Lemma 6.6 we need to show that the LOP discretization (6.3) satisfies (6.16) with respect to the PDS (6.4). □

Theorem 6.1 does not exploit the full generality of Lemma 6.6. In particular, Assumptions 6.1, 6.2, and 6.3 guarantee the existence of unique complete solutions for both the PDS (6.4) and the LOP discretization (6.3). Hence, there is no need for more fine-grained results (for local solutions) as presented in Section 5.3 for anti-windup approximation of PDS.

6.4 Semiglobal Practical Robust Stability

Analogously to Lemma 6.6, we prove the following lemma about semiglobal practical asymptotic stability of discrete approximations starting from Lemma 5.6. Albeit the same result can be obtained from simplifying [216, Thm. 5.3].

Lemma 6.7. *Consider $\mathcal{C} \subset \mathbb{R}^n$ and $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and let the inclusion*

$$\dot{x} \in H(x)\,, \quad x \in \mathcal{C} \tag{6.18}$$

be well-posed, let $\mathcal{C}$ be a viability domain for H, and assume that $\mathcal{A} \subset \mathcal{C}$ is a compact GAS set for (6.18), i.e., there exists $\beta \in \mathcal{KL}$ such that $d_{\mathcal{A}}(x(t)) \leq \beta(d_{\mathcal{A}}(x(0)), t)$ holds for some and any (complete) solution x and all $t \geq 0$. Further, consider a class of discrete-time systems, parametrized in $\delta > 0$, and given by

$$x^+ = H_\delta(x) \quad x \in \mathcal{C}_\delta \tag{6.19}$$

with $\mathcal{C}_\delta \subset \mathbb{R}^n$ and $H_\delta : \mathcal{C}_\delta \to \mathcal{C}_\delta$. Assume that there exists $\rho \in \mathcal{K}_\infty$ and $\delta^\star > 0$ such that for all $\delta < \delta^\star$ it holds that $\mathcal{C}_\delta \subset \mathcal{C} + \rho(\delta)\mathbb{B}$ and for every $x \in \mathcal{C}_\delta$ we have

$$H_\delta(x) \subset x + \delta \operatorname{co} H((x + \rho(\delta)\mathbb{B}) \cap \mathcal{C}) + \delta\rho(\delta)\mathbb{B}\,. \tag{6.20}$$

Then, $\mathcal{A}$ is SPAS for (6.19). Namely, for every compact $\mathcal{B} \subset \mathbb{R}^n$ and every $\zeta > 0$ there exists $\delta' < \delta^\star$ such that for every $\delta < \delta'$, every solution x_δ of (6.19) starting in $\mathcal{B} \cap \mathcal{C}_\delta$ satisfies $d_{\mathcal{A}}(x_\delta[k]) \leq \beta(d_{\mathcal{A}}(x_\delta[0], \delta k) + \zeta$ for all $k \in \mathbb{N}$.

Proof. Using the same ideas as in the proof of Lemma 6.6, we know that any (T, ϵ)-truncated interpolated solution of (6.19) is a (T, ϵ)-truncated solution of the $\delta\sigma$-perturbation of (6.18) which is given by

$$\dot{x} \in H(x + \delta\sigma\mathbb{B}) + \delta\sigma\mathbb{B} \quad x \in \mathcal{C} + \delta\sigma\mathbb{B}\,.$$

Hence, Lemma 5.6 applies directly and guarantees that $d_{\mathcal{A}}(\hat{x}_\delta(t)) \leq \beta(d_{\mathcal{A}}(\hat{x}_\delta(0), t) + \zeta$ for all $t \geq 0$, and in particular, for all $t = \delta k$ with $k \in \mathbb{N}$. □

Theorem 6.2. *Let Assumptions 6.1, 6.2, and 6.3 be satisfied for the LOP discretization (6.3). Further, let $\mathcal{A} \subset \mathcal{C}$ be compact and globally asymptotically stable for the PDS (6.4), i.e., there is $\beta \in \mathcal{KL}$ such that for every solution x it holds that*

$$d_{\mathcal{A}}(x(t)) \leq \beta(d_{\mathcal{A}}(x(0)), t) \quad \forall t \geq 0\,.$$

Then, for every $\zeta > 0$ and every compact $\mathcal{B} \subset \mathcal{C}$ there exists $\alpha^\star > 0$ such that for all $\alpha \in (0, \alpha^\star)$ every solution x_α of the LOP (6.3) with $x_\alpha(0) \in \mathcal{B}$ satisfies

$$d_{\mathcal{A}}(x_\alpha[k]) \leq \beta(d_{\mathcal{A}}(x_\alpha[0]), \alpha k) + \zeta \qquad \forall k \geq 0\,.$$

Proof. The proof is analogous to the proof of Theorem 6.1. Namely, note that the PDS (6.4) is not well-posed. However, the Krasovskii regularization is well-posed and $\mathcal{C}$ is a viability domain. Moreover, since $\mathcal{C}$ is assumed to be prox-regular (hence Clarke regular), and f and G are continuous, every Krasovskii solution is a (Carathéodory) solution of (6.4) by Corollary 4.4. Importantly, this means that if $\mathcal{A}$ is gas for Carathéodory solutions then $\mathcal{A}$ is gas for Krasovskii solutions. Furthermore, (6.17) holds, as before, and from Lemma 6.7 we conclude practical stability of $\mathcal{A}$ for the parametrized class of LOP discretizations (6.15). □

6.5 Notes & Comments

The idea of linearizing constraints around a current iterate and then computing a (descent) direction by solving a quadratic program is well-established and, in fact, the key feature of SQP methods in numerical optimization [250, 191, 31]

With the LOP discretization presented in this chapter, we can recover this kind of SQP methods by choosing the vector field f to be the negative gradient of a cost function. However, our method has a more general scope because it can approximate any kind of projected dynamics.

Moreover, the results presented in this chapter seem to be the first rigorous study of the continuous-time limit of this kind of algorithm and

bring together three key components: sensitivity analysis for programs, uniform convergence arguments (generally used to prove existence of solutions for ODEs and differential inclusions), and oblique projected dynamical system.

Part II

Feedback-Based Optimization: Convergence & Robustness

CHAPTER 7

Preliminaries on Gradient Flows

In this second part of the thesis, we turn to projected dynamical systems for solving constrained optimization problems. In particular, we explore the possibilities of implementing (or approximating) projected gradient flows as the feedback control system's closed-loop behavior. For this purpose, we study algorithms that are based on projected gradient flows of the form

$$\dot{x} \in \Pi_{\mathcal{C}}^{G} \left[-\operatorname{grad}_G \Phi \right](x) \,, \tag{7.1}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is closed and $G : \mathcal{C} \to \mathbb{S}_+^n$ is a locally weakly bounded metric. Further, $\Phi : \mathbb{R}^n \to \mathbb{R}$ is an objective function, continuously differentiable in a neighborhood of $\mathcal{C}$, and $\operatorname{grad}_G \Phi(x) := G^{-1}(x) \nabla \Phi(x)^T$ is its gradient at $x \in \mathcal{C}$ with respect to G.[1]

The results of Chapter 4 can be used to guarantee the existence and uniqueness of (Carathéodory or Krasovskii) solutions of (7.1) under appropriate conditions on $\mathcal{C}, G$ and Φ. In fact, (7.1) is well-defined even on subsets of abstract C^1-manifolds.

Dynamics of the form (7.1) serve to find local optimizers to the optimization problem

$$\text{minimize } \Phi(x) \quad \text{subject to } x \in \mathcal{C} \,. \tag{7.2}$$

[1] This definition of a gradient is inspired by Riemannian geometry where $\operatorname{grad}_G \Phi(x)$ is the unique vector that satisfies $\langle \operatorname{grad}_G \Phi(x), w \rangle_{G(x)} = D_x \Phi(w)$ for all $w \in T_x \mathbb{R}^n$, where $D_x \Phi$ denotes the differential of Φ at x and $T_x \mathbb{R}^n$ is the tangent space of $\mathbb{R}^n$ at x (which is isomorphic to $\mathbb{R}^n$).

Note that the metric G is a property of the dynamical system (7.1) only and does not affect the optimizers of (7.2). Furthermore, it is reasonable (but important to note) that, in general, the metric that defines the gradient has to be the same metric required for the projection.

In the remainder of this chapter, we discuss two topics that are informative for the behavior of projected gradient flows and allow us to make connections to existing work: First, in Section 7.1 we review basic results for unconstrained gradient flows. In particular, we discuss pathologies that can occur and how to avoid them. However, our discussion of this topic is very limited since we will not pursue it further in the context of projected gradient flows. Second, in Section 7.2, we consider convex constrained optimization problems and show how particular discretizations of projected gradient flows recover special cases of well-established "proximal" optimization algorithms.

The remaining chapters of this part of the thesis are structured as follows: In Chapter 8 we establish stability and convergence properties of the basic, yet general, projected gradient flow (7.1). In Chapters 9 and 10 we then study the convergence features of anti-windup approximations and LOP discretizations of (7.1). Finally, in Chapters 11 and 12 discuss timescale separation requirements and tracking performance for time-varying optimization problems. Both of these subjects are of particular importance in feedback-based optimization.

7.1 Unconstrained Gradient Flows

We quickly review unconstrained gradient flows in $\mathbb{R}^n$ in the Euclidean metric $G \equiv \mathbb{I}$ to illustrate the basic convergence properties and possible pathologies, such as non-pointwise convergence or unstable minimizers.

Let $\Phi : \mathbb{R}^n \to \mathbb{R}$ be continuously differentiable with compact sublevel sets, i.e., for every $\ell \in \mathbb{R}$ let $\mathcal{S}_\ell := \{x \,|\, \Phi(x) \leq \ell\}$ be compact. Then,

$$\dot{x} = -\nabla\Phi(x)^T \tag{7.3}$$

admits a complete solution for every initial condition $x(0) \in \mathbb{R}^n$ and the set of weak equilibria coincides with the critical points of Φ, i.e., the

points for which $\nabla\Phi(x) = 0$.

This fact can be seen as follows: By continuity of $\nabla\Phi$, (7.3) admits a (local) solution for every initial condition. Furthermore, we have

$$\tfrac{d}{dt}\Phi(x(t)) = -\nabla\Phi(x(t))\nabla\Phi(x(t))^T \leq 0\,.$$

It follows that Φ is non-increasing along trajectories of (7.3) and therefore the sublevel sets $\mathcal{S}_\ell$ are invariant. By compactness of $\mathcal{S}_\ell$ we conclude that trajectories of (7.3) cannot escape in finite time and therefore must be complete. Moreover, from Theorem 2.5 we know that the limit set of every trajectory lies on a level set $\{x \,|\, \Phi(x) = r\}$ of Φ for some r intersected with the set of critical points $\{x \,|\, \nabla\Phi(x) = 0\}$.

This reasoning about the limit behavior of gradient flows is easily generalized to more general settings, e.g., on Riemannian manifolds [126, Chap. C.12].

Although solutions of (7.3) converge the set of critical points of Φ,

(i) a trajectory of (7.3) may not converge to a singleton, i.e., a single critical point of Φ, and

(ii) minimizers of Φ may not be stable.

For instance, in [194], the authors give an example of a gradient flow with trajectories that are not pointwise convergent. Namely, the ω-limit set of any trajectory is the entire unit circle. The same example is also documented in [72]. Moreover, the possibility of non-pointwise convergence can also occur in discrete-time gradient methods [1].

In [2], the authors show, using a counterexample, that minimizers of Φ are not necessarily stable and strict local minimizers are not necessarily asymptotically stable.

Both of these counterexamples are highly pathological. However, they can both be ruled out if the so-called *Lojasciewicz inequality* holds (see below).

7.1.1 Pointwise Convergence

Pointwise Convergence under Convexity/Monotonicity

Suppose that Φ is strictly convex (and has compact sublevel sets) or Φ is strongly convex (which implies that Φ has compact sublevel sets). In that case, there exists a unique critical point $x^\star$, which is also the unique global minimizer of Φ. Hence, convergence of any trajectory of (7.3) is necessarily to $x^\star$.

If, on the other hand, we assume that Φ is convex with compact sublevel sets, the set of minimizers is not necessarily a singleton (but it is connected and convex). Nevertheless, we can conclude pointwise convergence.

Proposition 7.1. *Let $\Phi : \mathbb{R}^n \to \mathbb{R}^n$ be convex with compact sublevel sets. Then, every trajectory of (7.3) with $G \equiv \mathbb{I}$ converges to a point that is a global minimizer of Φ.*

Proof. For the sake of contradiction, let $\hat{x}_1, \hat{x}_2 \in \Omega(x_0)$ be two different limit points of the trajectory starting at x_0. Since all trajectories of (7.3) converge to the largest invariant subset of critical points of Φ and since every critical point of a convex function is a global minimizer, it follows that $\hat{x}_1, \hat{x}_2$ are both global minimizers of Φ.

Next, we can define $V_{\hat{x}_i}(x) := \frac{1}{2}\|x - \hat{x}_i\|^2$ for $i \in \{1, 2\}$ and note that

$$\tfrac{d}{dt} V_{x_i^\star}(x) = \nabla\Phi(x(t))(x(t) - \hat{x}_i) \leq 0\,.$$

In particular, we have that $\|x(t') - \hat{x}_i\| \leq \|x(t) - \hat{x}_i\|$ for all $t' \geq t$ and $i \in \{1, 2\}$, i.e., the distance to either point is non-increasing.

Since $\hat{x}_1$ is a limit point, there exist a monotonically increasing sequence $\{t_k^1\}$ with $t_k^1 \to \infty$ and such that $x(t_k^1) \to \hat{x}_1$.

Let $\epsilon < \frac{1}{2}\|\hat{x}_1 - \hat{x}_2\|$. There exists $K \in \mathbb{N}$ such that $\|x(t_k^1) - \hat{x}_1\| \leq \epsilon$ for all $k \geq K$ and consequently $\|x(t) - \hat{x}_1\| \leq \epsilon$ for all $t \geq t_K$. This, however, implies that there cannot exist a sequence $\{t_k^2\}$ with $t_k^2 \to \infty$ and such that $x(t_k^2) \to \hat{x}_2$. Therefore, $\hat{x}_2$ cannot be a limit point of x. □

The arguments in the proof of (7.1) have been applied, implicitly or explicitly, to prove pointwise convergence of monotone dynamics such as saddle-point flows [234, 58, 59, 117].

Whether the same line of reasoning can be extended to gradient flows with a variable metric remains open.

Pointwise Convergence under Real Analycity

If Φ is not convex, we can still prove pointwise convergence by showing that trajectories of the gradient flow (7.3) have finite length.[2] This is particularly true for real analytic functions for which the *Lojasiewicz inequality* holds. Namely, if Φ is real analytic and $\hat{x}$ is a critical point of Φ, then there exist $\theta \in [\frac{1}{2}, 1)$ and $C > 0$ such that

$$|\Phi(x) - \Phi(\hat{x})|^\theta \leq C\|\nabla\Phi(x)\| \tag{7.4}$$

for all x in neighborhood $\mathcal{N}$ of $\hat{x}$.

Next, let $x : [0, \infty) \to \mathbb{R}^n$ be a trajectory of the gradient flow (7.3) such that $x(t) \in \mathcal{N}$ for all $t \geq 0$. Then, we have

$$\begin{aligned} -\tfrac{C}{1-\theta}\tfrac{d}{dt}\left(\Phi(x(t))^{1-\theta}\right) &= -C\left(\Phi(x(t))^{-\theta}\right)\nabla\Phi(x(t))\dot{x}(t) \\ &= C\left(\Phi(x(t))^{-\theta}\right)\|\nabla\Phi(x(t))\|^2 \\ &\geq \|\nabla\Phi(x(t))\| = \|\dot{x}(t)\|\,. \end{aligned}$$

Hence, the length of the trajectory x can be bounded as

$$\int_0^\infty \|\dot{x}(t)\|dt \leq \tfrac{C}{1-\theta}\Phi(x(0))^{1-\theta} < \infty$$

and it follows that the ω-limit set of x is the singleton $\hat{x}$.

The Lojasciewicz inequality (7.4) is the starting point for the field of *tame optimization* [136, 45] which also extends to non-smooth objects.

We will not pursue this avenue further for projected gradient flows. Instead, note that non-pointwise convergence of gradient flows is a possibility (at least theoretically), and powerful tools exist to rule out this kind of pathology.

[2] The content of this subsection is inspired by [113, 72].

7.1.2 Stability of Minimizers

From the viewpoint of numerical optimization, minimizers of an optimization problem can be distinguished from other critical points by applying second-order optimality conditions such as the SSOSC (Definition 2.3).

Because we want to implement projected gradient flows as the closed-loop behavior of a physical control system, we are interested in identifying minimizers by their stability properties. In other words, if we can show that the only (asymptotically) stable equilibria are minimizers, we do not have worry about converging to critical points that are not minimizers. However, the relation between stability and optimality is not trivial as the following result from [2] shows:

Theorem 7.1. *For a critical point of* (7.3) *the following relations hold:*

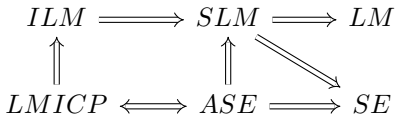

where (S)LM stands for (strict) local minimizer, (A)SE for (asymptotically) stable equilibrium, ILM for isolated local minimizer, and LMICP for local minimizer & and isolated critical point.

If, in addition, Φ is real analytic (or, more precisely, the Lojasiewicz inequality (7.4) *holds), then*

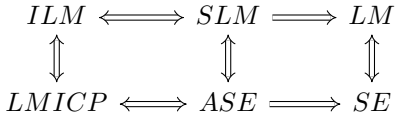

In particular, a local minimizer of Φ is not generally a stable equilibrium, and a strict local minimizer is not necessarily asymptotically stable. The following example, simplified from [2] illustrates this last point.

Example 7.1. Consider the function $h : \mathbb{R} \to \mathbb{R}$ defined as $\Phi(x) := e^{-1/x^2}(x^2 + 1 + \sin\frac{1}{x^2})$ for all $x \in \mathbb{R} \setminus \{0\}$ and $\Phi(0) = 0$. In particular, it can be shown that Φ is continuously differentiable with $\frac{d}{dt}\Phi(0) = 0$.

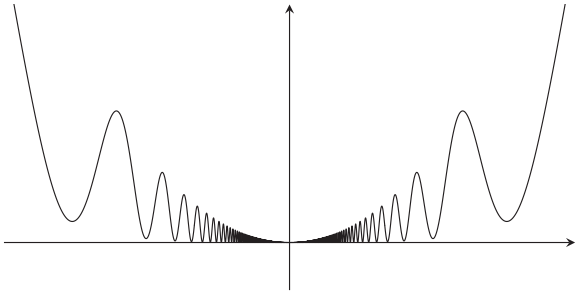

Figure 7.1: Qualitative behavior of Φ in Example 7.1

Figure 7.1 illustrates Φ qualitatively. Notice that $x = 0$ is a strict (global) minimizer of h, however, it is not isolated. Hence, in every neighborhood of 0 there exist infinitely many (local) minimizers, each of which is a stable equilibrium of the gradient flow $\dot{x} = -\nabla\Phi(x)^T$. Consequently, $x = 0$ cannot be asymptotically stable. ■

7.2 Related Work: Discretizations of Projected Gradient Flows

It is interesting to note that different discretizations of projected gradient flows result in special cases of well-established proximal optimization algorithms [23, 196, 21] such as the *proximal-point algorithm* [21, Ch. 27.1], *proximal gradient descent* [23, 25], *proximal Newton methods* [160], and *mirror descent* [150, 24, 188].

Our discussion of this topic is fairly superficial. It merely serves to illustrate how these optimization algorithms have continuous-time limits that can be captured by a single yet general model of oblique projected gradient flows.

For detailed results and theoretical guarantees (e.g., on the convergence rate) for iterative optimization algorithms, the reader is referred to the respective papers.

All of the algorithms considered in this section apply primarily to

convex optimization problems. Hence, we limit ourselves to the problem

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad x \in \mathcal{C}\,, \tag{7.5}$$

where $\mathcal{C}$ is closed convex and $\Phi : \mathbb{R}^n \to \mathbb{R}$ is a proper closed and convex (i.e., Φ has a non-empty, closed convex epigraph $\operatorname{epi}\Phi := \{(x, y) \,|\, y \geq \Phi(x)\}$).

Furthermore, a recurring object in this section in the proximal operator prox_f defined by

$$\operatorname{prox}_f(x) := \arg\min_y \left(f(y) + \tfrac{1}{2}\|x - y\|^2\right)$$

for a proper closed convex function $\mathbb{R}^n \to \mathbb{R}$. For a comprehensive treatment of the proximal mappings and related algorithms, the reader is referred to [196, 211].

7.2.1 Implicit Euler Discretization & Proximal Minimization

Arguably the simplest proximal optimization algorithm is the *proximal-point algorithm* ([196, Ch. 4.1] or [21, Ch. 27.1]) which is given by

$$x^+ = \operatorname{prox}_\Phi(x)\,, \tag{7.6}$$

where $\Phi : \mathbb{R}^n \to \mathbb{R}$ is a convex function with non-empty closed epigraph. If Φ has a minimizer, then any solution of (7.6) converges to the set of minimizers of Φ.

In [196] it was noted that if Φ is continuously differentiable (7.6) can be interpreted as an implicit Euler integration of the unconstrained gradient flow $\dot{x} = -\nabla\Phi(x)^T$.

It turns out that this analogy extends to projected gradient flows: Define the convex function $\Phi_\mathcal{C} := \Phi + I_\mathcal{C}$ where $I_\mathcal{C} : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ denotes the indicator function of $\mathcal{C}$ (i.e., $I_\mathcal{C}(x) = 0$ for all $x \in \mathcal{C}$ and $I_\mathcal{C} = \infty$ otherwise). Hence, $\Phi_\mathcal{C}$ is a proper closed convex function with non-empty closed epigraph. We can reformulate the definition of $\operatorname{prox}_{\Phi_\mathcal{C}}$

as

$$\begin{aligned}\operatorname{prox}_{\Phi_{\mathcal{C}}}(x) &= \arg\min_{y\in\mathbb{R}^n} \left\{\Phi(y) + I_{\mathcal{C}}(y) + \tfrac{1}{2\alpha}\|y-x\|^2\right\} \\ &= \arg\min_{y\in\mathcal{C}} \left\{\Phi(y) + \tfrac{1}{2\alpha}\|y-x\|^2\right\} .\end{aligned}$$

Note that $y^\star := \operatorname{prox}_{\Phi_{\mathcal{C}}}(x)$ satisfies the KKT conditions

$$-\nabla\Phi(y^\star)^T - \tfrac{1}{\alpha}(y^\star - x) \in N_{y^\star}\mathcal{C}$$

which can be rearranged to yield

$$(x - \alpha\nabla\Phi(y^\star)^T) - y^\star \in N_{y^\star}\mathcal{C} . \tag{7.7}$$

However, (7.7) is also satisfied by the (unique) solution of

$$\text{minimize } \|(x-\alpha\nabla\Phi(y)^T) - y\|^2 \qquad \text{subject to} \quad y\in\mathcal{C}$$

and therefore satisfies $y^\star = P_{\mathcal{C}}\left(x - \alpha\nabla\Phi(y^\star)^T\right)$. Hence, the iteration $x^+ = \operatorname{prox}_{\Phi_{\mathcal{C}}}(x)$ can rewritten implicitly as

$$x^+ = P_{\mathcal{C}}\left(x - \alpha\nabla\Phi(x^+)^T\right)$$

which, in turn, can be interpreted as implicit Euler integration of the (Euclidean) projected gradient flow

$$\dot{x} = \Pi_{\mathcal{C}}\left[-\nabla\Phi^T\right](x) \qquad x\in\mathcal{C} .$$

Note that, for general PDS an implicit Euler integration is usually not practical as it requires solving the implicit relation $x^+ = P_{\mathcal{C}}(x - \alpha f(x^+))$ where $f : \mathbb{R}^n \to \mathbb{R}^n$ is a vector field. The proximal-point algorithm marks a special case in which solving this system amounts to solving a (simple) convex optimization problem.

7.2.2 Forward Euler Discretizations

Projected/Proximal Gradient Method

As second proximal optimization algorithm we consider the *proximal gradient method* ([23, Ch. 10] or [196, Ch. 4.2]) which, in our case, reduces to a simple *projected gradient descent*.

Proximal gradient descent is primarily suited for minimizing a composite objective $\Psi := \Phi + \Lambda$ where $\Phi, \Lambda : \mathbb{R}^n \to [-\infty, \infty]$ are proper closed and convex, and, in addition, Φ is continuously differentiable. Then, the proximal gradient method takes the form

$$x^+ = \operatorname{prox}_\Lambda \left(x - \alpha_k \nabla \Phi(x)^T \right) , \tag{7.8}$$

where α_k denotes the step-size (chosen to be fixed or determined, e.g., by backtracking line search). If $\Lambda := \lambda \| \cdot \|_1$ with $\lambda > 0$ is an ℓ_1-regularization, (7.8) is also known as *iterative shrinkage-thresholding algorithm* (ISTA) [25].

If we choose $\Lambda := I_{\mathcal{C}}$ as the indicator function of the non-empty closed convex set $\mathcal{C}$, then the $\operatorname{prox}_{I_{\mathcal{C}}}$ reduces to a simple Euclidean projection onto $\mathcal{C}$ and we recover a simple projected gradient descent

$$x^+ = \operatorname{prox}_{I_{\mathcal{C}}} \left(x - \alpha_k \nabla \Phi(x)^T \right) = P_{\mathcal{C}} \left(x - \alpha \nabla \Phi(x)^T \right)$$

which a simple *forward Euler* integration (see also Remark 6.1) of the projected gradient flow

$$\dot{x} = \Pi_{\mathcal{C}} \left[-\nabla \Phi^T \right] (x) \qquad x \in \mathcal{C} \, .$$

This connection is not surprising and has been documented previously in [23, 186] and others.

Newton Gradient Flows & Proximal Newton Methods

More interesting than the Euclidean projected gradient descent/flow are algorithms that use the possibility of a variable metric. In this context, *proximal Newton methods* [160] have been developed to achieve superior convergence rates of the Newton method for the minimization of a composite objective function with a non-smooth cost component.

To describe this Newton-type method, we first define the *scaled proximal operator* as

$$\operatorname{prox}_f^Q(x) := \arg\min_y \left(f(y) + \tfrac{1}{2} \|x - y\|_Q^2 \right) ,$$

where $Q \in \mathbb{S}^n_+$ is a fixed positive definite matrix. Then, a simplified version (with fixed step-size) of the proximal Newton method for minimizing $\Phi + \Lambda$ takes the form

$$x^+ = \operatorname{prox}_{\Lambda}^{G(x)} \left(-\alpha G(x)^{-1} \nabla \Phi(x)^T \right) ,$$

where $G(x) = \nabla^2 \Phi(x)$ is the Hessian of Φ at x (assuming that Φ is twice continuously differentiable).

As before, we choose $\Lambda = I_{\mathcal{C}}$ to be the indicator function of $\mathcal{C}$ and hence we have

$$\begin{aligned} \operatorname{prox}_{I_{\mathcal{C}}}^{G(x)} &= \arg\min_{y \in \mathcal{C}} \left\{ \tfrac{1}{2} \|x - \alpha G(x)^{-1} \nabla \Phi(x)^T - y\|^2_{G(x)} \right\} \\ &= P_{\mathcal{C}}^{G(x)} \left(x - \alpha \operatorname{grad}_G(x) \right) , \end{aligned}$$

where $P_{\mathcal{C}}^{Q}$ denotes the projection onto $\mathcal{C}$ with respect to the metric induced by Q, i.e.,

$$P_{\mathcal{C}}^{Q}(x) := \arg\min_{y \in \mathcal{C}} \|x - y\|_Q \, .$$

Lemma 7.1. *Given a closed convex set $\mathcal{C} \subset \mathbb{R}^n$, a metric $G : \mathcal{C} \to \mathbb{S}^n_+$, and a vector $v \in \mathbb{R}^n$ we have, for every $x \in \mathcal{C}$ that*

$$\lim_{\delta \to 0^+} \frac{P_{\mathcal{C}}^{G(x)}(x + \delta v) - x}{\delta} = \Pi_{\mathcal{C}}^{G} \left[v \right] (x) \, .$$

Proof. The statement follows directly from the analogous result for the Euclidean case, i.e., $\lim_{\delta \to 0^+} \frac{P_{\mathcal{C}}(x+\delta v) - x}{\delta} = \Pi_{\mathcal{C}} \left[v \right] (x)$ [130, Prop. III.5.3.5], after applying the linear transformation $\hat{x} = G(x)^{1/2} x$. □

Hence, the proximal Newton method with $\Lambda = I_{\mathcal{C}}$ can be seen as a forward Euler discretization of the continuous-time projected Newton flow

$$\dot{x} = \Pi_{\mathcal{C}}^{G} \left[- \operatorname{grad}_G \Phi \right] (x) \qquad x \in \mathcal{C} \, .$$

Note that, to make this connection rigorous, having a proper definition of *oblique* PDS as introduced in [Ha2] (i.e, Chapter 4) is essential.

Further note that, in practical applications, proximal Newton methods are often implemented as "inexact" methods where only an approximation of $G(x) = \nabla^2 \Phi(x)$ and $G(x)^{-1}$ is evaluated at every iteration.

7.2.3 Mirror Descent

Mirror Descent is an optimization algorithm for convex optimization going back to [188] where it was described as a descent algorithm in which gradient steps are performed in the dual space of $\mathbb{R}^n$ and a correspondence between $\mathbb{R}^n$ and its dual is established via a so-called *mirror map*.

Later, in [23, 24], a new interpretation of mirror descent as a generalized projected subgradient method was given. In this context, the Bregman divergence of the mirror map serves as a non-Euclidean norm. This statement will be made more precise below.

In the following, we conjecture that mirror descent is a discretization of a projected gradient flow in the metric defined by the Hessian of the mirror map. Although our analysis is not rigorous, it leads us to an intriguing subclass of projected gradient flows, which defies some of our previous characterizations in Chapter 4.

Mirror descent requires a so-called *mirror map* $\psi : \mathbb{R}^n \to [0, \infty]$ with a convex domain $\mathcal{X} := \{x \,|\, \psi(x) < \infty\}$ and $\psi(\partial\mathcal{X}) = \infty$, i.e., ψ reaches ∞ at the boundary of the domain. Furthermore, ψ is β-strongly convex and continuously differentiable on $\operatorname{int}\mathcal{X}$. Hence, mirror descent is defined by the iteration

$$x^+ = \arg\min_{y \in \mathcal{C}} \left\{ \Phi(x) + \nabla\Phi(x)(y - x) + \tfrac{1}{\alpha} B_\psi(y, x) \right\} , \tag{7.9}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is closed convex and B_ψ denotes the *Bregman divergence* of ψ given by

$$B_\psi(x, z) := \psi(x) - \psi(z) - \langle \nabla\psi(z), x - z \rangle . \tag{7.10}$$

For an appropriate step-size rule, (7.9) is guaranteed to converge to a minimizer of the convex optimization problem

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad \mathcal{C} \cap \mathcal{X} . \tag{7.11}$$

If ψ is twice continuously differentiable and $\|x - z\|$ is small, we have that

$$B_\psi(x, z) \approx (x - z)^T \nabla^2\psi(z)(x - z)$$

which we can use to approximate (7.9) as

$$\begin{aligned} x^+ &= \arg\min_{y\in\mathcal{C}} \left\{\Phi(x) + \nabla\Phi(x)(y-x) + \tfrac{1}{\alpha}B_\psi(y,x)\right\} \\ &\approx \arg\min_{y\in\mathcal{C}} \left\{\left\langle \alpha(\nabla^2\psi(x))^{-1}\nabla\Phi(x)^T, y-x\right\rangle_{\nabla^2\psi(x)} + \|y-x\|^2_{\nabla^2\psi(x)}\right\} \\ &= \arg\min_{y\in\mathcal{C}} \left\|y - x + \tfrac{\alpha}{2}\operatorname{grad}_{\nabla^2\psi}\Phi(x)\right\|^2_{\nabla^2\psi(x)} \\ &= \arg\min_{y\in\mathcal{C}} \left\|y - \left(x - \tfrac{\alpha}{2}\operatorname{grad}_{\nabla^2\psi}\Phi(x)\right)\right\|^2_{\nabla^2\psi(x)} \\ &= P_{\mathcal{C}}^{\nabla^2\psi}\left(x - \tfrac{\alpha}{2}\operatorname{grad}_{\nabla^2\psi}\Phi(x)\right) . \end{aligned}$$

Using Lemma 7.1 we therefore conjecture that the continuous-time limit of (7.9) as $\alpha \to 0^+$ is given by the projected gradient flow

$$\dot{x} = \Pi_{\mathcal{C}}^{\nabla^2\psi}\left[-\operatorname{grad}_{\nabla^2\psi}\Phi\right](x) \qquad x \in \mathcal{C}\cap\operatorname{int}\mathcal{X}\,. \tag{7.12}$$

The system (7.12) is itself very interesting, even without considering a potential connection to mirror descent. Namely, because ψ is a mirror map, $\nabla^2\psi$ is well defined only on $\operatorname{int}\mathcal{X}$, the domain of $\mathcal{C}$. For this reason, (7.12) is well-defined only on $\mathcal{C}\cap\operatorname{int}\mathcal{X}$.

While Corollaries 4.1 and 4.3 guarantee the existence of local solutions of (7.12) for any initial condition in $\mathcal{C}\cap\operatorname{int}\mathcal{X}$, the existence of complete solutions cannot be easily guaranteed, since a solution may reach the boundary of $\mathcal{X}$ in finite time.

Moreover, because (7.12) is a projected gradient flow it can be shown (in the forthcoming Chapter 8) that Φ is non-increasing along trajectories of (7.12). This fact is, however, not enough to apply the invariance principle Theorem 2.5 which requires the existence of complete solutions. Therefore, one cannot easily conclude that trajectories of (7.12) converge to critical points of (7.11), although for unconstrained gradient flows this appears to be the case [8].

CHAPTER 8

Convergence of Projected Gradient Flows

Recall that for a differential inclusion $\dot{x} \in F(x)$, $\hat{x}$ is a weak equilibrium if $x(t) = \hat{x}$ for all $t \geq 0$ is a solution. Hence, $\hat{x}$ is a weak equilibrium if and only if $0 \in F(\hat{x})$. The point $\hat{x}$ is a *strong equilibrium* if $x(t) = \hat{x}$ for all $t \geq 0$ is the only solution starting at $\hat{x}$.

When considering the projected gradient flow (7.1) we need to distinguish between equilibrium points for Carathódory and Krasovskii solutions. In particular, we say that $x^\star$ is a *weak (strong) K-equilibrium*, if it is a weak (strong) equilibrium of the Krasovskii-regularized inclusion. Analogously, $x^\star$ is a *weak (strong) C-equilibrium* if it is an equilibrium for Carathéodory solutions (i.e., solutions of the unregularized inclusion).

Since every Carathéodory solution of (7.1) is also a Krasovskii solution, it follows that every strong K-equilibrium is also a strong C-equilibrium. On the other hand, a weak C-equilibrium is a weak K-equilibrium. Furthermore, recall from Definition 2.8 that a *critical point* of the optimization problem (7.2) is a point $x^\star \in \mathcal{C}$ satisfying

$$\nabla\Phi(x^\star)w \geq 0 \qquad \text{for all} \qquad w \in T_{x^\star}\mathcal{C}\,. \tag{8.1}$$

The results in this chapter generalize previous statements from [Ha15] and [Ha2, Sec. 5]. The main technical results are Proposition 8.1 and Theorem 8.1.

Lemma 8.1. *Every critical point of (7.2) and, in particular, every minimizer of (7.2), is a weak K-equilibrium of (7.1) and every weak C-equilibrium of (7.1) is a critical point of (7.2).*

Proof. We can reformulate (8.1) as $\langle -\operatorname{grad}_G \Phi(x^\star), w\rangle_{G(x^\star)} \leq 0$ for all $w \in T_{x^\star}\mathcal{C}$. Furthermore, by Lemma 4.3 we have, for all $x \in \mathcal{C}$,

$$\langle -\operatorname{grad}_G \Phi(x), w\rangle_{G(x)} \geq \|w\|^2_{G(x)} \quad \forall w \in \mathrm{K}\left[\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]\right](x)\,.$$

Combining these two statements we have

$$0 \geq \langle -\operatorname{grad}_G \Phi(x^\star), w\rangle_{G(x^\star)} \geq \|w\|^2_{G(x^\star)} \qquad \forall w \in T_{x^\star}\mathcal{C} \cap \mathrm{K}\left[\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]\right](x^\star)\,.$$

We know that $T_x\mathcal{C} \cap \mathrm{K}\left[\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]\right](x) \neq \emptyset$ for all $x \in \mathcal{C}$ since $\mathcal{C}$ is a viability domain for $\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]$. Therefore, we conclude that $w = 0 \in \mathrm{K}\left[\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]\right](x^\star)$.

Consider a weak C-equilibrium $x^\star \in \mathcal{C}$, i.e., $0 \in \Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right](x^\star)$. If $x^\star$ is not a critical point of (7.2), then there exists $v \in T_{x^\star}\mathcal{C}$ such that $\langle -\operatorname{grad}_G \Phi(x), v\rangle > 0$. This means that $0 \notin \Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right](x^\star)$. To see this, note that the projection of u onto v is given by $w := \frac{\langle u,v\rangle}{\|v\|^2} v$. Using the triangle inequality to the right triangle $\{0, u, w\}$, we have $\|u - w\| < \|u - 0\|$. Hence, 0 cannot be a projection of u onto $T_{x^\star}\mathcal{C}$ since it does not achieve the minimal the distance to $T_{x^\star}\mathcal{C}$. □

Lemma 8.2. *Along Krasovskii solutions of (7.1), Φ is nonincreasing. Moreover, the sublevel sets of Φ on $\mathcal{C}$, i.e., $S_\ell := \{x \mid \Phi(x) \leq \ell\} \cap \mathcal{C}$ for $\ell \in \mathbb{R}$, are invariant for Krasovskii solutions.*

Proof. Given any Krasovskii solution $x : [0, T] \to \mathcal{C}$ of (7.1), we have, for almost all $t \in [0, T]$

$$\tfrac{d}{dt}\Phi(x(t)) = D_{x(t)}\Phi(w) = \langle \operatorname{grad}_G \Phi(x), w(t)\rangle_{G(x(t))}$$

for some $w(t) \in \mathrm{K}\left[\Pi^G_{\mathcal{C}}\left[-\operatorname{grad}_G \Phi\right]\right](x(t))$. By Lemma 4.3 on Page 84, we then have

$$\tfrac{d}{dt}\Phi(x(t)) = D_{x(t)}\Phi(w) = -\langle -\operatorname{grad}_G \Phi(x(t)), w\rangle_{G(x(t))} \leq -\|w(t)\|^2_{G(x(t))} \leq 0\,. \quad (8.2)$$

Thus Φ is non-increasing along Krasovskii solutions of (7.1) and hence $\mathcal{S}_\ell$ is invariant. □

Lemma 8.3. *Minimizers of* (7.2) *are strong K-equilibria of* (7.1).

Proof. By Corollary 4.1 there exists a Krasovskii solution $x : [0, T] \to \mathcal{C}$ of (7.1) starting at the minimizer $x^\star \in \mathcal{C}$ of (7.2). Assume for the sake of contradiction that $x(0) = x^\star$ but $x(T) \neq x^\star$.

The sublevel set $\mathcal{S}_{\ell^\star}$ with $\ell^\star = \Phi(x^\star)$ is invariant and $x(t) \in \mathcal{S}_{\ell^\star}$ for all $t \in [0, T]$ by Lemma 8.2. Since $x^\star \in \mathcal{C}$ is a minimizer there exists a neighborhood $\mathcal{N} \subset \mathcal{C}$ of $x^\star$ such that $\Phi(x') \geq \Phi(x^\star)$ for all $x' \in \mathcal{N}$. If necessary, restrict the solution x such that $x : [0, T] \to \mathcal{N}$. It follows that $\Phi(x(t)) = \Phi(x^\star)$ for all $t \in [0, T]$ and therefore $\frac{d}{dt}\Phi(x(t)) = 0$ for almost all $t \in [0, T]$ which implies that

$$0 = \tfrac{d}{dt}\Phi(x(t)) = -\left\langle -\operatorname{grad}_G \Phi(x(t)), \dot{x}(t) \right\rangle_{G(x)} \leq -\|\dot{x}(t)\|\,,$$

where the inequality follows from Lemma 4.3. Consequently, we have $\dot{x}(t) = 0$ for almost all $t \in [0, T]$ and thus $x(t) = \int_0^t \dot{x}(t)dt = x_0$ which establishes the contradiction. □

Hence, Lemmas 8.1 and 8.3 together with Theorem 4.1 can be summarized in the following statements:

Proposition 8.1. *Consider the projected gradient flow* (7.1) *and the problem* (7.2). *If $\mathcal{C}$ is closed, G is locally weakly bounded and Φ is continuously differentiable in a neighborhood of $\mathcal{C}$, then the following inclusions hold:*

$$\begin{aligned} \textit{minimizer} \subset \textit{strong K-eq.} \subset \textit{strong C-eq.} \\ \subset \textit{weak C-eq.} \subset \textit{crit. pt.} \subset \textit{weak K-eq.} \end{aligned}$$

If, in addition, $\mathcal{C}$ is Clarke regular and G is continuous, then we have

$$\textit{minimizer} \subset \textit{strong eq.} \subset \textit{weak eq.} = \textit{crit. pt.}$$

As an example of a critical point that is a weak (C-)equilibrium, but not a strong (C-)equilibrium, we refer back to Example 4.2, which illustrates this case for a prox-regular set that induces non-unique solutions.

If solutions of (7.1) are unique, we do not have to distinguish between weak and strong equilibria, and Proposition 8.1 simplifies further to the fact that critical points are equivalent to equilibria.

Unfortunately, convergence is generally guaranteed only to weak K-equilibria as the following application of the invariance principle Theorem 2.5 shows.

Proposition 8.2. *Consider (7.1) and let $\Phi : \mathbb{R}^n \to \mathbb{R}$ have compact sublevel sets on $\mathcal{C}$, i.e., for every $\ell \in \mathbb{R}$ the set $\mathcal{S}_\ell := \{x \,|\, \Phi(x) \leq \ell\} \cap \mathcal{C}$ is compact. Then, (7.1) admits a complete Krasovskii solution $x : [0, \infty) \to \mathcal{C}$ for every initial condition $x(0) \in \mathcal{C}$. Furthermore, for some $r \in \Phi(\mathcal{S}_\ell)$, x converges to the largest weakly invariant subset of weak K-equilibrium points on $\Phi^{-1}(r) \cap \mathcal{C}$. If, in addition, $\mathcal{C}$ is Clarke regular and G is continuous, then convergence is to the set of critical points of (7.2).*

Proof. As before, the compactness and invariance of the level sets $\mathcal{S}_\ell$ of Φ on $\mathcal{C}$ implies that solutions cannot escape in finite time and therefore must be complete. Hence, Theorem 2.5 guarantees convergence to the largest weakly invariant subset for which $\frac{d}{dt}\Phi(x(t)) = 0$ (and which lies on a level set of Φ relative to $\mathcal{C}$) Using (8.2), we know that every limit point $\hat{x}$ of x satisfies $0 \in \mathrm{K}\left[\Pi_{\mathcal{C}}^{G}\left[-\operatorname{grad}_G \Phi\right]\right](\hat{x})$, i.e., $\hat{x}$ is a weak K-equilibrium of (7.1). Finally, under Clarke regularity of $\mathcal{C}$ and continuity of G, Proposition 8.1 implies that every weak equilibrium is a critical point. □

Although convergence is, in general, only to the set of weak equilibria, the following theorem establishes stability of minimizers analogously to Theorem 7.1, albeit for projected gradient flows and sets of minimizers (instead of individual points). Namely, in the following theorem $\hat{\mathcal{X}} \subset \mathcal{C}$ is a *(strict) minimizer of* (7.2) if $\Phi(x) = \Phi(x')$ for all $x, x' \in \hat{\mathcal{X}}$ and $\Phi(y) \geq (>)\Phi(x)$ for all $y \notin \hat{\mathcal{X}}$ (but $y \in \mathcal{C}$).

Theorem 8.1. *(stability of minimizers) Consider (7.1) and let Φ have compact sublevel sets on $\mathcal{C}$ as in Proposition 8.2. Let $\hat{\mathcal{X}} \subset \mathcal{C}$ be a connected component of the set of weak K-equilibria and a subset of $\{x \,|\, \Phi(x) = r\}$ for some r. Then, the following statements hold:*

(i) If $\hat{\mathcal{X}}$ is asymptotically stable for (7.1), then it is a strict minimizer of (7.2).

(ii) If $\hat{\mathcal{X}}$ is a strict minimizer of (7.2) then it is stable for (7.1).

Proof. Recall from Proposition 8.2 that the compactness of the sublevel sets guarantees the existence of complete solutions.

To show (i), let $\mathcal{V} \subset \mathcal{C}$ be a neighborhood of $\hat{\mathcal{X}}$ such any solution $x : [0, \infty) \to \mathcal{C}$ of (7.1) with $x(0) \in \mathcal{V}$ converges to $\hat{\mathcal{X}}$. Since Φ is C^1 and x is absolutely continuous, $\Phi \circ x$ is absolutely continuous, and we may write

$$\lim_{t \to +\infty} (\Phi \circ x)(t) = \Phi(x_0) + \int_0^{+\infty} D_x \Phi(\dot{x}(t)) dt = r \,.$$

Since $D_x \Phi(\dot{x}(t)) \leq 0$ holds for almost all $t \geq 0$, we can establish $\int_0^{+\infty} D_x \Phi(\dot{x}(t)) dt \leq 0$ and hence $r \leq \Phi(x(t)) \leq \Phi(x_0)$ for all $t \geq 0$. Because this reasoning applies to all x_0 in the region of attraction of $\hat{\mathcal{X}}$, it follows that $\hat{\mathcal{X}}$ is a local minimizer of Φ.

To see that $\hat{\mathcal{X}}$ is a strict minimizer, assume for the sake of contradiction that for some $\widetilde{x}$ in the region of attraction of $\hat{\mathcal{X}}$ it holds that $\Phi(\widetilde{x}) \leq r$. Every solution y to (7.1) with $y(0) = \widetilde{x}$ nevertheless converges to $\hat{\mathcal{X}}$ by assumption. Therefore, it must hold that $\int_0^{+\infty} D_y \Phi(\dot{y}(t)) = 0$ and since $D_y \Phi(\dot{y}(t)) \leq 0$, it follows that $D_y \Phi(\dot{y}(t)) = 0$ for almost all $t \geq 0$. But as a consequence of Proposition 8.2, all points x in the limit set are weak K-equilibrium points, this holds in particular for $\widetilde{x}$, and therefore $\hat{\mathcal{X}}$ cannot be asymptotically stable. For (ii), assume that $\hat{\mathcal{X}} \neq \mathcal{C}$ (otherwise stability is trivial). Hence, consider a bounded (relative) neighborhood $\mathcal{W} \subset \mathcal{C}$ of $\hat{\mathcal{X}}$ in which $\hat{\mathcal{X}}$ is a strict minimizer. Next, we construct a neighborhood $\mathcal{V} \subset \mathcal{W}$ such that all trajectories starting in $\mathcal{V}$ remain in $\mathcal{W}$. Namely, let α be such that $r < \alpha < \min_{x \in \partial \mathcal{W}} \Phi(x)$ where $\partial \mathcal{W}$ is the boundary of $\mathcal{W}$ relative to $\mathcal{C}$. Define $\mathcal{V} := \{x \in \mathcal{W} \,|\, \Phi(x) \leq \alpha\} \subseteq \mathcal{W}$ which has a non-empty interior because $r < \alpha$. Since for any trajectory, we have $D_x \Phi(\dot{x}(\tau)) \leq 0$ we conclude that $\mathcal{V}$ is strongly invariant and consequently remains in $\mathcal{V}$, thus establishing stability. □

CHAPTER 9

Robust Convergence of Anti-Windup Approximations

In this chapter we revisit the anti-windup approximations introduced in Chapter 5. Namely, we consider the special case of (5.1) when f depends only on $P_{\mathcal{C}}(x)$, i.e., we study the system

$$\dot{x} \in F_K(x) := f(P_{\mathcal{C}}(x)) - \tfrac{1}{K} G^{-1}(P_{\mathcal{C}}(x))(x - P_{\mathcal{C}}(x)) \,, \tag{9.1}$$

where, as before, $\mathcal{C} \subset \mathbb{R}^n$ is α-prox-regular, $G : \mathcal{C} \to \mathbb{S}^n_+$ is a continuous metric, $K > 0$ is a scalar, and $f : \mathcal{C} \to \mathbb{R}^n$ is a continuous vector field. All of the previous results from Chapter 5 for (5.1) also apply to (9.1). In particular, as $K \to 0^+$, trajectories of (9.1) converge uniformly to solutions of the PDS (5.2), i.e., to solutions of

$$\dot{x} = \Pi^G_{\mathcal{C}}\left[f\right](x) \,, \qquad x \in \mathcal{C} \,. \tag{9.2}$$

Also, the practical stability results of Section 5.4 apply, but we show in this chapter that stronger results can be derived for (9.1). In particular, we prove three results:

(i) We show in Section 9.1 that equilibria of (9.2) coincide with *projected* equilibria of (9.1), i.e., equilibrium points after projection

The material presented in this chapter corresponds to the second part of [Ha4] and Section 9.3 corresponds to [Ha5]. Theorems 9.1, 9.2, and 9.3 constitute the main technical results of this chapter.

onto $\mathcal{C}$.

(ii) In Section 9.2 we prove that for strongly monotone f, convex $\mathcal{C}$, $G \equiv \mathbb{I}$, and K small enough, trajectories of (9.1) converge to the unique projected equilibrium point of (9.2).

(iii) When f is a gradient vector field, $\mathcal{C}$ is prox-regular, and $G \equiv \mathbb{I}$, we prove in Section 9.3 that (projected) trajectories of (9.1) converge to the set of minimizers of the potential function on $\mathcal{C}$, i.e., solution of the underlying optimization problem.

Finally, in Sections 9.4 and 9.5 we give four control designs that fall into the class of (9.1). We provide simulations confirming our theoretical results and formulate open questions.

9.1 Preservation of Equilibria

An important advantage of (9.1) over the more general system (5.1) is that equilibria of (9.2) are preserved in the following sense (which generalizes [Ha6, Prop. 4]):

Proposition 9.1. *If $\overline{x}^\star \in \mathcal{C}$ is a weak equilibrium of the PDS (9.2), then there exists $K^\star > 0$ such that for all $K \in (0, K^\star)$ there exists a weak equilibrium point $x_K^\star \in \overline{x}^\star + N_{\overline{x}^\star}\mathcal{C} \cap \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ for the AWA (9.1). Conversely, if $x_K^\star \in \mathcal{C}_\alpha^\circ$ is a weak equilibrium of (9.1) for some K, then $P_{\mathcal{C}}(x^\star)$ is a weak equilibrium of (9.2).*

Proof. Given a weak equilibrium $\overline{x}^\star \in \mathcal{C}$ of (9.2), let $x_K^\star := \overline{x}^\star - KG(\overline{x}^\star)f(\overline{x}^\star)$. For $K \in (0, K^\star) := 1/(2\alpha\|G(\overline{x}^\star)f(\overline{x}^\star)\|)$, we have $x_K^\star \in \mathcal{C}_\alpha^\circ$.

Since $\overline{x}^\star$ is an equilibrium of (9.2) (by assumption) and by Lemma 4.1, we have $f(\overline{x}^\star) \in -N_{\overline{x}^\star}^G \mathcal{C}$. It follows from (3.2) that $-KG(\overline{x}^\star)f(\overline{x}^\star) \in N_{\overline{x}^\star}\mathcal{C}$ and consequently $x_K^\star \in \overline{x}^\star + N_{\overline{x}^\star}\mathcal{C}$. By Proposition 3.1, it follows that $P_{\mathcal{C}}(x_K^\star) = \overline{x}^\star$ and therefore

$$F_K(x_K^\star) = f(\overline{x}^\star) - \tfrac{1}{K} G^{-1}(\overline{x}^\star)\,(\overline{x}^\star - KG(\overline{x}^\star)f(\overline{x}^\star) - \overline{x}^\star) = 0\,.$$

Thus, $x_K^\star$ is a weak equilibrium of (9.1). The converse case follows the same ideas. □

Although equilibria of the PDS (9.2) are preserved by the AWA (9.1) (after projection), it is not clear whether convergence properties are preserved, especially since we are primarily interested in the convergence of $t \mapsto P_{\mathcal{C}}(x(t))$ rather than the convergence of the solution x of (9.1). Theorem 5.2 suggests that, in general, convergence is only within a neighborhood of asymptotically stable equilibria of the PDS (9.2).

However, as we shown below, under additional conditions on f, G and $\mathcal{C}$, the projected solutions $t \mapsto P_{\mathcal{C}}(x(t))$ do indeed converge to an equilibrium of (9.2).

9.2 AWA of Monotone Dynamics

Next, we show that if $-f$ is strongly monotone and $G \equiv \mathbb{I}$, then $-F_K$, as defined in (9.1), is strictly monotone for small enough K. This, in turn, allows us to conclude asymptotic stability for anti-windup approximations of the form (9.1).

Since we require only monotonicity of $-f$, the following results can be used not only when f is chosen as the gradient of a convex cost function, but also for saddle-point flows (see Section 9.5), and pseudo-gradients for Nash-equilibrium seeking [73, 186].

Given a set $\mathcal{C} \subset \mathbb{R}^n$, recall that a map $F : \mathcal{C} \rightrightarrows \mathbb{R}^n$ is *(strictly; β-strongly) monotone* if for all $x, x' \in \mathcal{C}$ and all $v \in F(x)$ and $v' \in F(x')$ it holds that

$$\langle v - v', x - x' \rangle \geq 0 \, (> 0; \geq \beta \|x - x'\|^2) \, .$$

Further, if $\mathcal{C}$ is α-prox-regular, the map $x \mapsto N_x\mathcal{C}$ has a *hypomonotone localization* (Lemma 3.4), i.e., for all $x, x' \in \mathcal{C}$, all $\eta \in N_x\mathcal{C} \cap \mathbb{B}$, and all $\eta' \in N_{x'}\mathcal{C} \cap \mathbb{B}$ we have

$$\langle \eta - \eta', x - x' \rangle \geq -2\alpha \|x - x'\|^2 \, .$$

In particular, if $\mathcal{C}$ is convex, we have $\langle \eta' - \eta, x' - x \rangle \geq 0$ and $x \mapsto N_x\mathcal{C}$ is monotone.

Proposition 9.2. *Consider F_K as defined in (9.1) with $G \equiv \mathbb{I}$ and $\mathcal{C}$ is assumed to be α-prox-regular. Let $-f$ be β-strongly monotone and globally ℓ-Lipschitz. Then $-F_K$ is strictly monotone on $\mathcal{C}_\alpha^\circ$ for all $0 < K < 4(\beta - 2\alpha)/\ell^2$.*

Proof. Given any $x, x' \in \mathcal{C}_\alpha^\circ$, let $\overline{x} := P_{\mathcal{C}}(x)$ and $\overline{x}' := P_{\mathcal{C}}(x')$. Further, let $\eta := x - \overline{x} \in N_{\overline{x}}\mathcal{C}$ and $\eta' := x' - \overline{x}' \in N_{\overline{x}'}\mathcal{C}$. We can work directly with the monotonicity of f, the hypomonotocity of $x \mapsto N_x\mathcal{C}$, and Cauchy-Schwarz to derive

$$
\begin{aligned}
&\langle x - x', F_K(x) - F_K(x') \rangle \\
&\quad = \left\langle x - x', f(\overline{x}) - f(\overline{x}') - \tfrac{1}{K}(x - \overline{x}) + \tfrac{1}{K}(x' - \overline{x}') \right\rangle \\
&\quad = \left\langle \overline{x} - \overline{x}' + \eta - \eta', f(\overline{x}) - f(\overline{x}') - \tfrac{1}{K}(\eta - \eta') \right\rangle \\
&\quad = \langle \overline{x} - \overline{x}', f(\overline{x}) - f(\overline{x}') \rangle - \tfrac{1}{K} \langle \eta - \eta', \eta - \eta' \rangle \\
&\qquad\qquad + \underbrace{\langle \eta - \eta', f(\overline{x}) - f(\overline{x}') \rangle}_{\ell \|\eta - \eta'\| \|\overline{x} - \overline{x}'\|} - \tfrac{1}{K} \underbrace{\langle \overline{x} - \overline{x}', \eta - \eta' \rangle}_{\geq -2\alpha \|\overline{x} - \overline{x}'\|^2} \\
&\quad \leq -(\beta - 2\alpha)\|\overline{x} - \overline{x}'\|^2 + \ell \|\overline{x} - \overline{x}'\| \|\eta - \eta'\| - \tfrac{1}{K}\|\eta - \eta'\|^2 .
\end{aligned}
$$

A sufficient condition for the right-hand side to be negative for all $\overline{x} \neq \overline{x}'$ is that $\beta - 2\alpha > 0$ and that the determinant $\frac{1}{K}(\beta - 2\alpha) - \frac{1}{4}\ell^2$ is positive, i.e., if $0 < K < 4(\beta - 2\alpha)/\ell^2$. □

This leads us the following result which establishes convergence of anti-windup approximations for strongly monotone dynamics on convex sets:

Theorem 9.1. *Consider the AWA (9.1) with $G \equiv \mathbb{I}$ and let $\mathcal{C}$ be closed convex. Assume that $-f$ is β-strongly monotone and globally ℓ-Lipschitz. Then, for all $K < 4\beta/\ell^2$, every trajectory of (9.1) converges to an equilibrium point $x^\star$ (which is unique) such that $P_{\mathcal{C}}(x^\star)$ is the unique equilibrium of the PDS (9.2).*

Proof. Because of convexity of $\mathcal{C}$, $P_{\mathcal{C}}(x)$ is single-valued and continuous for all $x \in \mathbb{R}^n$ and globally 1-Lipschitz (i.e., non-expansive). As a

consequence, F_K is globally Lipschitz continuous and there exists a unique complete solution of (9.1) for every initial condition $x(0) \in \mathbb{R}^n$. Furthermore, since $K < 4\beta/\ell^2$ and $\mathcal{C}$ is convex (which lets us take $\alpha \to 0^+$), Proposition 9.2 guarantees that $-F_K$ is strictly monotone on $\mathbb{R}^n$.

Next, recall that the strong monotonicity of $-f$ and convexity of $\mathcal{C}$ imply that (9.2) has a unique equilibrium $\overline{x}^\star$ [186, Thm. 2.3]. Consequently, Proposition 9.1 guarantees the existence of an equilibrium point $x^\star$ of (9.1) such that $P_{\mathcal{C}}(x^\star) = \overline{x}^\star$. Furthermore, $x^\star$ is unique by [186, Thm. 2.2]. In particular, strict monotonicity of $-F_K$ implies that $V(x) := \frac{1}{2}\|x - x^\star\|^2$ is a Lyapunov function for (9.1) which can be used to establish global asymptotic stability of $x^\star$. □

Theorem 9.1 can, presumably, be generalized to prox-regular sets as well as general metrics G. However, in that case, additional restrictions on $x(0)$ are required, the threshold value for K is less easily quantifiable, and convergence is likely only local.

9.3 Robust Convergence of Anti-Windup Gradient Flows

Although Theorem 9.1 establishes robust convergence for monotone dynamics with a quantitative bound on the anti-windup gain K, we show next that, in the special case of an AWA of a projected gradient flow, robust convergence is guaranteed under significantly weaker assumptions. More precisely, we consider the optimization problem

$$\text{minimize} \quad \Phi(x) \quad \text{subject to} \quad x \in \mathcal{C} \tag{9.3}$$

for which we make the following assumption:

Assumption 9.1. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular. Further, let $\Phi : \mathbb{R}^n \to \mathbb{R}$ be differentiable in a neighborhood of $\mathcal{C}$ with compact sublevel sets $\mathcal{S}_\ell := \{x \in \mathcal{C} \,|\, \Phi(x) \leq \ell\}$.*

Under Assumption 9.1, $x^\star \in \mathcal{C}$ is a *critical point of* (9.3) (i.e., *1st-order optimal*) if $\nabla\Phi(x^\star)^T \in -N_{x^\star}\mathcal{C}$. Namely, local optimizers of (9.3) are critical [211, Thm. 6.12].

To solve (9.3), we consider the *anti-windup gradient inclusion*

$$\dot{x} \in F(x) := -\nabla\Phi(P_{\mathcal{C}}(x))^T - \tfrac{1}{K}\left(x - P_{\mathcal{C}}(x)\right) \tag{9.4}$$

which falls into the class of (9.1). In particular, we consider only the case $G \equiv \mathbb{I}$. Since $P_{\mathcal{C}}$ is not necessarily single-valued outside $\mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$, (9.4) has to be treated as an inclusion. However, we will not concern ourselves with potential solutions outside of $\mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$. Instead, we define the sets of admissible initial conditions (which we later show to be invariant) as

$$\mathcal{C}_\ell := \left\{x \in \mathcal{C} + \tfrac{1}{2\alpha}\operatorname{int}\mathbb{B} \,|\, P_{\mathcal{C}}(x) \in \mathcal{S}_\ell\right\}$$

which is the pre-image of $\mathcal{S}_\ell$ restricted to $\mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$.

Remark 9.1. Note that the system (9.4) bears similarities with the gradient flow

$$\dot{x} = -\nabla\tilde{\Phi}(x)^T = -\nabla\Phi(x)^T - \tfrac{1}{K}\left(x - P_{\mathcal{C}}(x)\right), \tag{9.5}$$

where $\tilde{\Phi}(x) := \Phi(x) + \frac{1}{2K}d_{\mathcal{C}}^2(x)$ and $d_{\mathcal{C}}$ denotes the point-to-set distance to $\mathcal{C}$. Namely, $\tilde{\Phi}$ is a cost function augmented with a term penalizing the distance from the feasible set $\mathcal{C}$. However, the inconspicuous difference between (9.4) and (9.5) in the argument of $\nabla\Phi$ leads to two important contrasts: First, if $x^\star$ is an equilibrium of (9.4), then $P_{\mathcal{C}}(x^\star)$ is a optimizer of (9.3) (see Proposition 9.1). This is not the case for equilibria of (9.5); equilibria of (9.5) are minimizers of $\tilde{\Phi}$ but not necessarily optimizers of (9.3). Second, convergence to the set of global minimizers of $\tilde{\Phi}$ can be easily established for (9.5). However, proving convergence of solutions of (9.4) to optimizers of (9.3) is more challenging and the main topic of this section. ■

To study (9.4) it is convenient to think in terms of *projected trajectories*, i.e., the trajectory obtained from the (pointwise) projection of a solution of (9.4) as illustrated in Section 9.3. Although these projected

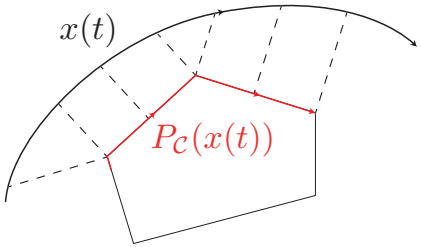

Figure 9.1: Construction of a projected trajectory

trajectories are not the solution of a dynamical system in the form of a differential inclusion, their limit behavior can be studied.

Our main result in this section guarantees that there always exists $K > 0$ such that the projected trajectories of the anti-windup gradient flow (9.4) converge to the critical points of (9.3), although, K may depend on the choice on ℓ and thereby on the set of initial conditions.

Theorem 9.2. *Under Assumption 9.1 and given $\ell \in \mathbb{R}$, there exists $K^\star > 0$ such that (9.4) admits a complete solution $x : [0, \infty) \to \mathcal{C}_\ell$ for all $K \in (0, K^\star)$ and all initial conditions $x(0) \in \mathcal{C}_\ell$. Further, for any such solution, the projected trajectory $\overline{x} := P_{\mathcal{C}} \circ x$ converges to the set of critical points of (9.3).*

Theorem 9.2 also applies to convex $\mathcal{C}$ since convex sets are α-prox-regular for any $\alpha > 0$. However, we can derive a stronger result without restrictions on the initial condition or on K.

Theorem 9.3. *If Assumption 9.1 holds and $\mathcal{C}$ is closed convex, (9.4) admits a complete solution x for all $K > 0$ and all $x(0) \in \mathbb{R}^n$. Further, for any such solution, the projected trajectory $\overline{x} := P_{\mathcal{C}} \circ x$ converges to the set of critical points of (9.3).*

The proofs of Theorems 9.2 and 9.3 are relatively involved, but ultimately an application of the invariance principle Theorem 2.5 based on the fact that $x \mapsto \Phi(P_{\mathcal{C}}(x))$ is non-increasing along trajectories of (9.4). To make this argument rigorous, we proceed as follows: First, we prove several properties of the directional derivative of $P_{\mathcal{C}}$ in Section 9.3.1. Then, in

Section 9.3.2 we show that the projected trajectories $x \mapsto P_{\mathcal{C}}(x(t))$ of (9.4) are well-defined and absolutely continuous. In Section 9.3.3, we establish convergence to the largest invariant set for which $\frac{d}{dt}P_{\mathcal{C}}(x(t)) = 0$ and in Section 9.3.4 we show that this largest invariant subset is equivalent to the set of critical points of (9.3), thus proving Theorem 9.3. Finally, in Section 9.3.5 we show the modifications that are required to prove Theorem 9.3.

9.3.1 Directional Derivatives of Projection Maps

Given a closed set $\mathcal{C} \subset \mathbb{R}^n$, we establish properties of the directional derivative of the projection $P_{\mathcal{C}}$.

The differentiability of $P_{\mathcal{C}}$ has been studied extensively, albeit—to the best of the author's knowledge—only for convex sets $\mathcal{C}$. Even if $\mathcal{C}$ is convex, $P_{\mathcal{C}}$ is in general not differentiable unless $\mathcal{C}$ has a smooth boundary [102]. Further, $P_{\mathcal{C}}$ is not generally directionally differentiable [220, 153] unless *second-order regularity* assumptions on $\mathcal{C}$ are satisfied [221, 47]. An up-to-date review of this subject including detailed examples can also be found in [254]. We avoid these technicalities because we require directional differentiability only along a trajectory (c.f. Lemma 9.8).

Hence, recall that the directional derivative of $P_{\mathcal{C}}$ at $x \in \mathbb{R}^n$ in direction $v \in \mathbb{R}^n$ is defined as

$$DP_{\mathcal{C}}(x; v) := \lim_{\delta \to 0^+} \frac{P_{\mathcal{C}}(x + \delta v) - P_{\mathcal{C}}(x)}{\delta}. \tag{9.6}$$

The classical result [130, Prop. III.5.3.5] states that for convex $\mathcal{C}$, $DP_{\mathcal{C}}(x; v)$ exists for all $x \in \mathcal{C}$ and all $v \in \mathbb{R}^n$ and is given as the projection of v onto the tangent cone at x. Its generalization to α-prox-regular sets is straightforward.

Lemma 9.1. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular for some $\alpha > 0$. Then, $DP_{\mathcal{C}}(x; v)$ exists for all $x \in \mathcal{C}$ and all $v \in \mathbb{R}^n$ and is given by*

$$DP_{\mathcal{C}}(x; v) = \Pi_{\mathcal{C}}[v](x) = \lim_{\delta \to 0^+} \frac{P_{\mathcal{C}}(x + \delta v) - x}{\delta}.$$

Characterizing $DP_{\mathcal{C}}(x;v)$ at $x \notin \mathcal{C}$ is harder and directional differentiability is in general not guaranteed (see [220, 153]). However, the forthcoming Lemma 9.7 guarantees that, along an absolutely continuous trajectory, the directional derivative of $P_{\mathcal{C}}$ exists for almost all t.

Assuming that $DP_{\mathcal{C}}(x;v)$ exists, one can establish various properties. First of all, it immediately follows from the definition of the tangent cone that $DP_{\mathcal{C}}(x;v)$ is *viable*:

Lemma 9.2. *If $\mathcal{C} \subset \mathbb{R}^n$ is α-prox-regular, $x \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$, $v \in \mathbb{R}^n$, and if $DP_{\mathcal{C}}(x;v)$ exists, then $DP_{\mathcal{C}}(x;v) \in T_{P_{\mathcal{C}}(x)}\mathcal{C}$.*

The next two lemmas exploit basic properties of $P_{\mathcal{C}}$.

Lemma 9.3. *Consider an α-prox-regular set $\mathcal{C} \subset \mathbb{R}^n$ and let $x \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$ and $v \in \mathbb{R}^n$ be such that $\overline{v} := DP_{\mathcal{C}}(x;v)$ exists. Then, we have $\langle v, \overline{v}\rangle \geq 0$.*

Proof. Recall that for a closed set $\mathcal{C}$, the projection $P_{\mathcal{C}}$ is monotone [211, Cor. 12.20]. It follows that this property also holds in the limit by continuity of $P_{\mathcal{C}}$ (Lemma 3.3), i.e.,

$$\langle \overline{v}, v\rangle = \lim_{h\to 0^+} \tfrac{\langle P_{\mathcal{C}}(x+hv)-P_{\mathcal{C}}(x),(x+hv)-x\rangle}{h^2} \geq 0\,. \qquad \square$$

Lemma 9.4. *Let $\mathcal{C} \subset \mathbb{R}^n$ be α-prox-regular, $x \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$, $v \in \mathbb{R}^n$, and assume that $\overline{v} := DP_{\mathcal{C}}(x;v)$ exists. Then, it holds that*

$$\langle \overline{v}, x - P_{\mathcal{C}}(x)\rangle = 0\,.$$

Proof. Define the map $\phi(h) := x+hv$ for all $h \geq 0$. Using Proposition 3.1 and the chain rule, we know that

$$\nabla\left(d^2_{\mathcal{C}} \circ \phi\right)\Big|_{h=0} = 2\,\langle v, x - P_{\mathcal{C}}(x)\rangle\ .$$

On the other hand, we can apply the chain rule to $d^2_{\mathcal{C}}(\phi(h)) = \|\phi(h) - P_{\mathcal{C}}(\phi(h))\|^2$ to arrive at

$$\nabla\left(\|\phi(h) - P_{\mathcal{C}}(\phi(h))\|^2\right)\Big|_{h=0} = 2\,\langle v - \overline{v}, x - P_{\mathcal{C}}(x)\rangle\ .$$

The difference of the expressions yields the result. $\square$

Lemmas 9.2 and 9.4 yield that $DP_{\mathcal{C}}(x;v)$, if it exists, lies in $\mathcal{K}(x) := T_x\mathcal{C} \cap \{v \mid \langle v, x - \overline{x}\rangle = 0\}$ which is known as the *critical cone at* x. This observation is in agreement with [221] and generalizes this insight from convex to prox-regular sets.

For the next crucial lemma we exploit the hypomonotone localization of $x \mapsto N_x\mathcal{C}$ according to Lemma 3.4.

Lemma 9.5. *Let* $\mathcal{C} \subset \mathbb{R}^n$ *be* α*-prox-regular,* $x \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$*,* $v \in \mathbb{R}^n$*, and assume that* $\overline{v} := DP_{\mathcal{C}}(x;v)$ *exists. Then,*

$$\langle v, \overline{v}\rangle = 0 \quad \Longleftrightarrow \quad \overline{v} = 0\,.$$

Proof. ($\Leftarrow$) is trivial. For ($\Rightarrow$), consider $h > 0$ such that $x_h := x + hv \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$. Further, let $\overline{x}_h = P_{\mathcal{C}}(x_h)$ and $\overline{x} := P_{\mathcal{C}}(x)$, as well as $\eta := x - \overline{x}$ and $\eta_h = x_h - \overline{x}_h$. Recall that $\eta \in N_{\overline{x}}\mathcal{C}$ and $\eta_h \in N_{\overline{x}_h}\mathcal{C}$ (Lemma 3.2). Using these facts and the definition of $\overline{v} = DP_{\mathcal{C}}(x;v)$ in (9.6) we can write

$$\begin{aligned}
\langle v, \overline{v}\rangle &= \lim_{h\to 0^+} \tfrac{1}{h^2}\,\langle x_h - x, \overline{x}_h - \overline{x}\rangle \\
&= \lim_{h\to 0^+} \tfrac{1}{h^2}\,\langle \overline{x}_h + \eta_h - \overline{x} + \eta, \overline{x}_h - \overline{x}\rangle \\
&= \|\overline{v}\|^2 + \lim_{h\to 0^+} \tfrac{1}{h}\,\langle \eta_h - \eta, \overline{x}_h - \overline{x}\rangle\,.
\end{aligned}$$

Since, by assumption, $x \in \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$, there exists $\epsilon > 0$ such that $x, x_h \in \mathcal{C} + \frac{1}{2\alpha+\epsilon}$ for small enough h. Therefore, $\|\eta\|$ and $\|\eta_h\|$ are both upper bounded by $\frac{1}{2\alpha+\epsilon}$.

To apply Lemma 3.4 we rescale $\hat{\eta} := (2\alpha+\epsilon)\eta$ and $\hat{\eta}_h := (2\alpha+\epsilon)\eta_h$ which satisfy $\hat{\eta}, \hat{\eta}_h \in \mathbb{B}$. It follows that

$$\begin{aligned}
\langle v, \overline{v}\rangle &= \|\overline{v}\|^2 + \lim_{h\to 0^+} \frac{1}{(2\alpha+\epsilon)h^2}\underbrace{\langle \overline{x}_h - \overline{x}, \hat{\eta}_h - \hat{\eta}\rangle}_{\geq -2\alpha\|\overline{x}_h - \overline{x}\|^2} \\
&\geq \|\overline{v}\|^2 - \lim_{h\to 0^+} \frac{2\alpha}{(2\alpha+\epsilon)h^2}\|\overline{x}_h - \overline{x}\|^2 \\
&= \left(1 - \tfrac{2\alpha}{2\alpha+\epsilon}\right)\|\overline{v}\|^2\,.
\end{aligned}$$

Since $\epsilon > 0$, we have $\frac{2\alpha}{2\alpha+\epsilon} < 1$ and thus $\langle v, \overline{v}\rangle = 0$ implies that $\overline{v} = 0$ which completes the proof. □

If $\mathcal{C}$ is closed convex, Lemmas 9.2-9.5 simplify to the following facts (see also [221, 102], and others):

Lemma 9.6. *Let $\mathcal{C} \subset \mathbb{R}^n$ be closed convex and let $x \in \mathbb{R}^n$ and $v \in \mathbb{R}^n$ be such that $\overline{v} := DP_{\mathcal{C}}(x; v)$ exists. Then,*

(i) $\overline{v} \in T_{P_{\mathcal{C}}(x)}\mathcal{C} \cap \{w \mid \langle w, x - P_{\mathcal{C}}(x)\rangle = 0\}$,

(ii) $\langle v, \overline{v}\rangle \geq 0$, *and*

(iii) $\langle v, \overline{v}\rangle = 0 \quad \Longleftrightarrow \quad \overline{v} = 0$.

9.3.2 Projected Trajectories

As mentioned before, establishing directional differentiability of $P_{\mathcal{C}}$, i.e., the existence of $DP_{\mathcal{C}}(x; v)$ for all $x \in \mathbb{R}^n$ and all directions $v \in \mathbb{R}^n$ is a major challenge and in general not possible without additional assumptions on $\mathcal{C}$. For our purposes, we do not require directional differentiability of $P_{\mathcal{C}}$ *everywhere* and in *all* directions because we consider only projected trajectories that come with a priori guarantees on the existence of their time derivative.

Lemma 9.7. *Consider an α-prox-regular set $\mathcal{C} \subset \mathbb{R}^n$ and an absolutely continuous map $x : [0, T] \to \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ for some $T > 0$. Then, $\overline{x} := P_{\mathcal{C}} \circ x$ is single-valued and absolutely continuous. Furthermore, $\dot{\overline{x}}(t)$ and $\dot{x}(t)$ exist and satisfy $\dot{\overline{x}}(t) = DP_{\mathcal{C}}(x(t); \dot{x}(t))$ for almost all $[0, T]$.*

Proof. Since $\mathcal{C}$ is α-prox-regular, Lemma 3.3 guarantees that $P_{\mathcal{C}}$ is Lipschitz in every closed neighborhood of $\mathcal{C}$ that is a subset of $\mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ (in particular $x([0, T])$ is compact by continuity of x). Since the composition of a Lipschitz map and an absolutely continuous function is absolutely continuous [212, Ex. 6.44], it follows that $\overline{x}$ is absolutely continuous and hence differentiable almost everywhere.

Since x and $\overline{x}$ are differentiable everywhere except on zero measure sets $\Xi_x, \Xi_{\overline{x}} \subset [0, T]$, respectively, it follows that $\dot{x}(t)$ and $\dot{\overline{x}}(t)$ both exist except on the zero-measure set $\Xi_x \cup \Xi_{\overline{x}}$ and $\dot{\overline{x}}(t) = DP_{\mathcal{C}}(x(t); \dot{x}(t))$ holds by definition of the time derivative of $\overline{x}$. □

Remark 9.2. The existence of $\dot{\overline{x}}(t)$ is in general independent of the existence of $\dot{x}(t)$. On one hand, even if $\dot{x}(t)$ exists, $\dot{\overline{x}}(t)$ might not exist because of a lack of directional differentiability. On the other hand, $\dot{\overline{x}}(t)$ might exist, even though $\dot{x}(t)$ does not. This can occur, for instance, if $\mathcal{C} = \{0\}$ in which case $\overline{x} \equiv 0$ is trivially differentiable everywhere. ■

9.3.3 Convergence to Invariant Set

We use Theorem 2.5 by showing that $t \mapsto \Phi(P_{\mathcal{C}}(x(t)))$ is non-increasing along any solution of (9.4). Then, we prove that the limit set contains only critical points of (9.3).

Throughout this section (and the next) we use the notation $\overline{x} := P_{\mathcal{C}}(x)$ for points and $\overline{x} := P_{\mathcal{C}} \circ x$ for trajectories.

Prox-regularity of $\mathcal{C}$ and continuity of $\nabla\Phi$ guarantee the existence of solutions of (9.4) in a neighborhood of $\mathcal{C}$:

Lemma 9.8. *Under Assumption 9.1, there exists a solution of* (9.4) *for every initial condition* $x(0) \in \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$. *More precisely, there exists a differentiable function* $x : [0, T] \to \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ *for some* $T > 0$ *that satisfies for all* $t \in [0, T]$

$$\dot{x}(t) = -\nabla\Phi(P_{\mathcal{C}}(x(t)))^T + \tfrac{1}{K}\left(x(t) - P_{\mathcal{C}}(x(t))\right) .$$

Proof. From Lemma 3.3 it follows that $P_{\mathcal{C}}$ is single-valued and continuous for all $x \in \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$. Further, since $\nabla\Phi$ is continuous, F is continuous. Hence, standard results for continuous ODEs guarantee the existence of a local solution for every initial condition on the open set $\mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$. □

To show that Φ is non-increasing along projected trajectories of (9.4) we use the lemmas in Section 9.3.1. Further, to apply Theorem 2.5 we need to show that (unprojected) trajectories of (9.4) are complete and bounded, which is possible, in general, only for small enough K (unless $\mathcal{C}$ is convex).

Lemma 9.9. *Let Assumption 9.1 hold. Given a solution* $x : [0, T] \to \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ *of* (9.4) *for some* $T > 0$, *the map* $t \mapsto \Phi(\overline{x}(t))$ *is non-increasing for all* $t \in [0, T]$.

Proof. Lemmas 9.3, 9.4, and 9.7 yield, for almost all $t \in [0, T]$,

$$\begin{aligned}\tfrac{d}{dt}\Phi(\overline{x}(t)) &= \nabla\Phi(\overline{x}(t))\dot{\overline{x}}(t)\\ &= \left\langle \nabla\Phi(\overline{x}(t))^T + \tfrac{1}{K}(x(t) - \overline{x}(t)), \dot{\overline{x}}(t)\right\rangle\\ &= \left\langle -\dot{x}(t), \dot{\overline{x}}(t)\right\rangle \leq 0\,,\end{aligned}$$

and we conclude that $\Phi \circ \overline{x}$ is non-increasing. □

Lemma 9.10. *Under Assumption 9.1, $\mathcal{C}_\ell$ is bounded $\forall \ell \in \mathbb{R}$.*

Proof. The set $\mathcal{C}_\ell$ is as the pre-image of $\mathcal{S}_\ell$ under $P_{\mathcal{C}}$ restricted to $\mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$. From Lemma 3.2 it follows that for any $x \in \mathcal{C}$ we have $P_{\mathcal{C}}^{-1}(x) = x + N_x\mathcal{C} \cap \frac{1}{2\alpha}\operatorname{int}\mathbb{B} \subset x + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$. Since $\mathcal{S}_\ell$ is compact, $\mathcal{C}_\ell \subset \mathcal{S}_\ell + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$ is bounded. □

Proposition 9.3. *Let Assumption 9.1 hold. Given $\ell \in \mathbb{R}$, there exists $K^\star > 0$ such that (9.4) admits a complete solution $x : [0, \infty) \to \mathcal{C}_\ell$ for every $x(0) \in \mathcal{C}_\ell$ and for all $K \in (0, K^\star)$.*

Proof. First, note that Lemma 9.8 guarantees the existence of a (local) solution $x : [0, T] \to \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$ for any initial condition $x(0) \in \mathcal{C}_\ell \subset \mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$ and for some $T > 0$.

Since Lemma 9.9 guarantees that $\Phi(\overline{x}(t)) \leq \Phi(\overline{x}(0)) \leq \ell$ for all $t \in [0, T]$, it follows that $\overline{x}(t) \in \mathcal{S}_\ell$ for all $[0, T]$.

By compactness of $\mathcal{S}_\ell$, there exists $M > 0$, such that $\|\nabla\Phi(y)\| \leq M$ for all $y \in \mathcal{S}_\ell$. Now, consider the Lie derivative of $d_{\mathcal{C}}^2$ along (9.4). For $x \in \mathcal{C}_\ell$ we have

$$\begin{aligned}\mathcal{L}_F d_{\mathcal{C}}^2(x) &= \left\langle x - \overline{x}, -\nabla\Phi(\overline{x})^T - \tfrac{1}{K}(x - \overline{x})\right\rangle\\ &\leq d_{\mathcal{C}}(x)\|\nabla\Phi(\overline{x})\| - \tfrac{1}{K}d_{\mathcal{C}}^2(x)\\ &\leq d_{\mathcal{C}}(x)\left(M - \tfrac{1}{K}d_{\mathcal{C}}(x)\right)\,.\end{aligned}$$

It follows that $\mathcal{L}_F d_{\mathcal{C}}^2(x) < 0$ for all $x \in \mathcal{C}_\ell$ for which $d_{\mathcal{C}}(x) > KM$. In particular, if $K < K^\star := \frac{1}{2\alpha M}$, any solution x of (9.4) starting in $\mathcal{C}_\ell$ cannot leave the neighborhood $\mathcal{C} + \frac{1}{2\alpha}\operatorname{int}\mathbb{B}$ on which $P_{\mathcal{C}}$ is single-valued. In addition, $\overline{x} = P_{\mathcal{C}}(x)$ remains in $\mathcal{S}_\ell$. Hence $\mathcal{C}_\ell$ is invariant. Together with the boundedness of $\mathcal{C}_\ell$, finite-time escape is precluded and thus guaranteeing the existence of a complete solution. □

Proposition 9.4. *Under Assumption 9.1 any complete solution* $x : [0, \infty) \to \mathcal{C}_\ell$ *of* (9.4) *converges to the largest weakly invariant subset* $\mathcal{S}$ *of* $\mathcal{M} := \operatorname{cl}\{x \in \mathcal{C}_\ell \,|\, DP_{\mathcal{C}}(x; F(x)) = 0\}$.

Proof. Note that $\Phi \circ P_{\mathcal{C}}$ is continuous on $\mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$ by continuity of Φ and Lemma 3.3. Hence, to apply Theorem 2.5, let $V : \mathbb{R}^n \to \mathbb{R}$ be any continuous function such that $V(y) = \Phi(P_{\mathcal{C}}(y))$ for all $y \in \mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$. Further, let $\mathcal{U} := \mathcal{C}_\ell$. The trajectory x is complete by assumption and bounded by Lemma 9.10. Hence, according to Theorem 2.5, x converges to the largest weakly invariant subset of $V^{-1}(r) \cap \mathcal{U} \cap \operatorname{cl} u^{-1}(0)$ for some r and where we have

$$\begin{aligned} u^{-1}(0) &= \{x \in \mathcal{U} \,|\, \langle F(x), DP_{\mathcal{C}}(x; F(x)) \rangle = 0\} \\ &= \{x \in \mathcal{U} \,|\, DP_{\mathcal{C}}(x; F(x)) = 0\}, \end{aligned}$$

where the second equality follows from Lemma 9.5. □

It is important to note that $\operatorname{cl}\{x \in \mathcal{C}_\ell \,|\, DP_{\mathcal{C}}(x; F(x)) = 0\}$ is not, in general, invariant itself. There can exist compact intervals $[t_1, t_2]$ on which $\overline{x}$ is constant (and hence $\dot{\overline{x}}(t) = 0$ for all $t \in [t_1, t_2]$), but on which x is not stationary. For example, in Figure 9.2, this is the case when $\overline{x}(t) = P_{\mathcal{C}}(x(t))$ is stuck in one of the vertices of the feasible polyhedron $\mathcal{C}$, while x is evolving outside of $\mathcal{C}$, moving "around the corner".

9.3.4 Characterization of Invariant Limit Set

Next, we show that the largest weakly invariant subset $\mathcal{S}$ in Proposition 9.4 is equivalent to the critical points of (9.3).

Lemma 9.11. *Consider the same setup as in Proposition 9.4 and let* $x : [0, \infty) \to \mathcal{S}$ *be a complete solution of* (9.4) *evolving on the weakly invariant set* $\mathcal{S}$. *Then,* $\overline{x}(t) = \overline{x}(0)$ *holds.*

Proof. Since $\overline{x}$ is absolutely continuous, it follows that $\overline{x}(\tau) - \overline{x}(0) = \int_0^\tau \dot{\overline{x}}(t) dt$. However, $\dot{\overline{x}}(t) = 0$ holds for almost all $t \geq 0$ since, by invariance, $\overline{x}(t) \in \mathcal{S} \subset \mathcal{M}$ and therefore $\overline{x}(\tau) = \overline{x}(0)$. □

Proposition 9.5. *Consider the setup of Proposition 9.4. Then, every $x^\star \in P_{\mathcal{C}}(\mathcal{S})$ is a critical point of (9.3).*

Proof. Consider a trajectory $x : [0, \infty) \to \mathcal{S}$ evolving on the weakly invariant set $\mathcal{S}$. By Lemma 9.11, we have that $\overline{x}(t) = \overline{x}(0) =: y$ for all $t \geq 0$. Therefore, x evolves on the pre-image $P_{\mathcal{C}}^{-1}(y)$ which, using Lemma 3.2, is given by $y + N_y\mathcal{C}$. In other words, $x(t) \in y + N_y\mathcal{C}$ for all $t \geq 0$. In particular, x satisfies $\dot{x}(t) = -\nabla\Phi(y)^T - \frac{1}{K}(x(t) - y)$ for all $t \geq 0$. Thus, x is also the solution of an asymptotically stable linear system and converges to a point $\hat{x}$ such that $-\nabla\Phi(y)^T = \frac{1}{K}(\hat{x} - y) \in N_y\mathcal{C}$ and $P_{\mathcal{C}}(\hat{x}) = y$ hold. In other words, y is a critical point. □

Theorem 9.2 now follows directly since Proposition 9.3 yields the existence of a complete solution and Propositions 9.4 and 9.5 guarantee the convergence of $\overline{x}$ to the set of critical points.

9.3.5 Proof Sketch for the Convex Case

Theorem 9.3 does not directly derive from Theorem 9.2 by letting $\alpha \to 0^+$, because $\lim_{\alpha \to 0^+} \mathcal{C}_\ell$ is not bounded. Instead, we need to adapt Proposition 9.3 as follows:

Proposition 9.6. *Let Assumption 9.1 hold and let $\mathcal{C}$ be convex. Then, (9.4) admits a complete and bounded solution for every initial condition $x(0) \in \mathbb{R}^n$ and all $K > 0$.*

Proof. The proof is analogous to the proof of Proposition 9.3. In particular, we have $\mathcal{L}_F d_{\mathcal{C}}^2(x) < 0$ for all $x \in \mathcal{C}_\ell$ for which $d_{\mathcal{C}}(x) > KM$. However, since $P_{\mathcal{C}}$ is globally single-valued, K does not need to be chosen small enough to guarantee the invariance of a neighborhood $\mathcal{C} + \frac{1}{2\alpha} \operatorname{int} \mathbb{B}$. Instead, we have

$$x(t) \in \mathcal{C} + \gamma\mathbb{B} \qquad \forall t \in [0, T]$$

for all $t \geq 0$ with $\gamma > \max\{KM, d_{\mathcal{C}}(x(0))\}$. More precisely,

$$x(t) \in \mathcal{C}_\ell^\gamma := \{y \in \mathcal{C} + \gamma\mathbb{B} \mid P_{\mathcal{C}}(y) \in \mathcal{S}_\ell\},$$

for all $t \geq 0$ and where $\ell := \Phi(P_{\mathcal{C}}(x(0)))$. This follows from Lemma 9.9 since $t \mapsto \Phi(P_{\mathcal{C}}(x(t)))$ is non-increasing. Using the same argument as for Lemma 9.10, we can show that $\mathcal{C}_\ell^\gamma$ is bounded. □

Finally, Propositions 9.4 and 9.5 can be adapted using $\mathcal{C}_\ell^\gamma$ instead of $\mathcal{C}_\ell$ and Theorem 9.3 follows similarly to Theorem 9.2.

9.4 Feedback-Based Gradient Schemes for Quadratic Programs

To illustrate the design opportunities for feedback-based optimization, we present three anti-windup schemes that approximate projected gradient flows for a *quadratic program* (QP). We consider the relatively simple problem of solving a QP as it allows for a concise presentation, easy implementation, and comparability. However, needless to say, our theoretical results in the previous sections cover much more general setups.

Our goal is to design a feedback controller that steers a plant to a steady state that solves the optimization problem

$$\begin{aligned} \text{minimize} \quad & \Phi(x) := \tfrac{1}{2} x^T Q x + c^T x + d \\ \text{subject to} \quad & x = h(u) := Hu + w \\ & u \in \mathcal{U} := \{v \,|\, A_u v \leq b_u\}, \end{aligned} \tag{9.7}$$

where $x \in \mathbb{R}^m$ and $u \in \mathbb{R}^p$ denote the system state and control input, respectively, and $Q \in \mathbb{S}_+^m$, $A_u \in \mathbb{R}^{r \times p}$ and the remaining parameters are of appropriate size. The map h denotes the steady-state input-to-state map of the plant subject to the disturbance w.[1] The set $\mathcal{U}$ defines constraints which are enforced by physical saturation.

[1] In contrast to (5.3), we assume for (9.7) that the physical plant is described by an steady-state input-to-state map $x = h(u)$ that satisfies $\tilde{f}(h(u), u) = 0$ for all $u \in \mathcal{U}$. This approximation can be motivated by singular perturbation ideas in Chapter 11 and [Ha3, Ha13] which stipulate that the interconnection of *fast decaying* plant dynamics and *slow* optimization dynamics is asymptotically stable.

For solving (9.7) we aim at approximating the projected gradient flow $\dot{u} = \Pi_{\mathcal{U}}^{G}\left[-\operatorname{grad}_G \hat{\Phi}\right](u)$, where we have defined $\hat{\Phi}(u) := \Phi(h(u))$ to eliminate the state variable x. In particular, we have $\nabla\hat{\Phi}(u)^T = H^T\nabla\Phi(h(u))^T$. In the following, the metric G will be either $G \equiv \mathbb{I}$ or $G \equiv Q$ (the latter yielding a projected Newton flow).

To approximate $\dot{u} = \Pi_{\mathcal{U}}^{G}\left[-\operatorname{grad}_G \hat{\Phi}\right](u)$, we consider three systems that fall into the class of anti-windup approximations defined by (5.1), two of which can be implemented in a feedback loop as in Figure 5.1. Their convergence behavior for the same problem instance and varying K is illustrated in Figure 9.2 and discussed below.

i) Penalty Gradient Flow: As a reference system we consider the gradient flow of the potential function $\Psi(u) := \hat{\Phi}(u) + \frac{1}{2K}d_{\mathcal{U}}^2(u)$ which is given by

$$\dot{u} = -\nabla\Psi(u)^T = -H^T\nabla\Phi(h(u))^T - \tfrac{1}{K}(u - P_{\mathcal{U}}(u))\,. \tag{9.8}$$

In this case, we have $G \equiv \mathbb{I}$ and $K > 0$ takes the role of a *penalty parameter* for the soft penalty term $d_{\mathcal{U}}^2$ that approximately enforces the input constraint $u \in \mathcal{U}$.[2] The system (9.8) is a special case of the AWA (5.1) and, as a consequence, Theorems 5.1 and 5.2 (uniform convergence and robust practical stability) and their corollaries apply as $K \to 0^+$. However, (9.8) is not of the special form (9.1) and convergence to the optimizer of the problem (9.7) is not guaranteed for positive $K > 0$. Neither does (9.8) lend itself to a feedback implementation, because $\nabla\Phi$ is evaluated at $h(u)$ rather than at $h(P_{\mathcal{U}}(u))$ (which is the actual system state for the saturated input).

ii) Anti-Windup Gradient Scheme: As a second type of dynamics we

[2] The penalty $d_{\mathcal{U}}^2$ is illustrative in the context of feedback-based optimization, however, it is not generally practical for numerical optimization, because evaluating $\nabla d_{\mathcal{U}}^2$ requires computing $P_{\mathcal{U}}$. Instead, in numerical applications, it is more common to use a penalty $\|\max\{A_u u - b_u, 0\}\|^2$.

consider

$$\underbrace{\dot{u} = -H^T \nabla \Phi(\overline{x})^T - \tfrac{1}{K}(u - \overline{u})}_{\text{controller}} \quad \underbrace{\overline{u} := P_{\mathcal{U}}(u) \quad \overline{x} := h(\overline{u})}_{\text{physical system}} \tag{9.9}$$

which can be implemented in closed loop because the quantities $\overline{u}$ and $\overline{x}$ are "evaluated" by the physical system at no computational cost (and are assumed to be measurable), which is one of the key features of feedback-based optimization.

Further, because $\mathcal{U}$ is convex and Φ is strongly convex (which implies strong monotonicity), Theorem (9.1) is applicable and guarantees that $\overline{x} = (\overline{u}, \overline{x})$ converges to the optimizer of (9.7). This is confirmed in Figure 9.2.

iii) *Anti-Windup Newton Scheme:* As the final gradient-based anti-windup scheme we consider an anti-windup approximation with $G \equiv Q$ and which is given by

$$\underbrace{\dot{u} = -Q^{-1}\left(H^T \nabla \Phi(\overline{x})^T - \tfrac{1}{K}(u - \overline{u})\right)}_{\text{controller}} \quad \underbrace{\overline{u} := P_{\mathcal{U}}(u) \quad \overline{x} := h(\overline{u})}_{\text{physical system}} . \tag{9.10}$$

The system (9.10) can be implemented in closed loop with a physical plant and approximates a *projected Newton flow* [Ha2, Ex. 5.6]. This fact is noteworthy, because, in general, projected Newton flows do not lend themselves to an easy implementation (e.g., as an iterative algorithm, c.f. Section 7.2).

Even though, as seen in Figure 9.2, $\overline{u}$ converges to the optimizer of (9.7), strictly speaking, Theorem (9.1) is not directly applicable because $Q \neq \mathbb{I}$.

The anti-windup gradient and Newton schemes defined above illustrate some of the key features of feedback-based optimization and anti-windup implementations:

[3] The anti-windup dynamics are simulated with MATLAB using a fixed-step-size forward Euler scheme. The projection on $\mathcal{U}$ is evaluated using `quadprog`. The *nominal* PDS is approximated using a projected forward Euler scheme as $u^{k+1} = P_{\mathcal{U}}(u^k + \alpha f(u^k))$ which is guaranteed to converge uniformly as $\alpha \to 0^+$ (Remark 6.1 and Section 7.2 and [186].

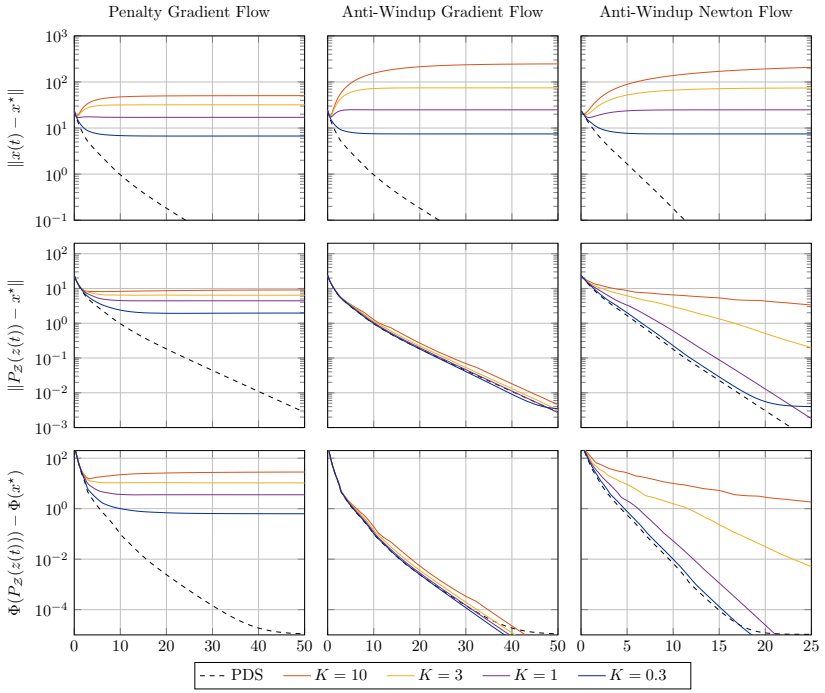

Figure 9.2: Convergence behavior of (9.8), (9.9), and (9.10) for a problem instance of (9.7) with $p = 100$ (input dimension) and $r = 300$ (# of input constraints).[3]

i) Under the conditions of Theorem 9.1, the actual system state and saturated control input converge to the optimizer $u^\star$ of (9.7), even though the internal control variable u does not in general converge to $u^\star$.

ii) In a feedback implementation exploiting input saturation, neither the set $\mathcal{U}$ nor the steady-state disturbance w needs to be known (or estimated). The only model information required is H. Furthermore, recent preliminary theoretical [67] and experimental results for power systems [Ha9] suggest that these feedback schemes are robust against uncertainties in H.

iii) The simulations in Figure 9.2 suggest that the convergence rate of the "projected trajectory" of (9.9) is not affected by the value of K and is equivalent to the convergence rate of the nominal projected gradient flow. In contrast, the convergence rate of the anti-windup Newton scheme (9.10) does depend on K, and one can recover the rate of projected Newton flow only in the limit $K \to 0^+$. An analysis of this observation remains, however, outside the scope of this thesis.

9.5 Anti-Windup Saddle-Point Flows

In feedback-based optimization, constraints on the system state (or output) cannot be directly enforced because they are not directly controllable and often subject to disturbances affecting the physical plant (e.g., an unknown value of w). To enforce state or output constraints, projected saddle-point flows have been proven effective [Ha9, 71, 242]. In this section, we indicate how anti-windup approximations can be combined with this type of dynamical system, even though this leads us slightly outside the scope of our theoretical results. We consider quadratic program

$$\begin{aligned} \text{minimize} \quad & \Phi(x) \\ \text{subject to} \quad & x = h(u),\, u \in \mathcal{U} \\ & x \in \mathcal{X} := \{x \,|\, A_x x \leq b_x\}\,, \end{aligned} \tag{9.11}$$

where Φ, h, and $\mathcal{U}$ are defined as in (9.7) and $\mathcal{X}$ denotes a set of state constraints with $A_x \in \mathbb{R}^{s\times m}$ and $b_x \in \mathbb{R}^s$. To solve (9.11), we consider the projected saddle-point flow

$$\dot{u} = \Pi_{\mathcal{U}}\left[-H^T \nabla\Phi(h(u))^T - H^T A_x^T \mu\right] \qquad \dot{\mu} = \Pi_{\mathbb{R}^s_{\geq 0}}[A_x h(u) - b_x]\,, \tag{9.12}$$

where $\mu \in \mathbb{R}^s$ denotes the dual multipliers associated with the output constraints. The system (9.12) (and special cases in which either primal or dual variables are not projected) has been extensively studied and

convergence is guaranteed, for instance, under strict convexity of Φ. We refer the reader to [58, 117] and references therein.

We approximate (9.12) with a (partial) anti-windup implementation as

$$\underbrace{\begin{aligned}\dot{u} &= -H^T \nabla\Phi(\overline{x})^T - H^T A_x^T \mu - \tfrac{1}{K}(u - \overline{u}) \\ \dot{\mu} &= \Pi_{\mathbb{R}^s_{\geq}0}[A_x \overline{x} - b_x]\end{aligned}}_{\text{controller}} \quad \underbrace{\overline{u} := P_{\mathcal{U}}(u) \quad \overline{x} := h(\overline{u})}_{\text{physical system}} . \tag{9.13}$$

We do not approximate the projected integration of the dual variables with an anti-windup term since the dual variables are often internal variables of the controller, and the projection on the non-negative orthant is easily implementable.

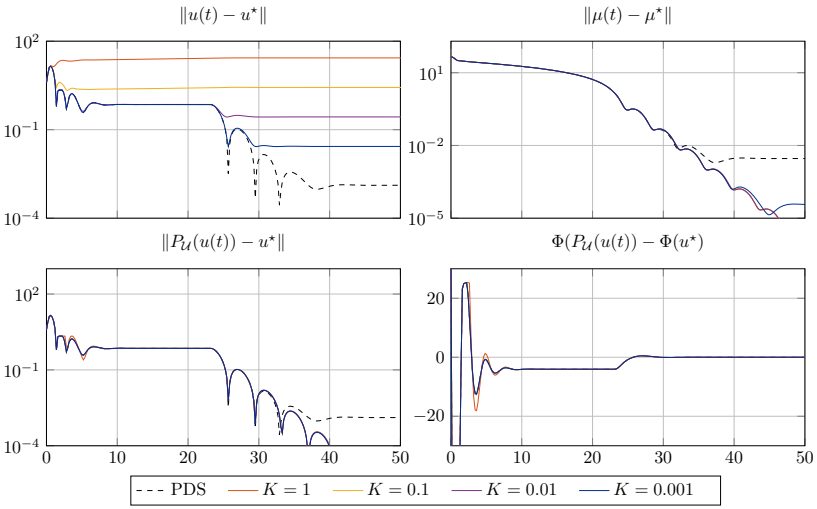

Figure 9.3: Convergence behavior of (9.13) (and the PDS (9.12)) for a problem instance of (9.11) with $p = 3$ (input dimension), $m = 5$ (state dimension), $r = 10$ (# of input constraints), and $s = 5$ (# of state constraints).

Figure 9.3 illustrates the behavior of (9.12) and (9.13). Similarly to

the results for the gradient anti-windup approximations, we observe that u does not, in general, converge to its optimal value. However, the saturated control input $P_{\mathcal{U}}(u)$ (and thereby the actual system state) and the dual variable μ converge to the solution of (9.11).

Theorem 9.1 (robust convergence) does not apply to (9.13). First, while the projected saddle-flow (9.12) is monotone, strong monotonicity is usually not guaranteed [58, 117]. Second, by applying only a partial anti-windup approximation, the vector field remains discontinuous because of the projection of μ on $\mathbb{R}^s_{\geq 0}$.

CHAPTER 10

Convergence of LOP Gradient Descent

In this short chapter we revisit the LOP discretization introduced in Chapter 6 and apply it to projected gradient flows of the form (7.1). Our primary concern is the convergence of trajectories of the LOP-discretized system to the minimizers of the underlying optimization problem. In particular, we aim at strengthening the practical stability results of Section 6.4 in the case of projected gradient flows.

For convenience, we recall all required objects and assumptions: We consider a feasible set

$$\mathcal{C} := \{x \,|\, g(x) \leq 0,\, k(x) \leq 0\}\,, \tag{10.1}$$

where $g : \mathbb{R}^n \to \mathbb{R}^p$ and $k : \mathbb{R}^n \to \mathbb{R}^q$ are continuously differentiable and we have $\mathcal{U} := \{x \,|\, g(x) \leq 0\}$. Further, we consider a metric $G : \mathcal{U} \to \mathbb{S}^n_+$ and a cost function $\Phi : \mathbb{R}^n \to \mathbb{R}$ continuously differentiable in a neighborhood of $\mathcal{U}$.

We study the discrete-time *LOP gradient system*

$$x^+ = x + \alpha\, \Sigma^G_{\mathcal{C}}[-\operatorname{grad}_G \Phi](x, \alpha)\,, \tag{10.2}$$

The content of this chapter and, in particular, the main result Theorem 10.1 is based on [Ha1]. The work in this chapter is a collaboration with V. Häberle.

where $\alpha > 0$ is a fixed step-size and $\Sigma_{\mathcal{C}}^{G}[f] : \mathcal{U} \times \mathbb{R} \to \mathbb{R}^n$ is defined as

$$\Sigma_{\mathcal{C}}^{G}[f](x, \alpha) := \arg\min_{w \in \mathbb{R}^n} \quad \|f(x) - w\|_{G(x)}^2 \tag{10.3a}$$

$$\text{subject to} \quad x + \alpha w \in \mathcal{U} \tag{10.3b}$$

$$k(x) + \alpha \nabla k(x) w \leq 0\,. \tag{10.3c}$$

The central claim of this chapter is that solutions of (10.2) converge to optimizers of

$$\text{minimize} \quad \Phi(x) \qquad \text{subject to} \quad x \in \mathcal{C} \tag{10.4}$$

under the following assumptions (Assumptions 6.1, 6.2, and 6.3):

Assumption 10.1. *For (10.2), the set $\mathcal{C}$ is non-empty and satisfies LICQ. Further, ∇k, ∇g, $\nabla \Phi$ and G are locally Lipschitz.*

Assumption 10.2. *For (10.2), the set $\mathcal{U} = \{x \,|\, g(x) \leq 0\}$ is compact.*

Assumption 10.3. *The set $\mathcal{U} = \{x \,|\, g(x) \leq 0\}$ is convex and, for all $x \in \mathcal{U}$ and all $\alpha > 0$, the feasible set of (10.3) defined as*

$$\tilde{\mathcal{C}}(x, \alpha) := \{w \,|\, g(x + \alpha w) \leq 0,\, k(x) + \alpha \nabla k(x) w \leq 0\}$$

is non-empty and satisfies LICQ.

Recall that under Assumption 10.3, the system (10.2) is well-defined and $\mathcal{U}$ is invariant (Lemma 6.1). Further, $\Sigma_{\mathcal{C}}^{G}[f](x, \alpha)$ is continuous in (x, α) and the map of Lagrange multipliers of (6.2) is continuous (Lemma 6.2).

Hence, the following main result of this chapter guarantees convergence of solutions of (10.2) to the set of critical points of (10.4) and mirrors the stability properties of minimizers from the continuous-time case (Theorem 8.1).

Theorem 10.1. *Under Assumptions 10.1, 10.2, and 10.3, there exists an $\alpha^\star > 0$ such that for all $\alpha \in (0, \alpha^\star)$ we have that*

(i) every solution of (10.2) starting in $\mathcal{U}$ converges to the set of critical points of (10.4),

(ii) every asymptotically stable equilibrium of (10.2), it is a strict local minimizer of (10.4).

Similarly to Theorem 8.1, one can also show that every strict local minimizer of (10.4) is a stable equilibrium of (10.2).

Proof of Theorem 10.1

For (i) we apply Theorem 2.6 and for (ii) we can proceed similarly as in the proof of Theorem 8.1.

Lyapunov Function

Given $x \in \mathcal{U}$, let $\mu_i^\star(x)$ be the Lagrange multiplier of (10.3) for the i-th constraint of (10.3b) with $i = 1, \dots, l$. Since $\mu_i^\star(x)$ is continuous on the compact set $\mathcal{U}$ (by Lemma 6.2), there exists $\gamma \geq \sup_{x \in \mathcal{U}; i=1,\dots,l} \{\mu_i^\star(x)\}$. We may consider the function $V : \mathbb{R}^n \to \mathbb{R}$, defined as

$$V(x) = \Phi(x) + \gamma \left[\textstyle\sum_{i=1}^{l} \max\{0, k(x)\} \right] \tag{10.5}$$

which we show next to be non-increasing along solutions of (6.3).

To prove this claim, note that the optimization problem (10.3) underlying $\Sigma_{\mathcal{C}}^{G}[f](x, \alpha)$ is equivalent to solving

$$\begin{aligned} \arg\min_{w \in \mathbb{R}^n} \quad & \tfrac{\alpha}{2} w^T G(x) w + \alpha \nabla \Phi(x) w && \text{(10.6a)} \\ \text{subject to} \quad & g(x + \alpha w) \leq 0 && \text{(10.6b)} \\ & \alpha \nabla k(x) \leq -k(x)\,, && \text{(10.6c)} \end{aligned}$$

where we have multiplied the objective with α and ignored the constant term in the objective.

Since the feasible set $\tilde{\mathcal{C}}(x, \alpha)$ is convex and satisfies LICQ, the KKT conditions are necessary and sufficient to certify optimality of a solution w of (10.6). Namely, $w \in \mathbb{R}^n$ is a solution if (10.6b)-(10.6c) are satisfied and, for some dual multipliers $\nu \in \mathbb{R}^p_{\geq 0}$ and $\mu \in \mathbb{R}^q_{\geq 0}$, *stationarity*

$$\alpha w^T G(x) + \alpha \nabla \Phi + \alpha \nu^T \nabla g(x + \alpha w) + \alpha \mu^T \nabla k(x) = 0 \tag{10.7}$$

holds and *complementary slackness* is satisfied, i.e., for all $j = 1, ..., p$ and $i = 1, ..., q$, we have

$$\begin{aligned} \nu_j g_j(x + \alpha w) &= 0 \\ \mu_i(k(x) + \alpha \nabla k(x)) &= 0\,. \end{aligned} \tag{10.8}$$

Lemma 10.1. *Under Assumptions 10.1, 10.2, and 10.3, let V be as in (10.5), where γ is an upper bound of the Lagrange multipliers of $\Sigma_{\mathcal{C}}^{G}[-\operatorname{grad}_G \Phi](x, \alpha)$ for all $(u, \alpha) \in \mathcal{U} \times [0, \alpha']$ for some $\alpha' > 0$. Further, assume that $\nabla\Phi$ is ℓ-Lipschitz and ∇k_i is ℓ_i-Lipschitz for all $i = 1, ..., q$ on $\mathcal{U}$. Then, $V(x[k+1]) \leq V(x[k])$ is satisfied for every solution of (10.2) with $x[0] \in \mathcal{C}$ if*

$$\alpha < \alpha^\star := \min\left\{\alpha', \frac{2\lambda_{\min}}{\ell + \gamma \sum_{i=1}^{q} \ell_i}\right\}, \tag{10.9}$$

where $\lambda_{\min} := \min_{x \in \mathcal{U}} \lambda_G^{\min}(x)$.

Proof. The proof is inspired by [22, Lemma 10.4.1]. In the following let $w := \Sigma_{\mathcal{C}}^{G}[-\operatorname{grad}_G \Phi](x, \alpha)$. With the Descent Lemma (Lemma 3.5) we can establish

$$\Phi(x^+) - \Phi(x) \leq \alpha \nabla\Phi(x) w + \tfrac{\ell}{2} ||\alpha w||^2\,. \tag{10.10}$$

Further, from Lemma 6.4, we have, for all $i = 1, \ldots, q$,

$$\max\{k_i(x^+), 0\} \leq \tfrac{\ell_i}{2} ||\alpha w||^2\,. \tag{10.11}$$

Next, we take the inner product of (10.7) and w, which results in

$$\begin{aligned} \alpha \nabla\Phi(x) w = -\alpha w^T G(x) w - \textstyle\sum_{j=1}^{p} \alpha \nu_j \nabla g(x + \alpha w) w \\ - \textstyle\sum_{i=1}^{q} \alpha \mu_i \nabla k_i(x) w\,. \end{aligned}$$

Using (10.8), we replace the summands, i.e.,

$$\begin{aligned} \alpha \nabla\Phi(x) w = -\alpha w^T G(x) w + \textstyle\sum_{j=1}^{p} \nu_j (g_j(x + \alpha w) \\ - \alpha \nabla g_i(x + \alpha w)) + \textstyle\sum_{i=1}^{q} \mu_i k_i(x)\,, \end{aligned}$$

which can be upper bounded by

$$\begin{aligned}\alpha\nabla\Phi(x)w &\leq -\alpha w^T G(x)w + \textstyle\sum_{i=1}^{q} \mu_i k_i(x) \\ &\quad + \textstyle\sum_{j=1}^{p} \nu_j \underbrace{(g_j(x+\alpha w) - \alpha\nabla g_i(x+\alpha w))}_{\leq g_i(x)\leq 0} \\ &\leq -\alpha w^T G(x)w + \textstyle\sum_{i=1}^{q} \mu_i \max\{k_i(x), 0\}\,, \end{aligned} \tag{10.12}$$

where we have exploited the convexity of g_i and the fact that $x \in \mathcal{U}$ and therefore $g_i(x) \leq 0$.

Finally, we can combine (10.10), (10.11) and (10.12) to obtain

$$\begin{aligned}V(x[k+1]) - V(x[k]) \leq &-\alpha\lambda_{\min}\|w\|^2 + \tfrac{\alpha^2}{2}\left[\ell + \gamma\textstyle\sum_{i=1}^{q}\ell_i\right]\|w\|^2 \\ &- \textstyle\sum_{i=1}^{q}(\gamma - \mu_i)\max\{k(x[k]), 0\}\,, \end{aligned} \tag{10.13}$$

where we have used the fact that $x[0] \in \mathcal{C}$ and thus $\max\{k_i(x[0], 0\} = 0$. Choosing α as in (10.9) guarantees $V(x[k+1]) \leq V(x[k])$. □

Convergence to critical points

To show (i) in Theorem 10.1, we can apply Theorem 2.6. Namely, Lemmas 6.1 and 6.2 guarantee that $H(x) := x + \alpha\Sigma_{\mathcal{C}}^{G}[-\operatorname{grad}_G \Phi](x, \alpha)$ is continuous in x and that $\mathcal{U}$ is invariant.

By Lemma 10.1, for the given α satisfying (10.9), the continuous function $V : \mathcal{U} \to \mathbb{R}$ in (10.5) is non-increasing along the solution of (10.2) for all $x \in \mathcal{U}$. Since $\mathcal{U}$ is compact and invariant, any solution of (10.2) is bounded and can be chosen to be complete. Hence, any solution of (10.2) converges to the largest invariant subset of $V^{-1}(r) \cap \{x \in \mathcal{U} \,|\, V(x^+) - V(x) = 0\}$ for some $r \in V(\mathcal{U})$.

To show that convergence is to the set of equilibrium points, we simply show that $V(x^+) - V(x) = 0$ implies $x^+ = x$. Namely, for $V(x^+) - V(x) = 0$, (10.13) reduces to

$$0 \leq \left(\tfrac{\alpha}{2}\left[\ell + \gamma\textstyle\sum_{i=1}^{q}\ell_i\right] - \lambda_{\min}\right)\alpha||w||^2 - \textstyle\sum_{i=1}^{q}(\gamma - \mu_i)\max\{k_i(x[0]), 0\}\,. \tag{10.14}$$

For α as in (10.9), the right-hand side of (10.14) is negative for all $w \neq 0$. Therefore, $V(x^+) - V(x) = 0$ implies that $w = 0$ and $k(x) \leq 0$, i.e.,

$x^+ \in \mathcal{C}$, and hence x is a feasible equilibrium. Note that, an uncertainty in ∇h does not affect the feasibility of equilibria (consider (10.6) for $w = 0$).

Finally, if $w^\star = 0$ solves (10.6) at $x^\star$, and $\nu^\star, \mu^\star$ are the associated Lagrange multipliers, then $(x^\star, \nu^\star, \mu^\star)$ satisfies the KKT conditions of (10.4). In particular, $x^\star$ is feasible for (10.4) and given LICQ[1], the KKT conditions are satisfied for (10.4).

Strict optimality of asymptotically stable equilibria

For (ii), consider the neighborhood $\mathcal{N}(x^\star) \subset \mathcal{U}$ of $x^\star \in \mathcal{C}$, such that any solution $x : \mathbb{N}_{\geq 0} \to \mathcal{U}$ of (10.2) starting at $x[0] \in \mathcal{N}(x^\star)$ converges to $x^\star$. For the given α satisfying (10.9), by Lemma 10.1, we have $V(x[0]) \geq V(x^\star)$, which implies either $\Phi(x[0]) \geq \Phi(x^\star)$, $x[0] \notin \mathcal{C}$, or both.

Hence, if $x[0] \in \mathcal{C}$, we have $V(x[0]) = \Phi(x[0])$ and $\Phi(x[0]) \geq \Phi(x^\star)$ follows. Since this reasoning applies to all $x[0] \in \mathcal{N}(x^\star) \cap \mathcal{C}$, it follows that $x^\star$ is a local minimizer of (10.4).

To see that $x^\star$ is a strict local minimizer of (10.4), assume for the sake of contradiction that for another minimizer $\hat{x} \neq x^\star$ in the region of attraction $\mathcal{N}(x^\star)$ of $x^\star$, such that $\hat{x} \in \mathcal{N}(x^\star) \cap \mathcal{C}$. We have $\Phi(\hat{x}) = \Phi(x^\star)$ because $\hat{x}$ is a minimizer. Consequently, by feasibility we also have $V(\hat{x}) = V(x^\star)$. Nevertheless, the solution $x' : \mathbb{N}_{\geq 0} \to \mathcal{U}$ starting at $x'[0] = \hat{x}$ converges to $x^\star$ by assumption. Since V is non-increasing along the trajectory of (6.3), it follows $V(x'[k]) = V(\hat{x}) = V(x^\star)$ for all k. However, as shown in the proof of (i), $V(x[k+1]) = V(x[k])$ implies that $x[k] \in \mathcal{C}$ and that $x[k]$ is an equilibrium. Consequently, $x^\star$ cannot be asymptotically stable on $\mathcal{N}(x^\star)$ which completes the contradiction.

[1] Assumption 6.3 implies LICQ of $\mathcal{C}$ for all $x \in \mathcal{C}$. This can be verified for $y = h(x)$ for $w = 0$ and $x \in \mathcal{C}$.

CHAPTER 11

Timescale Separation in Feedback-Based Optimization

One of the unique features of feedback-based optimization, compared to numerical optimization, is the interconnection with a physical system subject to various disturbances and uncertainties. So far, we have assumed that this physical system can be characterized by an algebraic steady-state input-to-output map, thus treating it simply as another constraint in our optimization problem.

In this chapter, we relax this assumption and study a plant with dynamics. More precisely, we consider *slow* optimization dynamics interconnected with *fast*, but exponentially stable plant dynamics, and we apply ideas from singular perturbation analysis to derive sufficient conditions for closed-loop stability.

Hence, in the remainder of this chapter we consider a physical plant modeled as

$$\dot{\zeta} = \tilde{f}(\zeta, u)\,, \tag{11.1}$$

where $\zeta \in \mathbb{R}^m$ is the system state, $u \in \mathcal{U} \subset \mathbb{R}^p$ where $\mathcal{U}$ is a non-empty

The results presented in this chapter have evolved from [Ha3]. Preliminary results in [Ha13] were a collaboration with S. Menta. The main result Theorem 11.1 generalizes [Ha3, Thm 3.2] by considering projected gradient flows and not requiring unique and complete solutions.

set of admissible inputs, and $\tilde{f} : \mathbb{R}^m \times \mathcal{U} \to \mathbb{R}^n$ is a vector field.

We will assume that (11.1) is exponentially stable for any fixed $u \in \mathcal{U}$ and admits a steady-state input-to-state map $h : \mathcal{U} \to \mathbb{R}^m$ such that $f(h(u), u) = 0$ for all $u \in \mathcal{U}$.

Hence, given a continuously differentiable cost function $\Phi : \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ we want to steer the plant (11.1) to a steady-state that is the solution of the problem

$$\begin{aligned} &\text{minimize} && \Phi(\zeta, u) \\ &\text{subject to} && \zeta = h(u) \\ & && u \in \mathcal{U}\,. \end{aligned} \tag{11.2}$$

For this purpose, we consider projected gradient dynamics, which define a *reduced* system

$$\dot{u} = \Pi_{\mathcal{U}}^{\tilde{G}} \left[-\operatorname{grad}_{\tilde{G}} \tilde{\Phi} \right](u)\,, \tag{11.3}$$

where $\tilde{G} : \mathcal{U} \to \mathbb{S}_+^m$ is a metric, and $\tilde{\Phi}(u) := \Phi(h(u), u)$ holds for all $u \in \mathcal{U}$.

Assuming that $\tilde{G}$ is continuous, that $\mathcal{U}$ is Clarke regular, and that Φ is continuously differentiable with compact sublevel sets with respect to $\mathcal{U}$, it follows from Proposition 8.2 that (11.3) admits a complete solution for every initial condition $u(0) \in \mathcal{U}$ and that every complete solution converges to the set of critical points of (11.2).

However, in this chapter, we consider the interconnection of (11.3) with (11.1). To realize this closed-loop system we replace any evaluation of $h(u)$ in the computation of $-\operatorname{grad}_{\tilde{G}} \tilde{\Phi}(u)$ with the measurement of the system state ζ. Namely, applying the chain rule, we have

$$-\operatorname{grad}_{\tilde{G}} \tilde{\Phi}(u) = -\tilde{G}(u)^{-1} H(u)^T \nabla \Phi(h(u), u)^T\,,$$

where $H(u)^T := \begin{bmatrix} \nabla h(u)^T & \mathbb{I} \end{bmatrix}$. Hence, the interconnection of (11.3) with (11.1), illustrated in Figure 11.1, is given by

$$\dot{\zeta} = \tilde{f}(\zeta, u) \tag{11.4a}$$

$$\dot{u} = \Pi_{\mathcal{U}}^{\tilde{G}} \left[-\tilde{G}(u)^{-1} H(u)^T \nabla \Phi(\zeta, u)^T \right](u)\,. \tag{11.4b}$$

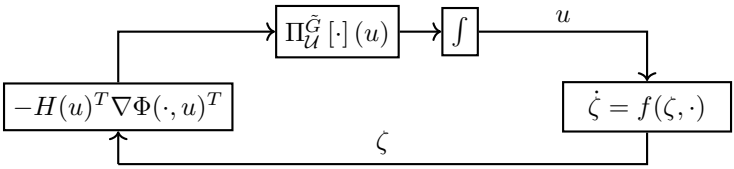

Figure 11.1: Interconnection of projected gradient controller and dynamic plant (dependence of $\nabla\Phi, H$ and $\Pi_{\mathcal{U}}^{\tilde{G}}[\cdot]$ on u is not drawn)

In the remainder of this chapter, we show that trajectories of (11.4) converge to steady states of (11.1) that are critical points of (11.2) whenever the eigenvalues of $\tilde{G}$ are small enough. In other words, we consider $\tilde{G}$ as a design parameter and (matrix-valued) control gain.

Throughout the chapter, we limit ourselves to projected gradient flows. Results for other types of algorithms, such as momentum methods and saddle-point flows, can be found in [Ha3]. However, we do generalize the results from [Ha3] by considering a general metric for the projected gradient flow.

Remark 11.1. The feedback law (11.4b) does not need to be implemented as a state-feedback controller. Assume that only output measurements $y = g(\zeta)$ are available, where $g : \mathbb{R}^m \to \mathbb{R}^q$ is continuously differentiable. This gives rise to a differentiable input-output steady-state map $h_{\mathrm{io}}(u) := g(h(u))$. Further, instead of (11.2), consider the problem

$$\begin{aligned} \underset{y,u}{\text{minimize}} \quad & \Phi_{\mathrm{io}}(y,u) \\ & y = h_{\mathrm{io}}(u) \\ & u \in \mathcal{U}\,, \end{aligned}$$

where Φ_{io} is a cost function depending only on the system output and the control input.

Then, by substituting y with $h_{\mathrm{io}}(u)$ in the objective function, as before, and computing the gradient of the reduced cost function, one arrives at the output-feedback law

$$\dot{u} = \Pi_{\mathcal{U}}^{\tilde{G}}\left[-\tilde{G}(u)^{-1}H_{\mathrm{io}}^T(u)\nabla\Phi_{\mathrm{io}}(y,u)^T\right](u)\,,$$

where $H_{\mathrm{io}}^T(u) := \begin{bmatrix} \nabla h_{\mathrm{io}}(u)^T & \mathbb{I}_p \end{bmatrix}$.

If, moreover, $g(\zeta) = C\zeta + d$ is an affine map, this feedback controller is equivalent to (11.4b). To see this, note that $\nabla h_{\mathrm{io}}(u)^T = \nabla h(u)^T \nabla g(h(u))^T$ with $\nabla g(h(u))^T = C^T$ and therefore

$$H_{\mathrm{io}}^T(u) \nabla \Phi_{\mathrm{io}}(g(\zeta), u)^T = H(u)^T \underbrace{\begin{bmatrix} C^T & 0 \\ 0 & \mathbb{I}_p \end{bmatrix} \nabla \Phi_{\mathrm{io}}(g(\zeta), u)^T}_{\nabla \Phi(\zeta, u)^T},$$

where $\Phi(\zeta, u) := \Phi_{\mathrm{io}}(g(\zeta), u)$.

Consequently, although (11.4b) is formulated in terms of the state ζ, it is not necessarily a state feedback controller and can be implemented as an output feedback law. The formulation in terms of the internal state is nevertheless important for the forthcoming stability analysis. ■

11.1 Assumptions & Preliminaries

First, we formalize the assumption that the fast system dynamics are exponentially stable and admit a well-defined steady-state map:

Assumption 11.1. *Given the system (11.1), $\tilde{f}$ is continuously differentiable, ℓ_ζ-Lipschitz in ζ, and ℓ_u-Lipschitz in u. There exists a differentiable, ℓ-Lipschitz continuous map $h : \mathcal{U} \to \mathbb{R}^m$ such that $\tilde{f}(h(u), u) = 0$ for all $u \in \mathcal{U}$. Finally, there exist $\tau, K > 0$ such that for every initial condition $\zeta_0 \in \mathbb{R}^n$ and every constant $\hat{u} \in \mathcal{U}$ it holds that*

$$\|\zeta(t) - h(\hat{u})\| \leq K \|\zeta_0 - h(\hat{u})\| e^{-\tau t},$$

where $\zeta(t)$ is a solution to (11.1) with $u(t) := \hat{u}$.

The existence of a well-defined steady-state map can, for instance, be guaranteed if $\tilde{f}$ is continuously differentiable and $\nabla_\zeta \tilde{f}(\zeta, u)$ is invertible for all $\zeta \in \mathbb{R}^m$ and all $u \in \mathcal{U}$. In this case, the Implicit Function Theorem guarantees the existence of $h : \mathcal{U} \to \mathbb{R}^m$ such that $\tilde{f}(h(u), u) = 0$ for all $u \in \mathcal{U}$. Lipschitz continuity of h is guaranteed since $\tilde{f}$ is Lipschitz continuous and if all eigenvalues of $\nabla_\zeta \tilde{f}(\zeta, u)$ are bounded away from 0 with some minimal distance for all $(\zeta, u) \in \mathbb{R}^m \times \mathcal{U}$.

Proposition 11.1. *Let Assumption 11.1 hold for (11.1). Then, for any fixed $u \in \mathcal{U}$ there exists a Lyapunov function $W : \mathbb{R}^n \times \mathcal{U} \to \mathbb{R}$ and parameters $\alpha, \beta, \gamma, \omega > 0$ such that*

$$\begin{aligned} \alpha\|\zeta - h(u)\|^2 \leq W(\zeta, u) &\leq \beta\|\zeta - h(u)\|^2 \\ \dot{W}(\zeta, u) &\leq -\gamma\|\zeta - h(u)\|^2 \\ \|\nabla_u W(\zeta, u)\| &\leq \omega\|\zeta - h(u)\| \,. \end{aligned}$$

Proposition 11.1 is a condensation of a standard converse Lyapunov theorem for exponentially stable systems. Only the definition of ω (which captures a Lipschitz-type property of W with respect to u) is non-standard and requires the Lipschitz continuity of h and $\tilde{f}$.

Proof. We use the change of coordinates $z := \zeta - h(u)$ such that (11.1) can be written as

$$\dot{z} = g(z, u) := \tilde{f}(z + h(u), u) \,. \tag{11.5}$$

By Lipschitz continuity of $\tilde{f}$ (in z and u) and h, we have

$$\begin{aligned} \|\nabla_z g\| &= \|\nabla_\zeta \tilde{f}\| \leq \ell_\zeta \\ \|\nabla_u g\| &= \|\nabla_\zeta \tilde{f} \nabla h + \nabla_u \tilde{f}\| \leq \ell_\zeta \ell + \ell_u =: \ell' \,, \end{aligned}$$

where the last bound follows from the triangle inequality and Cauchy-Schwarz. Let $\varphi(t, z, u)$ denote the solution of (11.5) at time t that starts in z for fixed u. Define

$$V(z, u) := \int_0^T \|\varphi(t, z, u)\|^2 dt$$

with $T = \frac{1}{2\tau} \ln(2K^2)$. Analogously to the proof of [143, Th. 4.14], it can be shown that $V(z, u)$ satisfies

$$\begin{aligned} \alpha\|z\|^2 \leq V(z, u) &\leq \beta\|z\|^2 \\ \dot{V}(z, u) &\leq -\gamma\|z\|^2 \\ \|\nabla_z V(z, u)\| &\leq \delta\|z\| \end{aligned}$$

with $\alpha = \frac{1}{2\ell_\zeta}(1-e^{-2\ell_\zeta T})$, $\beta = \frac{K^2}{2\tau}(1-e^{-2\tau T})$, $\gamma = 1/2$, and $\delta = \frac{2K}{\tau-\ell_\zeta}(1-e^{(\ell_\zeta-\tau)T})$.

Next, we proceed similarly as in the proof of [143, Lem. 9.8]. The *sensitivity function* $\varphi_u := \nabla_u \varphi(t,z,u)$ of the solution φ with respect to changes in u [143, Ch. 3.3], satisfies the ODE

$$\dot{\varphi}_u = \nabla_z g(\varphi(t,z,u),u)\varphi_u + \nabla_u g(\varphi(t,z,u),u)$$

with $\varphi_u(0,z,u) = 0$. Using Lipschitz continuity of g we get

$$\begin{aligned}\|\varphi_u(t,z,u)\| &\leq \int_0^t \ell_\zeta \|\varphi_u(s,z,u)\| ds + \int_0^t \ell' ds \\ &= \int_0^t \ell_\zeta \|\varphi_u(s,z,u)\| ds + \ell' t\,.\end{aligned}$$

Applying a special case of the Gronwall inequality [9, Cor. 6.2] (because $\ell' t$ is monotone increasing) yields the bound $\|\varphi_u(t,z,u)\| \leq \ell' t e^{\ell_\zeta t}$, and we have

$$\begin{aligned}\|\nabla_u V(z,u)\| &= \left\| \int_0^T 2\varphi(t,z,u)\varphi_u(t,z,a) dt \right\| \\ &\leq \int_0^T 2K\|z\| e^{-\tau t} \ell' t e^{\ell_\zeta t} dt = \omega' \|z\|\,,\end{aligned}$$

where we have used $\|\varphi(t,z,u)\| \leq K\|z\|e^{-\tau t}$ by exponential stability and $\omega' := \frac{2K\ell'}{(\ell_\zeta-\tau)^2}\left((\ell_\zeta T - \tau T - 1)e^{(\ell_\zeta-\tau)T}+1\right)$.

Finally, we can reverse the coordinate change by defining $W(\zeta,u) := V(\zeta - h(u),u)$. We immediately have the desired bounds

$$\alpha\|\zeta - h(u)\|^2 \leq W(\zeta,u) \leq \beta\|\zeta - h(u)\|^2$$

and the time derivative of W with respect to (3) as

$$\dot{W}(\zeta,u) = \dot{V}(z,u) \leq -\gamma\|\zeta - h(u)\|^2\,.$$

For the final bound, note that we have

$$\|\nabla_u W\| = \| - \nabla_z V \nabla h + \nabla_u V\| \leq \delta\ell\|z\| + \omega'\|z\| = \omega\|z\|\,,$$

where $\omega := \delta\ell + \omega'$. This completes the proof. □

Next, we need to assure that (11.4) is well-posed by making the following assumption:

Assumption 11.2. *For (11.4) under Assumption 11.1, the set $\mathcal{U}$ is Clarke regular, $\tilde{G}$ is continuous, f is continuous, Φ is continuously differentiable, and $u \mapsto \tilde{\Phi}(u) = \Phi(h(u), u)$ has compact sublevel sets $\mathcal{S}_\ell := \{u \in \mathcal{U} \,|\, \tilde{\Phi}(u) \leq \ell\}$.*

As a first consequence, Assumption 11.2 guarantees that the reduced PDS (11.3) admits complete solutions that converge to critical points of (11.2) by Proposition 8.2.

Secondly, similar to Example 5.1, the system (11.4) can be expressed as a general PDS, i.e., as a constrained differential inclusion of the form $\dot{x} \in \Pi_{\mathcal{C}}^{G}[f](x)$, $x \in \mathcal{C}$ with $\mathcal{C} := \mathbb{R}^m \times \mathcal{U}$,

$$x := \begin{bmatrix} \zeta \\ u \end{bmatrix}, \qquad G(x) := \begin{bmatrix} \mathbb{I} & 0 \\ 0 & \tilde{G}(u), \end{bmatrix}$$

and

$$f(x) := \begin{bmatrix} \tilde{f}(\zeta, u) \\ -\tilde{G}(u)^{-1} H(u)^T \nabla \Phi(\zeta, u)^T \end{bmatrix}.$$

Under Assumption 11.2, the set $\mathcal{C}$ is Clarke regular, G is continuous and f is continuous. Therefore, by Corollary 4.3, (11.4) admits a (local Carathéodory) solution for every initial condition $(\zeta(0), u(0)) \in \mathbb{R}^m \times \mathcal{U}$.

Thirdly, under Assumption 11.2, equilibria of (11.4) coincide with critical points of (11.2):

Proposition 11.2. *Every critical point of (11.2) is a weak equilibrium of (11.4) and every weak equilibrium $(y^\star, u^\star)$ with $y^\star = h(u^\star)$ is a critical point of (11.2).*

Proof. First, by Lemma 2.1, every critical point of (11.2) is also a critical point of

$$\text{minimize} \quad \tilde{\Phi}(u) \qquad \text{subject to} \quad u \in \mathcal{U}\,. \tag{11.6}$$

Hence, if $(\zeta^\star, u^\star)$ is a weak equilibrium with $\zeta^\star = h(u^\star)$ (which might not be the case if $\tilde{f}$ admits multiple steady states for the same input $u^\star$) we have

$$0 \in \Pi_{\mathcal{U}}^{\tilde{G}} \left[-\tilde{G}(u^\star)^{-1} H(u^\star)^T \nabla \Phi(\zeta^\star, u^\star)^T \right](u^\star) = \Pi_{\mathcal{U}}^{\tilde{G}} \left[-\operatorname{grad}_{\tilde{G}} \Phi \right](u^\star).$$

Consequently, $-\operatorname{grad}_{\tilde{G}} \Phi(u^\star) \in N_{u^\star}^{\tilde{G}} \mathcal{U}$ and therefore $u^\star$ is a critical point of (11.6) by Proposition 8.1. The reverse implication follows similarly. □

Finally, we need the following Lipschitz-like assumption on H and Φ:

Assumption 11.3. *Under Assumption 11.1, for the objective function Φ and the steady-state map h in (11.2) there exists $L > 0$ such that*

$$\left\| H(u)^T \left(\nabla \Phi(\zeta', u) - \nabla \Phi(\zeta, u) \right)^T \right\| \leq L \|\zeta' - \zeta\| \tag{11.7}$$

for all $\zeta', \zeta \in \mathbb{R}^m$ and all $u \in \mathcal{U}$.

Assumption 11.3 is, for instance, satisfied if $\nabla \Phi$ is $\tilde{L}$-Lipschitz continuous, in which case L can be chosen as $L := \ell \tilde{L}$ where ℓ is the Lipschitz constant of $H(u)$ (which exists by Assumption 11.1). However, in practice, a tighter bound can often be established by exploiting the structure of $H(u)$ and $\Phi(\zeta, u)$.

In Section 11.3 we will discuss the non-example of a subgradient flow interconnected with a dynamic plant. This setup does not satisfy Assumption 11.3 and is not asymptotically stable, thus illustrating the necessity of Assumption 11.3.

11.2 Main Stability Result

In the following main result of this chapter we consider the metric $\tilde{G}$ as a design parameter. In particular, the forthcoming stability bound illustrates the trade-off between the decay properties of the *fast* physical system and the gain of the *slow* optimization dynamics. This behavior will also be illustrated in Section 11.3 with the help of numerical examples. The following result is generalized from [Ha3, Thm 3.2]:

Theorem 11.1. *Consider* (11.4) *and let Assumptions 11.1, 11.2, and 11.3 hold. Assume that*

$$\sup_{u\in\mathcal{U}} \left\|\tilde{G}(u)^{-1}\right\| < \frac{\gamma}{\omega L} \tag{11.8}$$

holds, where $L > 0$ is a constant satisfying (11.7)*. Furthermore, γ and ω are constants associated with a Lyapunov function $W(\zeta, h(u))$ for* (11.1) *according to Proposition 11.1.*

Then, (11.4) *admits a complete solution for every initial condition $(\zeta(0), u(0)) \in \mathbb{R}^m \times \mathcal{U}$ and every complete solution converges to the set of critical points of* (11.2)*.*

Furthermore, the following statements hold:

(i) Every minimizer of (11.2) *is a strong equilibrium point of* (11.4)*.*

(ii) An asymptotically stable critical point of (11.4) *is a strict minimizer of* (11.2)*.*

(iii) Every minimizer of (11.2) *is stable for* (11.4)*.*

In particular, combining Proposition 11.2 and Theorem 11.1 we recover the same relations between minimizers, critical points and equilibria of (11.2) and (11.4) as in Proposition 8.1:

$$\text{minimizer} \implies \text{strong eq.} \implies \text{crit. pt.} \iff \text{weak eq.}$$

In many practical applications, the right-hand side of (11.8) can be estimated. The parameter L can be derived from model information. The parameters γ and ω can often be deduced from measurements of the decay rate of the open-loop system without explicitly formulating a Lyapunov function [217, Thm. 5.17].

More importantly, however, Theorem 11.1 guarantees that there always exist a small enough control gain such that the interconnected system is asymptotically stable. For instance, if $\tilde{G}(u)^{-1} = \epsilon G^{-1}$ where $G \succ 0$ is constant, the bound (11.8) expresses a design condition on the *global control gain* $\epsilon > 0$.

Corollary 11.1. *Consider the same setup as in Theorem 11.1 and assume* $\tilde{G}^{-1} \equiv \epsilon \mathbb{I}_n$. *Then, for all* $\epsilon < \epsilon^\star := \frac{\gamma}{\omega L}$ *the system* (11.4) *is asymptotically stable.*

Remark 11.2. If the integrator of the controller is grouped with the plant to make the feedback law purely proportional, then $\frac{\omega}{\gamma}$ is an estimate of the input-to-state (ISS) gain of the augmented plant and $\sup_{u \in \mathbb{R}^p} \|\tilde{G}(u)\| \cdot L$ is the ISS gain of the proportional feedback law. Hence, the condition (11.8) can also be interpreted as a small gain result: The product of the two gains has to be less than unity. ■

Proof of Theorem 11.1

Our proof is inspired by ideas from singular perturbation analysis [143, 146, 215] and we work towards an application of the invariance principle Theorem 2.5. For this, we consider functions of the form

$$V(\zeta, u) = (1 - \delta)\tilde{\Phi}(u) + \delta W(\zeta, u)\,,$$

where $0 < \delta < 1$ is a convex combination coefficient that is to be determined. In this context, recall that $\tilde{\Phi}(u) := \Phi(h(u), u)$ and that $W(\zeta, u)$ is of the form $W(\zeta, u) := W(\zeta - h(u), u)$, where $\zeta - h(u)$ is referred to as *boundary-layer* error coordinates in singular perturbation terminology and measures the deviation from the steady state.

First, we establish the requirement for V to be non-increasing along the trajectories of (11.4). We then show that the level sets of V are compact (and hence invariant), and therefore the invariance principle is applicable. Finally, we prove the connection between stability and optimality of equilibria.

Asymptotic Convergence

Throughout the rest of this section let $Q(u) := \tilde{G}^{\frac{1}{2}}(u)$ be the unique positive definite square root of $\tilde{G}(u) \in \mathbb{S}^m_+$ for all $u \in \mathcal{U}$ and define $\kappa := \sup_{u \in \mathcal{U}} \|Q(u)^{-1}\|$.

We first establish the key condition for V to be non-increasing along trajectories of (11.4) which is adapted from [Ha3, Lem 3.1]:

Lemma 11.1. *If for some $\delta \in (0,1)$, the 2-by-2-matrix*

$$\Lambda := \begin{bmatrix} -(1-\delta) & \frac{1}{2}\left(\kappa L(1-\delta) + \kappa\omega\delta\right) \\ \frac{1}{2}\left(\kappa L(1-\delta) + \kappa\omega\delta\right) & -\gamma\delta \end{bmatrix} \tag{11.9}$$

is negative definite, then $\frac{d}{dt}V(\zeta(t),u(t)) \leq 0$. Furthermore, if $\Lambda \prec 0$ then $\frac{d}{dt}V(\zeta,u) = 0$ implies that $\zeta = h(u)$ and, omitting the argument t, that $0 = \Pi_{\mathcal{U}}^{\tilde{G}}\left[-\tilde{G}(u)^{-1}H(u)^T\nabla\Phi(\zeta,u)^T\right]$.

Proof. By absolute continuity of any solution $(\zeta,u) : [0,T] \to \mathbb{R}^{m+p}$ of (11.4), the time derivative of V along (ζ,u) exists for almost all $t \in [0,T]$ and is given by

$$\begin{aligned}\tfrac{d}{dt}(\zeta(t),u(t)) = (1-\delta)\nabla\tilde{\Phi}(u)g(\zeta,u) \\ + \delta\nabla_\zeta W(\zeta,u)\tilde{f}(\zeta,u) + \delta\nabla_u W(\zeta,u)g(\zeta,u)\,,\end{aligned} \tag{11.10}$$

where we have omitted the argument t for brevity and defined

$$g(\zeta,u) := \Pi_{\mathcal{U}}^{\tilde{G}}\left[-\tilde{G}(u)^{-1}H(u)^T\nabla\Phi(\zeta,u)^T\right] = v - \eta$$

with $v = -\tilde{G}(u)^{-1}H(u)^T\nabla\Phi(\zeta,u)^T$ and $\eta \in N_u^{\tilde{G}}\mathcal{U}$. In particular, the decomposition of $g(\zeta,u)$ into v and η exists according to Lemma 4.1.

Each of the terms in (11.10) can be bounded. Namely, for the first term we can do a rearrangement, apply Cauchy-Schwarz (first inequality below), and use the definition of $\|\cdot\|_{\tilde{G}(u)}$ and Assumption 11.3 (second inequality below) to write

$$\begin{aligned}&\nabla\tilde{\Phi}(u)g(\zeta,u) \\ &\quad= -\nabla\Phi(h(u),u)H(u)g(\zeta,u) \\ &\quad= -\left(\nabla\Phi(h(u),u) - \nabla\Phi(\zeta,u)\right)H(u)Q(u)^{-1}Q(u)g(\zeta,u) \\ &\qquad\quad - \nabla\Phi(\zeta,u)H(u)g(\zeta,u) \\ &\quad\leq \left\|H(u)^T\left(\nabla\Phi(h(u),u) - \nabla\Phi(\zeta,u)\right)^T\right\| \left\|Q(u)^{-1}\right\| \|Q(u)g(\zeta,u)\| \\ &\qquad\quad - \nabla\Phi(\zeta,u)H(u)g(\zeta,u) \\ &\quad\leq \kappa L\|\zeta - h(u)\|\|g(\zeta,u)\|_{\tilde{G}(u)} - \nabla\Phi(\zeta,u)H(u)g(\zeta,u)\,.\end{aligned}$$

Next, according to Lemma 4.3 it we have $\langle v, g(\zeta,u)\rangle_{\tilde{G}(x)} \geq \|g(\zeta,u)\|^2_{\tilde{G}(u)}$ and therefore it holds that

$$\begin{aligned}
&\nabla\tilde{\Phi}(u)g(\zeta,u)\\
&\quad\leq \kappa L\|\zeta - h(u)\|\|g(\zeta,u)\|_{\tilde{G}(u)} - \nabla\Phi(\zeta,u)H(u)\tilde{G}(u)^{-1}\tilde{G}(u)g(\zeta,u)\\
&\quad= \kappa L\|\zeta - h(u)\|\|g(\zeta,u)\|_{\tilde{G}(u)} - v^T\tilde{G}(u)(v-\eta)\\
&\quad\leq \kappa L\|\zeta - h(u)\|\|g(\zeta,u)\|_{\tilde{G}(u)} - \|g(\zeta,u)\|^2_{\tilde{G}(u)}\,.
\end{aligned}$$

Thus, the first term in (11.10) is bounded in $\|\zeta - h(u)\|$ and $\|g(\zeta,u)\|_{\tilde{G}(u)}$. The second term in (11.10), is bounded, by Proposition 11.1, as

$$\nabla_\zeta W(\zeta,u)\tilde{f}(\zeta,u) \leq -\gamma\|\zeta - h(u)\|^2\,.$$

Finally, for the third term in (11.10) we can apply Cauchy-Schwarz, the definition of $\|\cdot\|_{\tilde{G}(u)}$, and the fact that $\|\nabla_u W(\zeta,u)\| \leq \omega\|\zeta - h(u)\|$ from Proposition 11.1 to arrive at

$$\begin{aligned}
\nabla_u W(\zeta,u)g(\zeta,u) &= \nabla_u W(\zeta,u)Q(u)^{-1}Q(u)g(\zeta,u)\\
&\leq \|\nabla_u W(\zeta,u)\|\left\|Q(u)^{-1}\right\|\|Q(u)g(\zeta,u)\|\\
&\leq \kappa\omega\|\zeta - h(u)\|\|g(\zeta,u)\|_{\tilde{G}(u)}\,.
\end{aligned}$$

Combining these three bounds, the time derivative of V along trajectories of (11.4) is bounded almost everywhere by a quadratic function that can be rewritten in matricial form (and omitting the argument t) as

$$\tfrac{d}{dt}V(\zeta,u \leq \begin{bmatrix}\|g(\zeta,u)\|_{\tilde{G}(u)}\\ \|\zeta - h(u)\|\end{bmatrix}^T \Lambda \begin{bmatrix}\|g(\zeta,u)\|_{\tilde{G}(u)}\\ \|\zeta - h(u)\|\end{bmatrix},$$

where Λ is given by (11.9). Clearly, if $\Lambda \prec 0$, then $\frac{d}{dt}V(\zeta(t),u(t)) \leq 0$.

Finally, we note that if $\Lambda \prec 0$, then $\frac{d}{dt}V(\zeta(t),u(t)) = 0$ holds if and only if $\|\zeta(t) - h(u(t))\| = 0$ and $\|g(\zeta(t),u(t))\| = 0$. In other words, when the system (11.4) is at steady state. This, in turn, implies by Proposition 11.2 that the system is at a critical point of (11.2) which completes the proof of Lemma 11.1. □

In order to choose an appropriate δ that guarantees $\Lambda \prec 0$ and therefore $\dot{V}(t) \leq 0$, we require the following result [146, pp.296]:

Lemma 11.2. *Consider a 2×2-matrix defined as*

$$\Lambda := \begin{bmatrix} -(1-\delta)\alpha_1 & \frac{1}{2}\left((1-\delta)\beta_1 + \delta\beta_2\right) \\ \frac{1}{2}\left((1-\delta)\beta_1 + \delta\beta_2\right) & -\delta(\alpha_2 - \xi) \end{bmatrix},$$

where $\beta_1, \beta_2, \xi \in \mathbb{R}$, $\delta \in (0,1)$, and $\alpha_1, \alpha_2 > 0$. If $\frac{\alpha_1\alpha_2}{\alpha_1\xi + \beta_1\beta_2} > 1$ and $\delta = \frac{\beta_1}{\beta_1+\beta_2}$, then Λ is negative definite.

By setting $\alpha_1 = 1$, $\alpha_2 = \gamma$, $\xi = 0$, $\beta_1 = \kappa L$ and $\beta_2 = \kappa\omega$, we conclude that in our case $\Lambda \prec 0$ whenever we choose

$$\frac{\gamma}{\kappa^2\omega L} > 1 \quad \text{and} \quad \delta = \frac{\omega}{\omega + L},$$

thus recovering the bound (11.8) in Theorem 11.1 which guarantees that V is non-increasing along trajectories of (11.4).

Next, we show that the sublevel sets of V are compact.

Lemma 11.3. *Consider a system (11.1) satisfying Assumption 11.1 with a Lyapunov function $W : \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ and a steady-state map $h : \mathbb{R}^p \to \mathbb{R}^m$. Further, let $V(\zeta, u) := (1-\delta)\Phi(u) + \delta W(\zeta, u)$ where $\Phi : \mathbb{R}^p \to \mathbb{R}$ is continuous and has compact level sets, and $\delta \in (0,1)$. Then, V has compact sublevel sets.*

Proof. Consider a sublevel set $\Omega_c := \{(\zeta, u) \,|\, V(\zeta, u) \leq c\}$, for some $c \in \mathbb{R}$. Since we have $W(\zeta, u) \geq 0$, $(\zeta, u) \in \Omega_c$ implies that $(1-\delta)\Phi(u) \leq c$. But since Φ has compact sublevel sets, there exist U such that $\|u\| \leq U$ for all $(\zeta, u) \in \Omega_c$.

On the compact set $\{u \,|\, \|u\| \leq U\}$ the continuous function $(1-\delta)\Phi(u)$ is also lower bounded by some value L. We therefore have that $\delta W(\zeta, u) \leq c - L$ in Ω_c. As W is positive definite, we must have that $\|\zeta - h(u)\|^2 \leq (c-L)/\gamma$. We then have

$$\begin{aligned} \|\zeta\| &\leq \|\zeta - h(u)\| + \|h(u)\| \\ &\leq \sqrt{\frac{c-L}{\gamma}} + \|h(u) - h(0)\| + \|h(0)\| \\ &\leq \sqrt{\frac{c-L}{\gamma}} + \ell(U-0) + \|h(0)\|, \end{aligned}$$

where ℓ is the Lipschitz constant of h, and therefore $\|\zeta\|$ is also bounded for all $(\zeta, u) \in \Omega_c$. □

It follows from Lemma 11.3 and if (11.8) holds, that the compact sublevel sets of V are invariant and consequently, solutions of (11.4) cannot escape to the horizon in finite time which guarantees the existence of complete solutions.

Furthermore, we can apply Theorem 2.5 and conclude that all complete solutions of (11.4) converge to the closure of the largest invariant subset for which $\frac{d}{dt}V = 0$. This, in turn, coincides with the set of critical points of (11.2) by Lemma 11.1.

Relation between Stability and Optimality

The proofs of (i), (ii), and (iii) in Theorem 11.1 are analogous to the proofs of the corresponding properties in Proposition 8.1.

For (i), let $(\zeta^\star, u^\star)$ be a minimizer of (11.2) and hence a weak equilibrium of (11.4) by Proposition 11.2. Assume for the sake of contradiction there exists a solution $(\zeta, u) : [0, T] \to \mathbb{R}^{m+p}$ such that $(\zeta(0), u(0)) = (\zeta^\star, u^\star)$ but $(\zeta(t), u(t)) \neq (\zeta^\star, u^\star)$ for some $t \in [0, T]$. The sublevel set $\mathcal{S}_{\ell^\star}$ of V with $\ell^\star = V(\zeta^\star, u^\star)$ is invariant and $x(t) \in \mathcal{S}_{\ell^\star}$ for all $t \in [0, T]$ by Lemma 11.1 as long as (11.8) holds. If $(\zeta^\star, u^\star)$ is a minimizer of (11.2), it is also a minimizer of V on $\mathbb{R}^m \times \mathcal{U}$ since $\zeta^\star = h(u^\star)$. Since V is non-increasing, it follows that $V(\zeta(t), u(t)) = V(\zeta^\star, u^\star)$ for all $t \in [0, T]$, thus $\frac{d}{dt}V(x(t)) = 0$ for almost all $t \in [0, T]$ which implies that $\dot{\zeta}(t) = 0$ and $\dot{u}(t) = 0$ for almost all $t \in [0, T]$ by Lemma 11.1, thus completing the contradiction.

To show (ii), let $(\hat{\zeta}, \hat{u})$ be an asymptotically stable point of (11.4) and let $\mathcal{N} \subset \mathbb{R}^m \times \mathcal{U}$ such any solution $(\zeta, u) : [0, \infty) \to \mathbb{R}^m \times \mathcal{U}$ of (11.4) with $(\zeta(0), u(0)) \in \mathcal{N}$ converges to $(\hat{\zeta}, \hat{u})$. Note, that V is non-increasing along (ζ, u) and $V(\hat{\zeta}, \hat{u}) \leq V(\zeta(0), u(0))$ for any $(\zeta(0), u(0)) \in \mathcal{N}$ and therefore $(\hat{\zeta}, \hat{u})$ is a minimizer of V on $\mathbb{R}^m \times \mathcal{U}$. However, any minimizer of V satisfies $\hat{\zeta} = h(\hat{u})$ and is therefore a minimizer of (11.2).

To see that $(\hat{\zeta}, \hat{u})$ is a strict minimizer of (11.2), assume for the sake of contradiction that for some $(\tilde{\zeta}, \tilde{u}) \in \mathcal{N}$ with $\tilde{\zeta} = h(\tilde{u})$ in the re-

gion of attraction of $(\hat{\zeta}, \hat{u})$ it holds that $\Phi(\tilde{u}) \leq r$. Every solution $(\zeta', u') : \mathbb{R}_{\geq 0} \to \mathbb{R}^m \times \mathcal{U}$ to (11.4) with $(\zeta'(0), u'(0)) = (\tilde{\zeta}, \tilde{u})$ nevertheless converges to $(\hat{\zeta}, \hat{u})$ by assumption. Therefore, it must hold that $\int_0^{+\infty} \frac{d}{dt} V(\zeta'(\tau), u'(\tau)) d\tau = 0$ and since $\frac{d}{dt} V(\zeta'(t), u'(t)) \leq 0$, it follows that $\frac{d}{dt} V(\zeta'(t), u'(t)) = 0$ for almost all $t \geq 0$. But as a consequence of Lemma 11.1, all points (ζ, u) in the limit set are weak equilibria points, this holds in particular for $(\tilde{\zeta}, \tilde{u})$ and therefore $(\hat{\zeta}, \hat{u})$ cannot be asymptotically stable. For (iii), i.e., to show that strict local minimizers of (11.6) are stable, let $\mathcal{V} \subset \mathbb{R}^m \times \mathcal{U}$ be any compact (relative) neighborhood of $(\zeta^\star, u^\star)$ in which $u^\star$ is a strict minimizer of $\tilde{\Phi}(u)$. In particular, without loss of generality assume that $\{u \,|\, (\zeta, u) \in \mathcal{V}\}$ is not a singleton. Otherwise, stability is trivially guaranteed by Assumption 11.1. We construct a neighborhood $\mathcal{W} \subset \mathbb{R}^m \times \mathcal{U}$ of $(\zeta^\star, u^\star)$ such that every trajectory starting in $\mathcal{W}$ remains in $\mathcal{V}$, thus proving stability.

Hence, consider the function V in the previous section, and let α be such that $V(\zeta^\star, u^\star) < \alpha < \min_{(\zeta,u)\in\partial\mathcal{V}} V(\zeta, u)$ where $\partial\mathcal{V}$ denotes the boundary of $\mathcal{V}$ (relative to $\mathbb{R}^m \times \mathcal{U}$). Define $\mathcal{W} := \{(\zeta, u) \in \mathbb{R}^m \times \mathcal{U} \,|\, V(\zeta, u) \leq \alpha\} \subset \mathcal{V}$ which has a non-empty (relative) interior because $V(\zeta^\star, u^\star) < \alpha$. Furthermore, as a sublevel set of V, the set $\mathcal{W}$ is invariant since $\dot{V}(\zeta, u) \leq 0$ (with the proper choice of δ according to Lemma 11.1). This establishes stability of $(\zeta^\star, u^\star)$.

Hence, the proof of Theorem 11.1 is complete.

11.3 Illustration of Gradient-Based Controllers and LTI Systems

In the following, we consider variations and extensions of the interconnected projected gradient flow (11.4) and discuss the stability bound in Theorem 11.1. For simplicity, throughout this section, instead of (11.1), we consider an *linear time-invariant* (lti) plant of the form

$$\dot{\zeta} = A\zeta + Bu + d\,, \tag{11.11}$$

where $A \in \mathbb{R}^{m\times m}$, $B \in \mathbb{R}^{m\times p}$ and $d \in \mathbb{R}^m$. Whenever A is Hurwitz, (11.1) is satisfied and parameters β, γ, ω of a quadratic Lyapunov func-

tion according to Proposition 11.1 can be easily computed.

Furthermore, for (11.11), the steady-state input-output map is given by

$$h(u) := -A^{-1}Bu - A^{-1}d$$

and thus $\nabla h(u) \equiv -A^{-1}B$ is constant.

Unconstrained Gradient Flows

In general, the conservativeness of the bound (11.8) depends largely on the specific problem. Figures 11.2a and 11.2b illustrate this fact based on two random problem instances for an unconstrained gradient flow (i.e., $\mathcal{U} = \mathbb{R}^p$). In both examples, we consider, for simplicity, the case where $\tilde{G} \equiv \epsilon \mathbb{I}_n$ (i.e., as in Corollary 11.1), the cost function Φ is convex quadratic, and the plant is lti. In each case, we have $n = 20$ (state dimension) and $p = 5$ (input dimension).

In both cases, the interconnected gradient system (11.4) is stable for values of ϵ larger than $\epsilon^\star = \frac{\gamma}{\zeta L}$. For $\epsilon = \epsilon^\star$, the feedback interconnection illustrated in Figure 11.2a exhibits a similar convergence rate as the reduced system. However, for ϵ larger than $10\epsilon^\star$ instability of the interconnected system occurs.

For the second example (Figure 11.2b), the stability bound on ϵ is more conservative. For $\epsilon = 200\epsilon^\star$, the interconnected system is stable. However, the convergence rate compared to the reduced system significantly deteriorates. For this problem instance, instability occurs for values of ϵ larger than $290\epsilon^\star$.

These examples illustrate the variable degree of conservativeness of our stability bound and the gradual performance degradation as the stability limit of the interconnected system is reached.

Projected Gradient Flows

One of the surprising aspects of Theorem 11.1 is the fact that the stability bound (11.1) is independent of whether the reduced system (11.3) is an

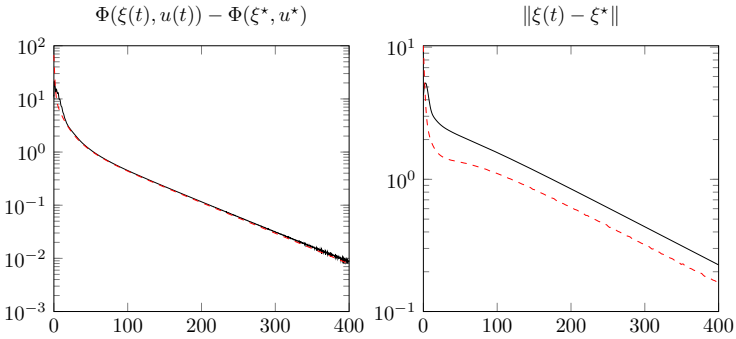

(a) $\epsilon = \epsilon^\star$ (according to Corollary 11.1). Both, the reduced gradient flow (11.3) (dashed) and the system (11.4) (solid) converge to the unique optimizer of (11.2) with similar convergence rate.

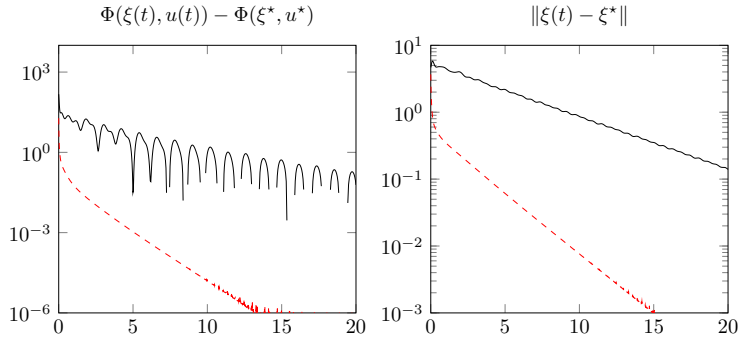

(b) $\epsilon = 200\epsilon^\star$. Both, the reduced (11.3) (dashed) and the interconnected system (11.4) (solid) converge. However, the convergence rate of is significantly worse for the interconnected system.

Figure 11.2: Examples of continuous-time unconstrained feedback gradient flow illustrating the varying degrees of convervativeness of $\epsilon^\star$ (Corollary 11.1) for two different problem instances.

unconstrained or a projected gradient flow.[1] Figure 11.3 illustrates this

[1] In fact, considering an unconstrained gradient flow makes Assumption 11.1 more restrictive, because a (nonlinear) plant needs to be exponentially stable for any fixed

fact by comparing a random instance of the interconnected system (11.4) (with stable LTI plant) in two cases: with $\mathcal{U} = \mathbb{R}^p$ (i.e., resulting in an unconstrained gradient flow), and with $\mathcal{U} = [-1, 1]^p$. The bound $\epsilon^\star$ in Corollary 11.1 is the same in both cases.

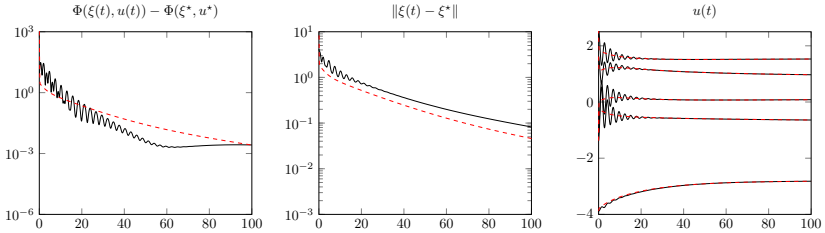

(a) Unconstrained feedback gradient flow.

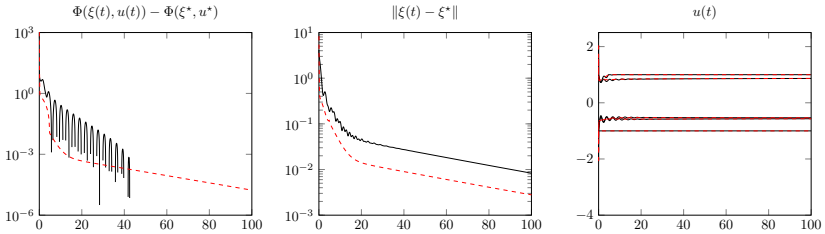

(b) Projected feedback gradient flow where u is constrained to the unit cube.

Figure 11.3: Comparison between an unconstrained and projected feedback gradient flow for the same LTI system, cost function, and the same $\epsilon = \epsilon^\star$ (according to Corollary 11.1). The right panels show the components of the unsaturated and saturated control signal, respectively.

In particular, the simulations in Figure 11.3 indicate that the stability of the projected feedback gradient flow implementation does not deteriorate compared to the unconstrained feedback gradient flow. If anything, the saturation of the control signal u onto $[-1, 1]^p$ reduces transient oscillations.

input $u \in \mathbb{R}^p$.

Newton Gradient Flows

The classical Newton method finds widespread application in numerical optimization as a second-order method (i.e., requiring information about second-order derivatives) with superlinear convergence [191, Ch. 3.3]. The continuous-time limit of the Newton method is given by a simple gradient flow of the form (11.3) (with $\mathcal{U} = \mathbb{R}^p$), namely,

$$\dot{u} = -\epsilon \operatorname{grad}_{\nabla^2 \tilde{\Phi}} \tilde{\Phi}(u) = -\epsilon (\nabla^2 \tilde{\Phi}(u))^{-1} \nabla \tilde{\Phi}(u)^T , \tag{11.12}$$

where $\epsilon > 0$ serves to adjust the convergence rate.

For (11.12) to be well-defined, we assume that $\tilde{\Phi}$ is μ-strongly convex and twice continuously differentiable such that the metric $\nabla^2 \tilde{\Phi}(u)$ is well-defined for all $u \in \mathbb{R}^p$. Hence, convergence to the unique equilibrium is exponential and moreover *isotropic*, i.e., trajectories approach the equilibrium from all directions with the same speed. In other words, the linearization around the equilibrium $u^\star$ is given by $\dot{u} = -\epsilon(u - u^\star)$.

In terms of stability, Newton flows are well-suited for the implementation as feedback controllers. Although the evaluation (or estimation) of the inverse Hessian of $\tilde{\Phi}$ can be computationally demanding.

Theorem 11.1 can be directly applied to give a condition for asymptotic stability in closed loop. Namely, since $\tilde{\Phi}$ is μ-strongly convex, we have that $\sup_{u \in \mathbb{R}^p} \|\epsilon (\nabla^2 \tilde{\Phi}(u))^{-1}\| \leq \epsilon/\mu$ and therefore the following holds.

Corollary 11.2. *Consider the same setup as in Theorem 11.1 and assume that $\tilde{\Phi}$ is μ-strongly convex and twice continuously differentiable. With the metric $\tilde{G}(u) := \epsilon \nabla^2 \tilde{\Phi}(u)$, the closed-loop system (11.4) is asymptotically stable and converges to the set of critical points of (11.6) whenever*

$$\epsilon < \frac{\gamma \mu}{\omega L} .$$

Corollary 11.2 applies, in fact, to *projected Newton flows* (as already encountered in Section 7.2). If, addition, $\mathcal{U}$ is convex then convergence is to the unique minimizer of (11.6).

Compared to the previous results, the above bound on ϵ is invariant for a uniform scaling of $\tilde{\Phi}$ by a constant $\alpha > 0$ since α scales both L and

μ by the same factor. Furthermore, the requirement that $\tilde{\Phi}$ is strongly convex implies the optimizer's uniqueness, but it does not necessarily require that the problem (11.6) is itself convex.

Figure 11.4 illustrates, similarly to Figures 11.2a and 11.2b, the interconnection of an LTI plant with a Newton flow for a quadratic function. In this case $\tilde{G} \equiv \epsilon(\nabla^2\Phi)^{-1}$ is constant. As before, the interconnected system is stable even for ϵ larger than the theoretical bound in Corollary 11.2. However, the convergence rate gradually worsens compared to the reduced system.

Subgradient Flow (Non-Example)

Subgradient flows are the continuous-time version of subgradient descent and generalize gradient flows to the case where $\tilde{\Phi}$ is not differentiable. Namely, assuming that $\tilde{\Phi}$ is convex, its subgradient at $u \in \mathbb{R}^p$ is defined as the set

$$\partial\tilde{\Phi}(u) := \{\eta \in \mathbb{R}^p \,|\, \forall v \in \mathbb{R}^p : \; \tilde{f}(v) - \tilde{f}(u) \geq \eta^T(v-u)\}\,.$$

As a set-valued map, $\partial\tilde{\Phi}$ gives rise to a dynamical system in the form

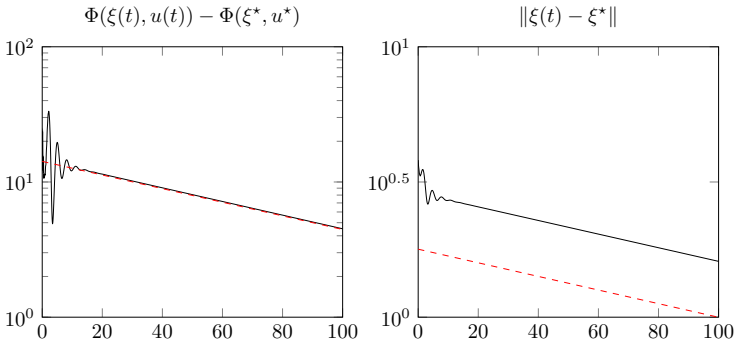

Figure 11.4: Continuous-time unconstrained feedback Newton flow with $\epsilon = \frac{\gamma\mu}{\omega L}$ (according to Corollary 11.2). Both, the reduced Newton gradient flow (11.12) (dashed) and the system (11.4) (solid) converge to the unique optimizer of (11.2) with similar convergence rate.

of a differential inclusion $\dot{u} \in -\partial\tilde{\Phi}$.

Subgradient inclusions are well-defined (i.e., existence of generalized solutions is guaranteed under technical assumptions), and convergence to critical points is also assured. However, subgradient flows are, in general, not appropriate for feedback-based optimization.

Apart from issues relating to the physical implementability, Theorem 11.1 is not applicable since Assumption 11.3 is in general not satisfied. Namely, if $\tilde{\Phi}$ is not continuously differentiable, then its gradient cannot be Lipschitz continuous.

In fact, subgradient flows in closed loop with a dynamical system are in general not asymptotically stable. To see this, consider a one-dimensional physical system in the form

$$\dot{\zeta} = -a\zeta + bu\,,$$

with $a > 0$ and steady-state map $\zeta = h(u) = \frac{b}{a}u$. Further, as an objective we consider the absolute value $\Phi(\zeta) := |\zeta|$ that gives rise to a *subgradient control law*

$$\dot{u} \in -\nabla h(u)\partial\Phi(\zeta) = -\frac{b}{a}\begin{cases} 1 & \text{if } \zeta > 0 \\ -1 & \text{if } \zeta < 0 \\ [-1,1] & \text{if } \zeta = 0 \end{cases}.$$

It is easy to see that this control law exhibits a *bang-bang* behavior that will not allow the closed-loop system to converge to the optimizer $\zeta^\star = 0$.

Figure 11.5 illustrates this behavior for a higher dimensional setup where we minimize an objective function $\Phi(\zeta, u) := \Phi(\zeta, u) + \rho\|\zeta\|_1$ with an ℓ_1-regularization term in an attempt to promote sparsity of the minimizing state variables.

11.4 Notes & Comments

The statements presented in this chapter are only a subset of the results in [Ha3] where other algorithms such as momentum methods and saddle-

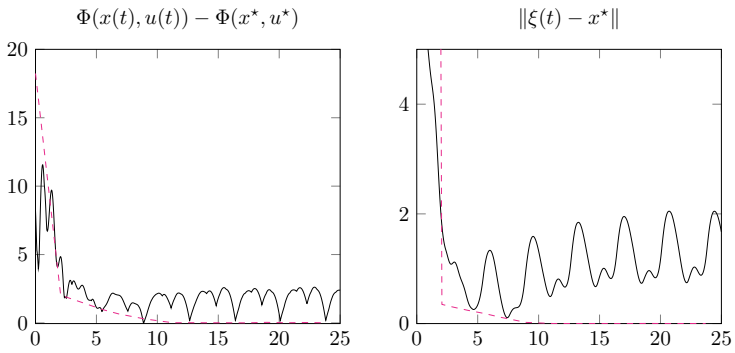

Figure 11.5: Subgradient flow induced by an ℓ_1-regularization on ζ. While the reduced dynamics (dashed) converge to $\zeta^\star$, the interconnection with a dynamical plant (solid) is not asymptotically convergent.

point flows where analyzed from a singular perturbation perspective. However, the proof of Theorem 11.1 was only sketched in [Ha3].

In contrast to LMI-based stability certificates [66, 187], the stability bounds presented in this chapter and based on singular perturbations, may be more conservative. However, they apply to a broader range of optimization algorithms (including non-convex problems), nonlinear plants, and non-smooth controllers (such as projected gradient flows). Moreover, these bounds guarantee the stability for small enough control gains, thus providing guidance for control design.

CHAPTER 12

Tracking for Time-Varying Constrained Optimization

In this chapter, we turn our attention to time-varying optimization problems of the form

$$\text{minimize} \quad \Phi(x,t) \qquad \text{subject to} \quad x \in \mathcal{X}(t)\,, \tag{12.1}$$

where $\Phi : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ is continuously differentiable in x and $\mathcal{X} : \mathbb{R}_{\geq} \rightrightarrows \mathbb{R}^n$ maps every $t \geq 0$ to a closed non-empty set.

Time-varying optimization problems are clearly of high interest for feedback-based optimization schemes, where physical plants are subject to (slowly) time-varying perturbations. In this context, one generally wishes to bound the tracking performance of a running algorithm, i.e., give a bound on the distance between the actual system/algorithm state and the (instantaneous) optimizer of the underlying optimization problem at time t.

In this chapter, we ask whether projected dynamical systems can be adapted for time-varying optimization. This inquiry will lead us to consider *perturbed sweeping processes*, which we use to define *sweeping gra-*

This chapter is based on the second part of [Ha10]. The work in this chapter was performed in collaboration with I. Subotić. The main technical results are Proposition 12.1 and Theorem 12.2.

dient flows, a generalization of projected gradient flows to time-varying optimization.

Sweeping gradient flows are of conceptual importance since they can be interpreted as the continuous-time limit of an online projected gradient descent [178, 269]. Moreover, sweeping processes of this type, which essentially model a continuously evolving projection onto a time-varying set, have been considered in feedback optimization [Ha4, Ha12, 242] where constraint enforcement may rely on input saturation in physical plants.

Although sweeping processes are extremely well-studied (see [6, 90, 248, 154, 55, 184] and references therein), known formulations cannot (yet) accommodate a variable metric as introduced for PDSs in Chapter 4. This seems to remain an open problem. Moreover, existence and uniqueness results for sweeping processes are more involved because they generally require additional assumptions on the time-varying nature of the set $\mathcal{X}(t)$. For these reasons, our results in this chapter on sweeping gradient flows are relatively specific, and there is, presumably, room for major generalizations.

To establish meaningful guarantees on the asymptotic tracking performance of sweeping gradient flows, we derive a tracking result that applies to *any* continuous-time algorithm with absolutely continuous trajectories that satisfy a common monotonicity-like property. Although the same property has previously been exploited for specific algorithms, our tracking result is more general since it does not require any additional assumptions on the (running) algorithm. For instance, our performance bound can be applied to algorithms whose trajectories are Krasovskii (Chapter 4 or [118, Ch. 5]) or Filippov [100] solutions of differential inclusions.

Another key point is that this tracking result is particularly powerful in combination with the quantitative sensitivity bounds for nonlinear optimization problems presented in Appendix A, which can be used to bound the rate of change of time-varying nonlinear optimization problems.

Hence, in Section 12.1, we first establish our general tracking result

before applying it to sweeping gradient flows, which we introduce in Section 12.2. In Section 12.3 we present two computational examples, illustrating our results.

12.1 Tracking under Monotonicity

Generally, in this section, we consider a time-varying optimization problem (12.1) for which we make the following assumption:

Assumption 12.1. *The problem (12.1) admits a unique (global) optimizer for every $t \in \mathbb{R}_{\geq 0}$ and the solution map $\tilde{x}^\star : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ is ℓ-Lipschitz.*

The following proposition guarantees the bounded asymptotic behavior of the running algorithm, i.e., the distance $\|x(t) - \tilde{x}^\star(t)\|$ between the trajectories of the running algorithm $x(t)$ and instantaneous optimizer $\tilde{x}^\star(t)$ is asymptotically bounded.

Proposition 12.1. *Consider (12.1) and let Assumption 12.1 be satisfied. Furthermore, let $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ be absolutely continuous and assume that*

$$\langle \dot{x}(t), x(t) - \tilde{x}^\star(t) \rangle \leq -a\|x(t) - \tilde{x}^\star(t)\|^2 \tag{12.2}$$

holds for almost every $t \in \mathbb{R}_{\geq 0}$ where $a \in \mathbb{R}_{>0}$. Then, $\limsup_{t\to\infty} \|x(t) - \tilde{x}^\star(t)\| \leq \frac{\ell}{a}$ holds. Moreover, if $\|x(0) - \tilde{x}^\star(0)\| \leq \frac{\ell}{a}$, then $\|x(t) - \tilde{x}^\star(t)\| \leq \frac{\ell}{a}$ holds for all $t \geq 0$.

In the next section, we provide the necessary technical results and the proof of Proposition 12.1.

12.1.1 Proof of Tracking Bound

Before we prove the proposition, we derive a key technical result from Barbalat's lemma [101, 143]. Recall that the function $\Phi : S \to \mathbb{R}$ is said to be *uniformly continuous on* $S \subseteq \mathbb{R}$ if and only if for all $\epsilon > 0$ there exists $\delta > 0$ such that for all $x', x \in S$ we have

$$|x' - x| < \delta \implies |\Phi(x') - \Phi(x')| < \epsilon\,.$$

Hence, Barbalat's lemma [143, Lemma 8.2] reads as follows:

Lemma 12.1 (Barbalat)**.** *Let $W : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be uniformly continuous. If $\lim_{t\to\infty} \int_0^t W(\tau)d\tau$ exists and is finite, then $\lim_{t\to\infty} W(t) = 0$.*

Moreover, we also require the following lemma:

Lemma 12.2. *If $W : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is continuous, lower bounded, and non-increasing, then W is uniformly continuous.*

Proof. Let $L = \inf_{t\geq 0} W(t)$ and consider any $\epsilon > 0$. Let $T > 0$ be such that $W(t) \leq L + \epsilon$ for all $t > T$. Since a continuous function is absolutely continuous on a compact interval, W is uniformly continuous on the compact interval $[0, T']$ with $T' > T$. In particular, for the given ϵ, there exists $\delta > 0$ such that $|t-\tau| \leq \delta$ implies $|W(t)-W(\tau)| \leq \epsilon$ for any $t, \tau \in [0, T']$. Without loss of generality, let $\delta < T' - T$. Next, note that for any $t, \tau \in [T, \infty)$ (and, in particular, for $|t - \tau| \leq \delta$), we have that $|W(t) - W(\tau)| \leq \epsilon$. Therefore, for any $t, \tau \geq 0$ with $|t - \tau| \leq \delta$ we have $|W(t) - W(\tau)| \leq \epsilon$. Since ϵ is arbitrary, W is uniformly continuous. □

Finally, we can present the key technical result necessary for the proof of Proposition 12.1:

Lemma 12.3. *Let $y : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ be absolutely continuous such that for almost every $t \geq 0$, some $a > 0$, and $b \geq 0$ it holds*

$$\langle \dot{y}(t), y(t) \rangle \leq -a\|y(t)\|^2 + b\|y(t)\| \,. \tag{12.3}$$

Then, $\limsup_{t\to\infty} \|y(t)\| \leq \frac{b}{a}$ holds. Moreover, if $\|y(T)\| \leq \frac{b}{a}$ holds for some $T \geq 0$, then $\|y(t)\| \leq \frac{b}{a}$ holds for all $t \geq T$.

Proof. First we note that the absolutely continuous function on a compact interval has a bounded variation and hence, it is differentiable almost everywhere. Moreover, we define $V : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ as

$$V(t) := \max\left\{\tfrac{1}{2}\left(\|y(t)\|^2 - (\tfrac{b}{a})^2\right), 0\right\} \,,$$

and let $W : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be given as

$$W(t) := \max\left\{a\|y(t)\|\left(\|y(t)\| - \tfrac{b}{a}\right), 0\right\} \,.$$

Note that both V and W are absolutely continuous. Furthermore, $V(t) > 0$ and $W(t) > 0$ holds for all t for which $\|y(t)\| > \frac{b}{a}$ and $V(t) = W(t) = 0$ otherwise. Moreover, for almost all $t \geq 0$ for which $\|y(t)\| \geq \frac{b}{a}$ we have

$$\dot{V}(t) = \langle \dot{y}(t), y(t) \rangle \leq -a\|y(t)\|^2 + b\|y(t)\| = -W(t) \leq 0\,.$$

Otherwise, for almost all t for which $\|y(t)\| < \frac{b}{a}$, we have $\dot{V}(t) = W(t) = 0$. Hence, the growth of V is bounded by $\dot{V}(t) \leq -W(t) \leq 0$ for almost all $t \geq 0$.

To prove the lemma, we apply Lemma 12.1 to W. Hence, we need to show that *(i)* the integral $\lim_{t\to\infty} \int_0^t W(\tau)d\tau$ exists and is finite, and that *(ii)* W is uniformly continuous.

To show *(i)*, we exploit the fact that the growth of V is bounded by $-W$ to establish the bound

$$\int_0^t W(\tau)d\tau \leq -\int_0^t \dot{V}(\tau)d\tau = V(0) - V(t) \leq V(0)\,,$$

where the last inequality follows from the fact that $V(t)$ is non-negative for all $t \in \mathbb{R}_{\geq 0}$, by definition. Further, the left-hand side integral is non-decreasing in t (since $W(t) \geq 0$ for all t). Hence, as $t \to \infty$, the limit exists and is finite.

For *(ii)* it suffices to show that W is non-increasing itself. Then, since W is continuous and lower bounded, it follows from Lemma 12.2 that W is uniformly continuous.

For this purpose, simply note that, for almost all t for which $\|y(t)\| \leq b/a$, we have $\dot{W}(t) = 0$ and, for almost all t for which $\|y(t)\| > b/a$, it holds that

$$\dot{W}(t) = \underbrace{\left(2a - \tfrac{b}{\|y(t)\|}\right)}_{\geq a > 0} \underbrace{\langle \dot{y}(t), y(t) \rangle}_{= \dot{V}(t) < 0} < 0\,.$$

Hence, we have shown that W is uniformly continuous and that $\lim_{t\to\infty} \int_0^t W(\tau)d\tau$ exists. Therefore, from Lemma 12.1, it follows that $W(t) \to 0$ for $t \to \infty$ and consequently, from the definition of W, we have $\limsup_{t\to\infty} \|y(t)\| \leq b/a$.

We show the last part of the lemma by contradiction. Assume that there exists $T \geq 0$ such that $\|y(T)\| \leq \frac{b}{a}$ and that there exists $t_1 \geq T$ such that $\|y(t_1)\| > \frac{b}{a}$. That means $V(t_1) > V(T) = 0$, but $\dot{V}(T) = 0$, hence $V(T+\delta) = 0$, for all $\delta \geq 0$. Consequently $V(t_1) = 0$ which is contradiction to $V(t_1) > V(T)$, hence $\|y(t)\| \leq \frac{b}{a}$, for all $t \geq T$. □

Note that the second part of Lemma 12.3, that guarantees the invariance of $\frac{b}{a}\mathbb{B}$, is essentially a consequence of Nagumo's theorem [18, Thm. 1.2.1]. Likewise, Lemma 12.3 and its proof relate to Lyapunov functions, barrier functions, etc.

Lastly, to show that Proposition 12.1 follows from Lemma 12.3, we first note x and $\tilde{x}^\star$ are both (individually) differentiable for almost all $t \geq 0$ and thus they are (jointly) differentiable for almost all $t \geq 0$ (since the union of two zero-measure sets has measure zero). Hence, the instantaneous optimizer trajectory $\tilde{x}^\star$ satisfies Assumption 12.1 for almost all $t \geq 0$, and therefore, by Cauchy-Schwarz, it holds that $|\langle \dot{\tilde{x}}^\star(t), x(t) - \tilde{x}^\star(t)\rangle| \leq \ell \|x(t) - \tilde{x}^\star(t)\|$ for almost all $t \geq 0$. Using (12.2) and letting $y(t) := x(t) - \tilde{x}^\star(t)$ (which is absolutely continuous) and $b := \ell$, the inequality (12.3) is satisfied and the proof of Proposition 12.1 follows immediately.

The *strong monotonicity* property (12.2) is, for example, satisfied with $a = \alpha$ for an unconstrained gradient flow $\dot{x} = -\nabla_x \Phi(x,t)$ if Φ is α-strongly convex for all $t \geq 0$, which is further detailed in Section 12.3.1.

12.2 Sweeping Gradient Flows

12.2.1 Perturbed Sweeping Processes

Sweeping processes [248, 154, 184] have originally been formulated to describe the sweeping effect of a moving, impenetrable boundary on a mobile object. When the object itself is subject to a perturbation (e.g., a force), we can consider a *perturbed sweeping process* [90, 55] which is formally defined as

$$\dot{x}(t) \in f(x,t) - N_x \mathcal{X}(t), \text{ with } x(0) \in \mathcal{X}(0)\,, \tag{12.4}$$

where $f : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$ is a time-varying vector field, and, for every $t \geq 0$, $\mathcal{X}(t) \subset \mathbb{R}^n$ is a non-empty closed, convex set. The idea behind (12.4) is that x evolves according to f in the interior of $\mathcal{X}(t)$ (where $N_x\mathcal{X}(t) = \{0\}$), but when x touches the boundary, a normal "force" is exerted on x to keep the state feasible, even as set $\mathcal{X}(t)$ varies.

A *solution of* (12.4) is an absolutely continuous map $x : [0, T] \to \mathbb{R}^n$ for some $T > 0$ such that, for almost all $t \in [0, T]$ (12.4) is satisfied, and $x(t) \in \mathcal{X}(t)$ holds for all $t \in [0, T]$. A *complete solution* of (12.4) in an absolutely continuous map $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ such that the restriction to any compact interval $[0, T]$ for $T > 0$ is a solution of (12.4).

The following theorem is based on [90, Thm. 1] and simplified[1] for our specific purposes.

Theorem 12.1. *For some $T > 0$, let $\mathcal{X} : [0, T] \rightrightarrows \mathbb{R}^n$ satisfy*

(i) $\mathcal{X}(t)$ is nonempty, closed and convex for all $t \in [0, T]$,

(ii) $\ell_{\mathcal{X}} > 0$ exists such that for any $z \in \mathbb{R}^n$ and $t', t \in [0, T]$,

$$|d_{\mathcal{X}(t')}(z)) - d_{\mathcal{X}(t)}(z))| \leq \ell_{\mathcal{X}}|t' - t|\,. \tag{12.5}$$

Further, let $f : [0, T] \times \mathbb{R}^n \to \mathbb{R}^n$ be measurable in t and

(iii) there exists an integrable function $\gamma : [0, T] \to \mathbb{R}_{\geq 0}$ such that

$$\|f(x, t) - f(y, t)\| \leq \gamma(t)\|x - y\|$$

holds for all $t \in [0, T]$ and any $x, y \in \mathbb{R}^n$.

Then, for any initial condition $x(0) \in \mathcal{X}(0)$, the perturbed sweeping process

$$\dot{x} \in f(x, t) - N_x\mathcal{X}(t) \quad x \in \mathcal{X}(t)$$

admits a unique solution $x : [0, T] \to \mathbb{R}^n$.

[1] The original result applies to prox-regular sets in a Hilbert space instead of convex sets in $\mathbb{R}^n$. Furthermore, in the original theorem (ii) is more general and requires $\mathcal{X}$ to vary in a *absolutely continuous* way instead of a *Lipschitz continuous* way. Finally, in the original statement (iii) is split into a *local Lipschitz* and *linear growth* condition. For simplicity, we assume global Lipschitz continuity to meet both requirements.

12.2.2 Sweeping Gradient Flows

As a new running algorithm for constrained time-varying optimization we consider a perturbed sweeping process of the form (12.4) where $f(x,t) := -\nabla_x \Phi(x,t)^T$ is the negative gradient of a time-varying cost function $\Phi : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$, thus the resulting *sweeping gradient flow* is given by

$$\dot{x}(t) \in -\nabla_x \Phi(x,t)^T - N_x \mathcal{X}(t)\,, \qquad x(0) \in \mathcal{X}(0)\,. \tag{12.6}$$

In the following, in order to be able to provide an explicit bound on the asymptotic tracking performance, we consider sweeping gradient flows for the special problem structure Corollary A.3 in Appendix A for which we can give easy-to-interpret results. However, (12.6) is well-defined for much more general setups.[2]

The key insight from this section is that our sensitivity bounds and the generalized tracking results of the previous sections can be applied to non-trivial discontinuous optimization dynamics involving time-varying constraints. Namely, we establish the following result:

Theorem 12.2. *Consider the problem*

$$\begin{aligned} \text{minimize} \quad & \hat{\Phi}(x - c(t)) \\ \text{subject to} \quad & Ux \leq v(t)\,, \end{aligned} \tag{12.7}$$

where $t \geq 0$. Define $\hat{\mathcal{X}}(t) := \{x \,|\, Ux \leq v(t)\}$ and let

(i) $\hat{\Phi} : \mathbb{R}^n \to \mathbb{R}$ be twice continuously differentiable, α-strongly convex, and $\nabla\hat{\Phi}$ is β-Lipschitz,

(ii) $c : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ and $v : \mathbb{R}_{\geq 0} \to \mathbb{R}^m$ be ℓ_c- and ℓ_v-Lipschitz continuous,

[2] For instance, existence of trajectories for sweeping gradient flows is guaranteed even if $\mathcal{X}$ is not convex, but *prox-regular*, as long as $\nabla_x \Phi$ is measurable in t, $\mathcal{X}(t)$ is non-empty for all $t \geq 0$ and, roughly-speaking, $\mathcal{X}$ varies in an absolutely continuous way [90]. To guarantee a global asymptotic bound on the tracking error, however, we generally require Assumptions A.2, A.3, A.4, and A.5 to hold such that Corollary A.2 and Proposition 12.1 can be applied.

(iii) $U \in \mathbb{R}^{n\times m}$ *and* $\omega > 0$ *such that for every* $\xi \in \Xi$ *and ever* $x \in \hat{\mathcal{X}}(\xi)$ *one has* $\omega^2 \mathbb{I} \preceq U_{\mathbf{I}_x^\xi} U_{\mathbf{I}_x^\xi}^T$.

If $\hat{\mathcal{X}}(t) := \{x \,|\, Ux \leq v(t)\} \neq \emptyset$ *for all* $t \geq 0$, *then the sweeping gradient flow*

$$\dot{x} \in -\nabla\hat{\Phi}(x - c(t))^T - N_x\hat{\mathcal{X}}(t) \tag{12.8}$$

admits a complete solution $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ *for every initial condition* $x(0) \in \hat{\mathcal{X}}(0)$ *and it holds that*

$$\limsup_{t\to\infty} \|x(t) - x^\star(t)\| \leq \frac{\ell_t}{\alpha} := \frac{\beta^{1/2}}{\alpha^{3/2}}\left(\frac{\beta\ell_c}{\alpha} + \frac{\ell_v}{\omega}\right), \tag{12.9}$$

where $x^\star(t)$ *is the unique optimizer of* (12.7) *at time* t.

Furthermore, if $\|x(0) - x^\star(0)\| \leq \frac{\ell_t}{\alpha}$, *then* $\|x(t) - x^\star(t)\| \leq \frac{\ell_t}{\alpha}$ *holds for all* $t \geq 0$.

To show Theorem 12.2 we first need to prove the existence of complete solutions for (12.6).

Lemma 12.4. *Given the setup of Theorem 12.2,* (12.8) *admits a unique complete solution for any initial point* $x(0) \in \hat{\mathcal{X}}(0)$.

Even though Lemma 12.4 is an existence result, we can prove it using the sensitivity results in Appendix A.

Proof. The proof is an application of Theorem 12.1. In order to apply Theorem 12.1 to (12.8), note that for the setup in Theorem 12.2 the requirements (i) in Theorem 12.1 (non-empty closed convex $\hat{\mathcal{X}}(t)$) and (iii) (Lipschitz vector field) hold by assumption on any compact interval $[0, T] \subset \mathbb{R}_{\geq 0}$. Also, $\nabla\Phi(x,t)^T = \nabla\hat{\Phi}(x - c(t))^T$ is measurable in t since it is continuous in t. It remains to show (ii) by establishing (12.5).

For this purpose, we consider $d(z, \mathcal{X}(t))$ as the solution of the parametrized optimization problem

$$\begin{aligned} \underset{y}{\text{minimize}} \quad & \tfrac{1}{2}\|y - z\|_2^2, \\ \text{subject to} \quad & Uy \leq v(t), \end{aligned} \tag{12.10}$$

where t varies but z is fixed. As a problem parametrized in t, (12.10) falls into the class of problems to which Corollary A.3 applies. Namely, we have $\alpha = \beta = 1$ (since $\hat{\Phi}(x) := \frac{1}{2}\|x\|^2$), $\ell_c = 0$ (since $c \equiv 0$), and ω as defined in Theorem 12.2. Further, by the assumption in Theorem 12.2, the feasible set of (12.10) (which is equivalent to $\hat{\mathcal{X}}(t)$) is non-empty for all $t \geq 0$.

Therefore, by Corollary A.3, the solution map $t \mapsto y_z^\star(t)$ of (12.10) is Lipschitz continuous with a bound on the Lipschitz constant that is independent of z and given by $\ell_{y^\star} = \ell_v/\omega$. Therefore, we have $|y_z^\star(t') - y_z^\star(t)| \leq \ell_{y^\star}|t' - t|$ for any $z \in \mathbb{R}^n$ and $t, t' \geq 0$, and (ii) in Theorem 12.1 holds on any compact interval $[0, T]$.

Hence, Theorem 12.1 guarantees the existence of a unique solution of (12.8) for every initial condition $x(0) \in \hat{\mathcal{X}}(0)$, and for every compact interval $[0, T]$ and hence, by definition, a complete solution on $\mathbb{R}_{\geq 0}$. □

Proof of Theorem 12.2. By Corollary A.3, (12.7) admits a unique primal solution $x^\star(t)$ for every $t \geq 0$. In particular, the solution map $t \mapsto x^\star(t)$ of (12.7) is ℓ_t-Lipschitz with $\ell_t := \sqrt{\frac{\beta}{\alpha}}\left(\frac{\beta\ell_c}{\alpha} + \frac{\ell_v}{\omega}\right)$. Consequently, Assumption 12.1 is satisfied. Furthermore, Lemma 12.4 guarantees the existence of a unique complete solution $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ of (12.8) for every initial condition $x(0) \in \hat{\mathcal{X}}(0)$. Next, we verify that (12.2) holds, i.e., we show that $\langle \dot{x}(t), x(t) - x^\star(t)\rangle \leq -\alpha\|x(t) - x^\star(t)\|^2$.

In the following let $\Phi(x, t) := \hat{\Phi}(x - c(t))$. Recall that, for almost all $t \geq 0$, we have $\dot{x}(t) = -\nabla_x \Phi(x(t), t)^T - \eta(t)$ for some $\eta(t) \in N_{x(t)}\hat{\mathcal{X}}(t)$. Further, because $x^\star(t)$ solves (12.7), it satisfies the KKT conditions which is equivalent to saying that $-\nabla_x \Phi(x^\star(t), t)^T - \eta^\star(t) = 0$ for some $\eta^\star(t) \in N_{x^\star(t)}\hat{\mathcal{X}}(t)$. Putting these two insights together, we have that

$$\dot{x}(t) = -\nabla_x \Phi(x(t), t)^T - \eta(t) + \nabla_x \Phi(x^\star(t), t)^T + \eta^\star(t)\,.$$

Next, we omit the argument t for x, $\tilde{x}^\star$, η, $\eta^\star$ and $\hat{\mathcal{X}}$. By definition of the normal vectors $\eta \in N_x\hat{\mathcal{X}}$ and $\eta^\star \in N_{x^\star}\hat{\mathcal{X}}$, for $\hat{\mathcal{X}}$, we have $\langle \eta, x - x^\star\rangle \geq 0$ and $\langle \eta^\star, x^\star - x\rangle \geq 0$, respectively. Since Φ is α-strongly convex we have

$$\left\langle \nabla\Phi(x, t)^T - \nabla\Phi(x^\star, t)^T, (x - x^\star)\right\rangle \geq \alpha\|x - x^\star\|^2\,. \tag{12.11}$$

Combining these facts, we get

$$\begin{aligned}\langle \dot{x}, x-x^\star\rangle &= \big\langle -(\nabla\Phi(x,t)^T - \nabla\Phi(x^\star,t)^T) - (\eta-\eta^\star), x-x^\star\big\rangle \\ &= -\langle \eta, x-x^\star\rangle - \langle \eta^\star, x^\star - x\rangle \\ &\quad -\big\langle \nabla\Phi(x,t)^T - \nabla\Phi(x^\star,t)^T, x-x^\star\big\rangle \\ &\leq -\alpha\|x-x^\star\|^2 .\end{aligned}$$

Consequently, by taking $a=\alpha$, (12.2) holds. Thus, Proposition 12.1 is applicable and yields the desired asymptotic tracking bound and completes the proof. □

In the forthcoming Section 12.2, we require the generality of this result for our novel discontinuous sweeping gradient flow.

12.3 Numerical Examples

In this section, we provide two numerical examples to illustrate our results for simple time-varying optimization setups. First, we consider an unconstrained optimization problem in one dimension, with non-smooth change in time to demonstrate our bound's tightness in that case. Second, we consider a constrained time-varying optimization problem to illustrate the behavior of sweeping gradient flows.

12.3.1 Time-Varying Unconstrained Optimization

Consider the problem of minimizing $\Phi(x,\tilde{\xi}(t)) = \|x-\tilde{\xi}(t)\|^2$, where $\tilde{\xi}$ is a triangular wave (hence non-smooth) with period $\tau = 4$s. That is, for all $n\in\mathbb{N}$, $\tilde{\xi}$ is defined as

$$\tilde{\xi}(t) = \begin{cases} (t-\tau(n-1))-1, & t\in[\tau(n-1),\tau(n-0.5)) \\ -(t-\tau(n-1))+3, & t\in[\tau(n-0.5),\tau n)\end{cases}.$$

In particular, in reference to (A.16) we have $\psi = t$ and $\Phi(x,\xi) := \|x-\xi\|^2$ is α-strongly convex with $\alpha = 2$. Furthermore, the solution map given by $\tilde{x}^\star(t) = \tilde{\xi}(t)$ is $\ell_{\tilde{\xi}}$-Lipschitz with $\ell_{\tilde{\xi}} = 1$. From (A.11) we have

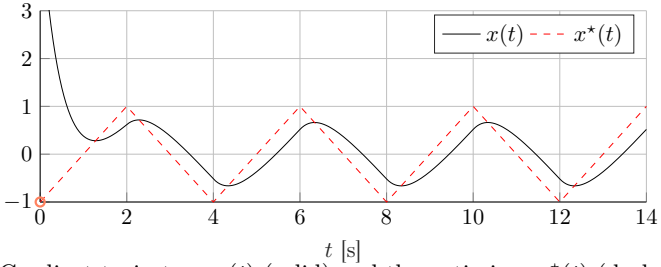

(a) Gradient trajectory $x(t)$ (solid) and the optimizer $x^\star(t)$ (dashed).

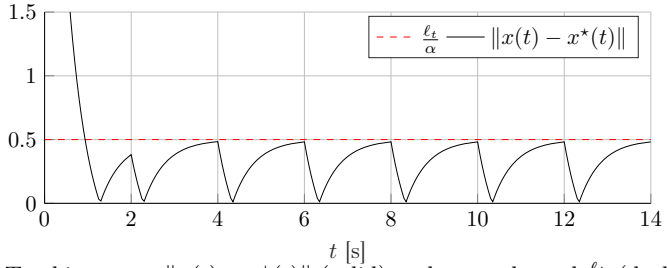

(b) Tracking error $\|x(t) - x^\star(t)\|$ (solid) and upper bound $\frac{\ell_t}{\alpha}$ (dashed).

Figure 12.1: Gradient flow for time-varying unconstrained optimization

$\sup_{\xi\in[-1,1]}\|\nabla^2_{\xi x}\Phi\| = 2$, the solution trajectory $x^\star$ is $\ell_{x^\star}$-Lipschitz with respect to ξ, where $\ell_{x^\star} = 1$. Therefore, the solution trajectory $\tilde{x}^\star$ is ℓ_t-Lipschitz continuous in t with $\ell_t := \ell_{x^\star}\ell_{\tilde{\xi}} = 1$.

For all $t \geq 0$, the unconstrained gradient flow is given by

$$\dot{x}(t) = -\nabla_x \Phi(x(t)^T, \tilde{\xi}(t)) = -2(x(t) - \tilde{\xi}(t))\,.$$

Finally, combining the zero cost function gradient, at the unconstrained optimizer (i.e., $\nabla_x \Phi(\tilde{x}^\star(t), \tilde{\xi}(t))^T = 0$) with the α-strongly convex property (12.11), for almost all $t \geq 0$ we have

$$\langle \dot{x}, x(t) - \tilde{x}^\star(t)\rangle \leq -\alpha\|x(t) - \tilde{x}^\star(t)\|^2\,.$$

Hence, from Proposition 12.1, $\limsup_{t\to\infty}\|x(t) - \tilde{x}^\star(t)\| \leq 0.5$.

Figure 12.1a shows the trajectory $x(t)$ (solid) obtained from the gradient descent algorithm with initial condition $x(0) = 5$ and the instantaneous optimizer trajectory $\tilde{x}^\star(t)$ (dashed).

Figure 12.1b shows the distance between the trajectories $\|x(t) - \tilde{x}^\star(t)\|$ (solid), and the analytical sensitivity bound $\frac{\ell_t}{\alpha} = 0.5$ (dashed). Easily, it can be noticed that as $t \to \infty$, the distance always stays within the analytical sensitivity bound. Additionally, in this example, the analytical sensitivity bound is indeed tight. Moreover, this example perfectly illustrates the invariance property of a ball of radius $\frac{\ell_t}{\alpha}$ around the time-varying unconstrained optimizer $\tilde{x}^\star(t)$. Namely, the algorithm trajectory never perfectly tracks the unconstrained optimizer trajectory, yet always stays within $\frac{\ell_t}{\alpha}$ from the optimizer.

12.3.2 Time-Varying Constrained Optimization

In this example, we consider the same problem setup as in Theorem 12.2 with smooth changes in time. Namely, let

$$Q := \begin{bmatrix} 12 & -8 \\ -8 & 10 \end{bmatrix} \succ 0, \; P(t) := -\begin{bmatrix} 3\sin(t+3) + 1.3t \\ 2\tanh(t-3) + 0.71t \end{bmatrix},$$

let $c(t) := -\frac{1}{2}Q^{-1}P(t)$, and the cost function is defined as

$$\hat{\Phi}(x(t) - c(t)) = (x(t) + \tfrac{1}{2}Q^{-1}P(t))^T Q(x(t) + \tfrac{1}{2}Q^{-1}P(t)).$$

In particular, $\hat{\Phi}$ is α-strongly convex with $\alpha = 2\sigma^{\min}(Q)$, and gradient $\nabla\hat{\Phi}^T$ is β-Lipschitz with $\beta = 2\sigma^{\max}(Q)$. Furthermore, we can compute $\ell_c = \sup_{t\geq 0}\|\dot{P}(t)\|$.

Next, the constraint set is given by

$$\hat{\mathcal{X}}(t) = \{x(t) \in \mathbb{R}^n \mid Ux(t) \leq v(t)\},$$

where $v(t) := (V_1 t + V_2)$ and

$$U := \begin{bmatrix} -2 & 1 \\ 1 & -1 \\ 0.5 & 1 \\ -3 & -1 \end{bmatrix}, \quad V_1 := \begin{bmatrix} -0.05 \\ -0.3 \\ 0.25 \\ -0.5 \end{bmatrix}, \quad V_2 := \begin{bmatrix} -2 \\ 5 \\ 4 \\ 3 \end{bmatrix}.$$

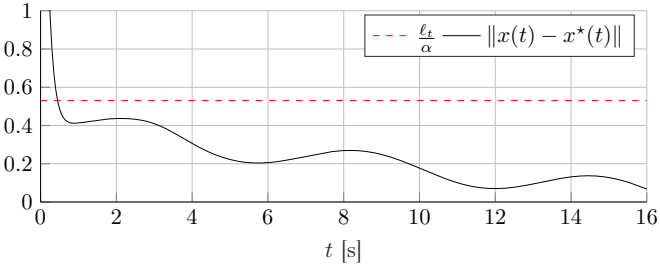

(a) Tracking error $\|x(t) - x^\star(t)\|$ (solid), and upper bound $\frac{\ell_t}{\alpha}$ (dashed).

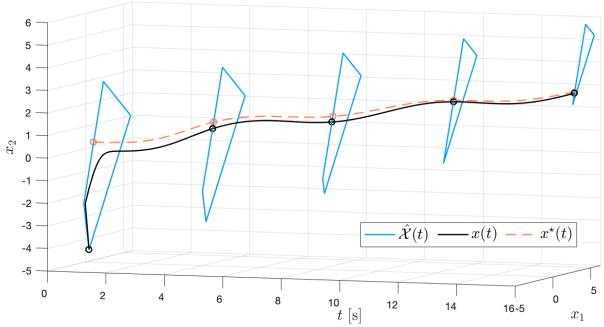

(b) Sweeping gradient trajectory $x(t)$ (solid), instantaneous optimizer trajectory $x^\star(t)$ (dashed), and time-varying feasible set.

Figure 12.2: Sweeping gradient flow for constrained problem

In particular, we have $\ell_v = \sigma^{\max}(V_1)$ and in this special case, we can determine ω by inspection.

Hence, we can compute the bound on the Lipschitz constant of the constrained optimizer $\tilde{x}^\star$ as $\ell_t = \sqrt{\frac{\beta}{\alpha}}\left(\frac{\ell_p}{\alpha} + \frac{\ell_C}{\omega}\right)$, and we can compute the analytical tracking bound as $\frac{\ell_t}{\alpha} \approx 0.5302$.

Figure 12.2a presents the tracking error $\|x(t) - \tilde{x}^\star(t)\|$ (solid) and the tracking error upper bound $\frac{\ell_t}{\alpha} \approx 0.5302$ (dashed). It can be seen that the tracking error is always within the bound and that it is decreasing within the considered time interval.

In contrast to the unconstrained time-varying optimization, for the

constrained time-varying optimization, the movement of both the cost function and the constraint set has to be considered. Figure 12.2b illustrates the sweeping gradient flow trajectory $x(t)$ (solid), instantaneous optimizer trajectory $x^\star(t)$ (dashed) and the time-varying constraint set $\hat{\mathcal{X}}(t)$ (blue, solid), captured in a different time steps, i.e., $t = 0$s, $t = 3.95$s, $t = 7.95$s, $t = 11.95$s, and $t = 16$s. Note, however, that contrary to the theoretical requirement, the problem becomes eventually infeasible with $\hat{\mathcal{X}}(t)$ collapsing to an empty set.

12.4 Conclusions

In this chapter, we have illustrated how projected gradient flows have a natural generalization to time-varying problems in terms of perturbed sweeping processes. However, we do not (yet) reach the same level of generality as in the previous chapters.

Namely, using a variable metric to induce oblique projections for sweeping processes has not been studied, thus limiting their application to Euclidean setups.

Although, the tracking bound Proposition 12.1 exploits the monotonicity found in strongly convex problems, it is more general than previous work [223, 207, 222, 201] because it applies to general continuous-time systems, including constrained differential inclusions such as sweeping gradient flows.

Finally, the tracking guarantees in this chapter, as shown in Theorem 12.2, work very well in combination with the sensitivity results in Appendix A which can be used to express the rate of change of an optimizer in terms of cost and constraint variations, rather than assuming an a-priori bound.

Part III

Feedback-Based Online Optimization for Power System Operations

CHAPTER 13

Nonlinear Power System Modeling

In this third part of the thesis, we turn our attention to the feedback-based optimization of power systems. Our emphasis is on establishing a strong connection between the abstract theoretical results of the previous parts and their application to the power systems context. Further, we illustrate the design opportunities and pitfalls for feedback-based optimization schemes for power system operations, we provide simulation results to showcase both characteristic as well as pathological behavior, and finally we discuss some potential use cases in the conclusions of the thesis. First, however, in this chapter, we review modeling basics of power systems at steady state.

13.1 Notational Conventions

For the remaining chapters, we adopt a slightly different notation to conform to conventions in power systems analysis, yet retain some consistency with the notation in Parts I and II. In particular, we continue to denote the inputs, and outputs of an (abstract) physical plant by u and y, respectively, and use x as a generic variable (often denoting the power system steady state) where necessary.

Otherwise, calligraphic letters like $\mathcal{M}, \mathcal{C}$, etc. continue to denote (uncountable) sets (e.g., subsets of $\mathbb{R}^n$) and bold capital letters $\mathbf{V}, \mathbf{E}$ denote (finite) index sets or matrices, like the admittance matrix $\mathbf{Y}$.

In the following $j := \sqrt{-1}$ is reserved for the imaginary unit, e denotes Euler's number, and $(\cdot)^*$ denotes the complex conjugate. Operations such as $\sin(\cdot), \cos(\cdot), \mathrm{e}^{(\cdot)}, |\cdot|$, etc. apply component-wise.

Constant physical quantities such as impedances, admittances, susceptance, turn ratios, etc. are denoted by roman (non-italic) symbols, e.g., $\mathrm{z}, \mathrm{y}, \mathrm{b}, \mathrm{n}$, etc.

Most electrical quantities are in *per unit* [p.u.], all angles are in radians [rad], and cost is measured in [$/h], even if it does not necessarily express an economic quantity.

13.2 AC Power Flow

An electricity grid is modeled as a connected directed graph $(\mathbf{V}, \mathbf{E})$ where $\mathbf{V}$ denotes the set of vertices and $\mathbf{E}$ denotes the set of edges with cardinality $N = |\mathbf{V}|$ and $M = |\mathbf{E}|$, respectively. Nodes are denoted by integers from $1 \leq k \leq N$ and edges are denoted by tuples $km := (k, m)$ describing an edge from node k to node m. Furthermore, $\mathbf{E}^{\mathrm{f}}(k)$ and $\mathbf{E}^{\mathrm{t}}(k)$ denote the outgoing and incoming edges of node k, respectively. The neighbors of node k, i.e., the nodes connected to k by either an incoming or outgoing line, are denoted by $\mathbf{N}(k)$.

We consider only balanced power flow conditions and therefore reduce the grid to its single-phase equivalent. Furthermore, unless noted otherwise, we assume that $(\mathbf{V}, \mathbf{E})$ is connected (in the sense that any two nodes can be connected by an undirected path). Although parallel transmission lines are common in reality, we assume, for simplicity, that any two nodes are directly connected by at most one line.

Hence, each node $k \in \mathbf{V}$ has an associated voltage phasor $e_k := v_k \angle\theta \in \mathbb{C}$ where $v_k \geq 0$ and $\theta_k \in [-\pi, \pi)$ denote the voltage magnitude. Furthermore, at every node k, the apparent power $s_k = p_k + jq_k \in \mathbb{C}$ is injected into the grid, where $p_k, q_k \in \mathbb{R}$ denote the real and reactive power components, respectively.

Each edge $km \in \mathbf{E}$ from node k to m has an associated *from-current* $i^{\mathrm{f}}_{km} \in \mathbb{C}$ denoting the current flowing out of node k into the line, and

a *to-current* $i^{\mathrm{t}}_{km} \in \mathbb{C}$ which is the current flowing out of node m into the line. Note that, in general, $i^{\mathrm{f}}_{km} \neq -i^{\mathrm{t}}_{km}$ (unless line shunts and tap-changing transformers can be ignored, as explained in Section 13.2.1 below).

Throughout, $e, s \in \mathbb{C}^N$, $i^{\mathrm{t}}, i^{\mathrm{f}} \in \mathbb{C}^M$, and $v, \theta, p, q \in \mathbb{R}^N$, etc. denote vectors obtained from stacking the respective variables and parameters.

The flow of power and current through in the electricity grid at steady state can be captured entirely by the *bus admittance matrix* $\mathbf{Y} \in \mathbb{C}^{N \times N}$ and the *line admittance matrices* $\mathbf{Y}^{\mathrm{f}}, \mathbf{Y}^{\mathrm{t}} \in \mathbb{C}^{M \times N}$. Namely, the quantities $s, p, q, e, v, \theta, i^{\mathrm{f}}, i^{\mathrm{t}}$ are related by the algebraic equations

$$s = p + jq = \mathrm{diag}(e)(\mathbf{Y}e)^* \qquad i^{\mathrm{f}} = \mathbf{Y}^{\mathrm{f}} e \qquad i^{\mathrm{t}} = \mathbf{Y}^{\mathrm{t}} e\,. \tag{13.1}$$

We refer to (13.1) as the *AC power flow* (ACPF) equations. The following subsection details the construction of these admittances matrices from the parameters of power transmission lines.

13.2.1 Line Models & Admittance Matrices

We model tap-changing and phase-shifting transformers and power transmission lines using a unified branch model based on a Π-equivalent circuit. We follow the modeling and the conventions used in MATPOWER [268] since this software package was used for most simulations in this thesis. The reader is also referred to [181, Chap. 2] for alternative representations of the same model.

Every branch $km \in \mathbf{E}$ is modeled, as illustrated in Figure 13.1, by a *line admittance* y_{km} which can be decomposed into a *line conductance* g_{km} and a *line susceptance* b_{km}. Equivalently, the line admittance can be expressed as the inverse of the *line impedance* which can, in turn, be decomposed into the *line resistance* r_{km} and the *line reactance* x_{km}. A *charging susceptance* $\mathrm{b}^{\mathrm{sh}}_{km}$ is split between both branch ends.

Furthermore, an ideal phase-shifting transformer with tap ratio $\mathrm{n}_{km} > 0$ and phase shift $\theta^{\mathrm{shift}}_{km} \in [-\pi, \pi)$ is located at the from-bus, i.e., at bus k. If the branch models a simple power line without transformer, then $\mathrm{n}_{km} = 1$ and $\theta^{\mathrm{shift}}_{km} = 0$.

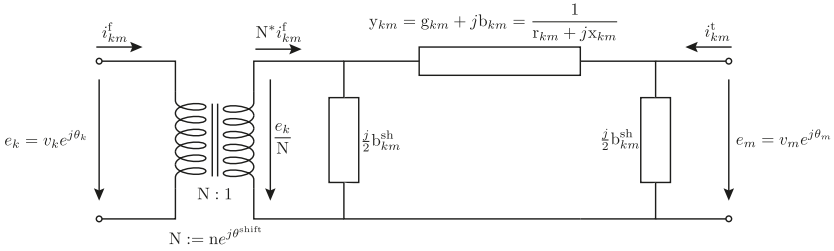

Figure 13.1: Unified branch model

Hence, the line currents for a line $km \in \mathbf{E}$ are given by

$$i^{\mathrm{f}}_{km} := \tfrac{\mathrm{y}_{km}}{\mathrm{n}^2}\left(e_k - e_m \mathrm{n} \mathrm{e}^{-j\theta^{\mathrm{shift}}}\right) + \tfrac{j}{2\mathrm{n}^2}\mathrm{b}^{\mathrm{sh}}_{km} e_k$$

and

$$i^{\mathrm{t}}_{km} := \mathrm{y}_{km}\left(e_m - e_k \tfrac{1}{\mathrm{n}} \mathrm{e}^{-j\theta^{\mathrm{shift}}}\right) + \tfrac{j}{2}\mathrm{b}^{\mathrm{sh}}_{km} e_m \,.$$

The line flow admittance matrices $\mathbf{Y}^{\mathrm{f}}$ and $\mathbf{Y}^{\mathrm{t}}$ follow directly by noting that i^{t} and i^{f} are linear in e.

Finally, we assume that every bus k has a shunt admittance $\mathrm{y}^{\mathrm{sh}}_k$ which may be used to model fixed capacitor banks for voltage support or constant impedance loads.

Given these parameters for every transmission line in the network, we may write the bus admittance matrix[1] as

$$[\mathbf{Y}]_{km} := \begin{cases} \mathrm{y}^{\mathrm{sh}}_k + \sum\limits_{kl \in \mathbf{E}^{\mathrm{f}}(k)} \frac{1}{\mathrm{n}^2_{kl}}\left(\mathrm{y}_{kl} + j\frac{\mathrm{b}^{\mathrm{sh}}_{kl}}{2}\right) + \sum\limits_{lk \in \mathbf{E}^{\mathrm{t}}(k)} \left(\mathrm{y}_{lk} + j\frac{\mathrm{b}^{\mathrm{sh}}_{lk}}{2}\right) & \text{if} \quad k = m \\ -\frac{\mathrm{y}_{km}}{\mathrm{n}_{km}} \mathrm{e}^{j\theta^{\mathrm{shift}}_{km}} & \text{if} \quad km \in \mathbf{E} \\ -\frac{\mathrm{y}_{km}}{\mathrm{n}_{km}} \mathrm{e}^{-j\theta^{\mathrm{shift}}_{km}} & \text{if} \quad mk \in \mathbf{E} \\ 0 & \text{otherwise,} \end{cases}$$

where $[\mathbf{Y}]_{km}$ denotes the component of $\mathbf{Y}$ in the k-th row and m-th column.

[1] Note that, without shunts ($\mathrm{y}^{\mathrm{sh}} = 0$, $\mathrm{b}^{\mathrm{sh}} = 0$), resistances ($\mathrm{g} = 0$) and without transformers ($\mathrm{n} = 1$, $\theta^{\mathrm{shift}} = 0$), $\mathbf{Y}$ reduces to a weighted graph Laplacian of $(\mathbf{V}, \mathbf{E})$.

13.3 Canonical AC Optimal Power Flow

Throughout the remaining chapters, we will usually consider a standard *AC optimal power flow* (ACOPF) problem

$$\begin{aligned}
\text{minimize} \quad & \Phi(v, \theta, p, q) && \text{(13.2a)}\\
\text{subject to} \quad & \operatorname{diag}(e)(\mathbf{Y}e)^* = p + jq && \text{(13.2b)}\\
& \underline{p} \le p \le \overline{p} && \text{(13.2c)}\\
& \underline{q} \le q \le \overline{q} && \text{(13.2d)}\\
& \underline{v} \le v \le \overline{v} && \text{(13.2e)}\\
& |\mathbf{Y}^{\mathrm{f}} e| \le \overline{i} \qquad |\mathbf{Y}^{\mathrm{t}} e| \le \overline{i}\,, && \text{(13.2f)}
\end{aligned}$$

where $p, q \in \mathbb{R}^N$ denote the real and reactive power injection at each bus, and $\Phi : \mathbb{R}^{4N} \to \mathbb{R}$ is a differentiable objective function.

The problem (13.2) is an abstraction and simplification of practical ACOPF problems, although widely accepted as a benchmark problem structure in academic research [181, 179, 105].

The power flow equations (13.2b), although written here as a set of $2N$ complex equations, are usually split into $4N$ real-valued equations for numerical computations.

Furthermore, (13.2c–f) denote upper and lower limits on real and reactive power injection, bus voltage magnitudes, and line currents, respectively. Fixed power loads at each bus can be modeled by offsetting the power injection constraints (13.2c–d) accordingly. More precisely, we generally assume limits of the form

$$\underline{p} := \underline{p}^{\mathrm{G}} - p^{\mathrm{L}} \quad \overline{p} := \overline{p}^{\mathrm{G}} - p^{\mathrm{L}} \quad \underline{q} := \underline{q}^{\mathrm{G}} - q^{\mathrm{L}} \quad \overline{q} := \overline{q}^{\mathrm{G}} - q^{\mathrm{L}}\,, \qquad (13.3)$$

where $p^{\mathrm{L}}, q^{\mathrm{L}}$ denote fixed power loads and $\underline{p}^{\mathrm{G}}, \underline{q}^{\mathrm{G}}, \overline{p}^{\mathrm{G}}, \overline{q}^{\mathrm{G}}$ denote limits on the generation capacity. Without loss of generality, if no generator is placed at node i, we have $\underline{p}^{\mathrm{G}} = \underline{q}^{\mathrm{G}} = \overline{p}^{\mathrm{G}} = \overline{q}^{\mathrm{G}} = 0$.

Although from a physical perspective, it is more natural to formulate constraints on the power generation rather than on power injections into the grid, the latter will be more convenient for the discussion in Chapter 14.

Note that (13.2) falls within the class of nonlinear optimization problems discussed in Section 2.2. In particular, if constraint qualifications such as LICQ hold, every (local) optimizer of (13.2) satisfies the KKT conditions and KKT points that satisfy the SSOSC are local minimizers of (13.2).

Although (13.2) is a non-convex optimization problem and generally intractable, it is still routinely solved. In particular, off-the-shelf (local) solvers often perform well [106, 107, 268] and convex relaxation techniques have recently been studied extensively to certify global optimality of solutions [169, 180, 182, 92, 157].

The requirement that constraint qualifications have to hold is often overlooked in the study of ACOPF because degenerate constraints do not seem to pose a problem in practice. This phenomenon has been rigorously explained in [Ha11] using the notion of generic transversality from differential topology. In particular, it was shown that the feasible set of (13.2) satisfies LICQ generically (i.e., for almost all problem instances) under very mild assumption.

CHAPTER 14

A Geometric Perspective on AC Power Flow

This chapter provides a geometric interpretation of the nonlinear steady-state power flow model described in the previous chapter. In particular, we discuss how power system models can be seen as embeddings or parametrizations of an abstract *power flow manifold* (PFM), which represents all physically possible steady states of a power system in the given modeling framework.

This insight then allows us to rigorously apply the results from Parts I and II in the power systems context. In particular, we discuss two projected gradient flows on the PFM, which both define a desired closed-loop behavior for feedback-based optimization schemes. We will propose numerical implementations in the subsequent chapter.

In this chapter, we also identify three metrics on the PFM, which lead to different algorithm variations to which the same theoretical results apply. Each of these metrics has been used implicitly in the design of existing algorithms, although they have not been identified as such. With a small example, we also illustrate how the choice of metric can dramatically affect the algorithms' behavior when the problem is ill-conditioned, and the grid is close to voltage collapse.

14.1 Preliminaries on Smooth Manifolds

We quickly review some important notions from differential geometry. Although we will use the following terms fairly loosely, they have rigorous definitions that can be found in standard texts on differential geometry like [162, 48].

A *smooth n-dimensional manifold* $\mathcal{M}$ is a set of points that satisfies certain technical properties and, more importantly, around every point $x \in \mathcal{M}$, $\mathcal{M}$ is *locally Euclidean*. That is, around x, $\mathcal{M}$ "looks like" $\mathbb{R}^n$. This concept is captured by the notion of a *chart*, which maps a chart domain $\mathcal{U} \subset \mathcal{M}$ to a subset of $\mathbb{R}^n$ and back (similar to a map of a geographic area which is a one-to-one representation of a small portion of the earth's surface). A set of charts that cover the entire manifold and which are "compatible" form an *atlas*. By compatibility, we mean a smooth coordinate transformation between any two chart domains that overlap.

The definition of a manifold in terms of local charts is very general and is particularly useful to extend definitions of objects that are well-defined locally on $\mathbb{R}^n$ to manifolds. In particular, this approach has been taken in Chapter 4 to define oblique projected dynamical systems on manifolds.

However, smooth manifolds are often visualized and treated as "hypersurfaces" *embedded* in a higher dimensional vector space. In this context, the most relevant result is that level sets of smooth functions with constant rank are smooth manifolds. More precisely, if $F : \mathbb{R}^m \to \mathbb{R}^n$ is continuously differentiable and $\nabla F(x)$ has rank $k \leq n$ for all x for which $F(x) = r$, then the set $\mathcal{M} := \{x \,|\, F(x) = r\}$ is a smooth (C^1) manifold of dimension $m - -k$ [48, Thm. 5.8].

Embeddings can be *local* in the sense that if $\nabla F(x)$ has rank k for all x in an open set $\mathcal{V} \subset \mathbb{R}^n$, then $\mathcal{M} := \{x \in \mathcal{V} \,|\, F(x) = r\}$ is a manifold (albeit an "open" manifold).

Embeddings also make it easy to visualize the tangent space $T_x\mathcal{M}$ of a manifold $\mathcal{M}$ at $x \in \mathcal{M}$: Namely, for $\mathcal{M} := \{x \in \mathcal{V} \,|\, F(x) = r\}$ the tangent space at x is simply given by the kernel of ∇F, i.e., $T_x\mathcal{M} =$

$\ker \nabla F(x) = \{v \mid \nabla F(x)v = 0\}$ [211, Ex. 6.8].

Concerning the topology of manifolds, we say that two manifolds are *homeomorphic* if there exists a continuous bijection, i.e., an invertible one-to-one map between the two manifolds.

14.2 The Power Flow Manifold

Given the bus admittance matrix $\mathbf{Y}$, we define the power flow manifold as the (abstract) set of all steady states of the power system that are consistent with the power flow model of the previous chapter. For instance, the following may define a point on the power flow manifold:

(i) A complex N-dimensional vector $e \in \mathbb{C}^N$ of voltage phasors: Given e, we can immediately compute real and reactive power injections $p, q \in \mathbb{R}^N$, the line currents $i^{\mathrm{f}}, i^{\mathrm{t}} \in \mathbb{C}^M$ (using $\mathbf{Y}^{\mathrm{f}}, \mathbf{Y}^{\mathrm{t}}$), as well as other quantities such as line power flows, power factors, etc.

(ii) A point $(v, \theta, p, q) \in \mathbb{R}^{4N}$ that satisfies the ACPF (13.1).

(iii) Two N-dimensional vectors $p, q \in \mathbb{R}^N$ of real and reactive power injections, if there exist vectors $v, \theta \in \mathbb{R}^N$ such that (13.1) is satisfied.

Point (i) above allows us to draw the following conclusion: The (steady) state of a power system is uniquely determined by the complex N-dimensional vector e of voltage phasors. Conversely, every vector $e \in \mathbb{C}^N$ describes a physical state of the power system at steady state. Hence, the PFM is homeomorphic to $\mathbb{C}^N$, and consequently to $\mathbb{R}^{2N}$.[1] In particular, the PFM has dimension $2N$ and *per se* does not have any holes or a boundary. The map $\mathcal{M} \to \mathbb{R}^{2N}$ that maps every point of the power flow manifold to the real and imaginary part of its voltage phasors constitute a *global* chart for the manifold.

[1] The equivalence to $\mathbb{R}^{2N}$ is stems from a Cartesian decomposition of any complex number. Polar coordinates are not suitable to establish this equivalence, because a polar representation maps $\mathcal{C}$ to $[0, \infty) \times [0, 2\pi)$ which is (1) not $\mathbb{R}^2$, and (2) not one-to-one because $0 \in \mathbb{C}$ has no unique polar representation.

Point (ii) hints at an embedding of the PFM in $\mathbb{R}^{4N}$: By defining

$$F(v,\theta,p,q) := \begin{bmatrix} \mathrm{Re}\left\{\mathrm{diag}(e)(\mathbf{Y}e)^*\right\} \\ \mathrm{Im}\left\{\mathrm{diag}(e)(\mathbf{Y}e)^*\right\} \end{bmatrix} - \begin{bmatrix} p \\ q \end{bmatrix}$$

the PFM can be represented as the level set $\{(v,\theta,p,q) \mid F(v,\theta,p,q) = 0\}$. In the forthcoming section we will study this formulation further by showing that F has full rank (and thus, indeed, defines an embedded manifold) and by deriving the expression for the tangent space.

Point (iii) above corresponds to a classical *AC power flow calculation*: Given the power injections at every bus, we need to compute the corresponding voltage magnitude and angles at every bus. This setup, however, does not correspond to a clear description of the power flow manifold:

First, if (v,θ) is a voltage phasor satisfying (13.1), then every voltage phasor $v\angle(\theta+\delta\mathbb{1})$ satisfies (13.1). This rotational symmetry is commonly eliminated by designating a slack bus to serve as an angle reference.

Second, even with a fixed angle reference, (13.1) often has multiple admissible voltage profiles for a given set of power injections. This property is well-known and simply due to the fact that, as a nonlinear system of equations in (v,θ), (13.1) may admit multiple solutions or none.

Finally, the bus admittance matrix $\mathbf{Y}$ is generally very ill-conditioned, which makes it difficult to solve (13.1) numerically. In particular, in the absence of any (nodal or line) shunts impedances and with purely inductive lines, $\mathbf{Y}$ reduces to a weighted Laplacian of the graph $(\mathbf{V},\mathbf{E})$ which is naturally rank-deficient. As a remedy, in numerical power flow simulations, the slack bus serves not only as an angle reference but also as a voltage reference with a fixed voltage. Consequently, the real and reactive power injections at the slack bus cannot be fixed in advance and need to be treated as variables for (13.1) to be a well-posed system of $2N$ equations with $2N$ variables.

14.3 A Simple PFM Embedding

As noted in the previous section, the N complex equations that make up the AC power flow equations (13.1) can be split into $2N$ real equations and written in the implicit form

$$0 = F(v,\theta,p,q) := \begin{bmatrix} \operatorname{Re}\{\operatorname{diag}(e)(\mathbf{Y}e)^*\} \\ \operatorname{Im}\{\operatorname{diag}(e)(\mathbf{Y}e)^*\} \end{bmatrix} - \begin{bmatrix} p \\ q \end{bmatrix},$$

where $F : \mathbb{R}^{4N} \to \mathbb{R}^{2N}$.

It is straightforward to see that $\nabla F = \begin{bmatrix} \nabla_{(v,\theta)} F & -\mathbb{I}_{2N} \end{bmatrix} \in \mathbb{R}^{2N \times 4M}$ has full rank $2N$, and therefore the set

$$\mathcal{M} := \{(v,\theta,p,q) \,|\, F(v,\theta,p,q) = 0\}$$

is indeed an embedding of the $2N$-dimensional PFM in $\mathbb{R}^{4N}$. Henceforth, for simplicity, we will thus refer to $\mathcal{M} \subset \mathbb{R}^{4M}$ as the PFM (although it is, strictly speaking, only an embedding of the *abstract* PFM).

The tangent space of $\mathcal{M}$ at $x \in \mathcal{M}$ is given by the linear space

$$T_x\mathcal{M} := \{v \,|\, \nabla F(x) v = 0\}.$$

An explicit expression for $\nabla_{(v,\theta)} F$ was derived in [41] and takes the form

$$\nabla_{(v,\theta)} F(v,\theta,p,q) = \left(\lfloor \operatorname{diag}(\mathbf{Y}e)^\star \rceil + \lfloor \operatorname{diag}(e^\star) \rceil \begin{bmatrix} \mathbb{I}_N & 0 \\ 0 & -\mathbb{I}_N \end{bmatrix} \lfloor \mathbf{Y} \rceil \right) \mathbf{R}(e), \tag{14.1}$$

where $e = v\angle\theta$, $\lfloor \cdot \rceil$ is a "realification"-operator given by

$$\lfloor A \rceil := \begin{bmatrix} \operatorname{Re} A & -\operatorname{Im} A \\ \operatorname{Im} A & \operatorname{Re} A \end{bmatrix},$$

and where we have defined

$$\mathbf{R}(e) := \begin{bmatrix} \operatorname{diag}(\cos\theta) & -\operatorname{diag}(v)\operatorname{diag}(\sin\theta) \\ \operatorname{diag}(\sin\theta) & \operatorname{diag}(v)\operatorname{diag}(\cos\theta) \end{bmatrix}.$$

Choosing $e = \mathbb{1}$ (which is always possible) leads to the *flat-voltage* solution. In the absence of phase-shifting transformers (i.e., n = 1 and

$\theta^{\text{shift}} = 0$) and shunts (i.e. $y^{\text{sh}} = 0$ and $b_{kl}^{\text{sh}} = 0$ for all $kl \in \mathbf{E}$) we further have $\operatorname{diag}(\mathbb{1}\angle 0)(\mathbf{Y}(\mathbb{1}\angle 0))^* = p + jq = 0$. In other words, no power is injected into or extracted from the grid.

For this case, it was noted in [41] that, at the flat-voltage solution, we have

$$\nabla F(\mathbb{1}, 0, 0, 0) = \begin{bmatrix} \operatorname{Re}\mathbf{Y} & -\operatorname{Im}\mathbf{Y} & -\mathbb{I} & 0 \\ \operatorname{Im}\mathbf{Y} & \operatorname{Re}\mathbf{Y} & 0 & -\mathbb{I} \end{bmatrix}.$$

The system of equations $\nabla F(\mathbb{1}, 0, 0, 0)w = 0$ hence corresponds to the *linear coupled power flow* model introduced in [44]. If we further neglect line resistances (i.e., $\mathrm{r} = 0$) we arrive at

$$\nabla F(\mathbb{1}, 0, 0, 0) = \begin{bmatrix} 0 & -\operatorname{Im}\mathbf{Y} & -\mathbb{I} & 0 \\ \operatorname{Im}\mathbf{Y} & 0 & 0 & -\mathbb{I} \end{bmatrix}$$

which corresponds to the common *DC power flow* model [235].

Remark 14.1. The embedding presented in this chapter is by far not the only representation of the PFM. On one hand, one can switch between Cartesian and polar coordinates for both voltages and powers [181, 105]. On the other hand, for radial grids, one can also express the same physical laws using the *DistFlow* model [49, 19, 20], or even as the submanifold of rank-1 matrices of $\mathbb{C}^{N\times N}$ as a starting point for a convex relaxation [169].

Since we have shown in Chapter 4 that projected dynamical systems are coordinate-free, we can state that any projected gradient flow on the PFM can be expressed in any valid coordinate representation listed above. ■

14.4 Input-Output Map & Singularities

We have seen in the previous section that the $2N$-dimensional PFM has a straightforward embedding $\mathcal{M} \subset \mathbb{R}^{4N}$. To perform closed-loop optimization on the PFM, we need to identify an input-output map of the power system.

For this purpose, we need to partition the state vector $x := (v, \theta, p, q)$ into three subvectors, each containing a specific type of variable. Namely, $x_c \in \mathbb{R}^r$ denotes the *controllable variables*, and $x_d \in \mathbb{R}^s$ contains all variables that are governed by external processes, but that are not controllable. Together x_c and x_d form the *exogenous variables* x_{ex}. As a fact, we always have $r + s = 2N$. That is, the number of exogenous variables matches exactly the dimension of the PFM. The remaining $2N$ variables are called *endogenous* and denoted by x_{end}.

The partition of the state variables into x_c, x_d, and x_{end} is done by bus type, as summarized in Table 14.1. For a *PQ bus* i with flexible generation, active and reactive power injections p_i and q_i are controllable variables since they can be controlled directly via appropriate set-points, whereas v_i and θ_i will adapt autonomously to the current state of the grid. Hence, they are endogenous. Similarly, for a *PV bus* i (with flexible generation) the controllable variables are p_i and v_i while q_i and θ_i are endogenous. For a *PQ load bus* without generation p and q are uncontrollable, but exogenous.

Finally, one node (without loss of generality, we choose Bus 1) will act as the *slack bus* of the system. The voltage v_1 is controllable whereas θ_1 is uncontrollable since it will serve as the angle reference and be set to 0. Consequently, p_1 and q_1 are strictly speaking endogenous because by definition the slack bus power injection compensates for power imbalance rather than following a set-point.

Table 14.1: Partition of state variables by bus type

	Exogenous		Endogenous
	Controllable	Uncontrollable	
PQ generation	p_i, q_i		v_i, θ_i
PQ load		p_i, q_i	v_i, θ_i
PV generation	p_i, v_i		q_i, θ_i
slack bus	v_1	θ_1	p_1, q_1

From the input-output perspective we consider the exogenous variables as inputs, i.e., $u = x_{\text{ex}}$, even though some inputs are fixed and cannot

be adjusted. This applies in particular to the angle reference $\theta_1 = 0$ and the power injection at a PQ load bus i, which is given by $p_i = \underline{p}_i = \overline{p}_i$ in view of (13.2). The endogenous variables form the system outputs $y = x_{\text{end}}$.

Given these definitions of the inputs u and the output variables y, we can apply the Implicit Function Theorem [213] to conclude the local existence of an input-output map h whenever $\nabla_y F(y,u) \in \mathbb{R}^{2N\times 2N}$ has full rank.

More precisely, for every $(y,u) \in \mathcal{M}$ for which $\nabla_y F(y,u)$ has full rank, there exists a neighborhood $\mathcal{V}$ of u, a neighborhood $\mathcal{W}$ of y, and a continuously differentiable map $h : \mathcal{V} \to \mathcal{W}$ such that $F(h(u),u) = 0$ for all $u \in \mathcal{V}$. Moreover, the Jacobian of h is given by

$$\nabla h(u) := -\left(\nabla_y F(y,u)\right)^{-1} \nabla_u F(y,u)\,.$$

Since $\nabla_y F(y,u)$ has generally full rank in a neighborhood of the flat-voltage solution we can define h around this point. The largest (connected) neighborhood of the flat-voltage solution on which an input-output map exists will be referred to as the *high-voltage PFM* $\mathcal{M}^{\text{high}}$. More precisely, $\mathcal{M}^{\text{high}}$ is defined as the largest connected component of $\{(y,u) \in \mathcal{M} \,|\, \operatorname{rank} \nabla_y F(y) = 2N\}$ containing the flat-voltage solution. The set $\mathcal{M}^{\text{high}}$ is itself a manifold as a (relatively) open subset of $\mathcal{M}$.

In the subsequent chapters, we will study projected gradient flows and other dynamical systems on $\mathcal{M}$ and $\mathcal{M}^{\text{high}}$. As we can see already, a dynamical system defined on $\mathcal{M}^{\text{high}}$ is not guaranteed to admit complete solutions because trajectories may reach the boundary of $\mathcal{M}^{\text{high}}$ in finite time, resulting in the loss of a well-defined input-output map.

From a power systems perspective, we will refer to $\nabla_y F(y,u)$ as the *power flow Jacobian*, which is dependent on the bus types. A singular power flow Jacobian (also referred to as an *algebraic singularity*) is generally associated with *voltage collapse* conditions [83, 252]. The subset of $\mathcal{M}$ on which $\nabla_y F(y,u)$ is singular has been referred to as *impasse surface* [128, 155] as it splits $\mathcal{M}$ into multiple submanifolds, but it cannot generally be crossed. The following well-known example illustrates this idea.

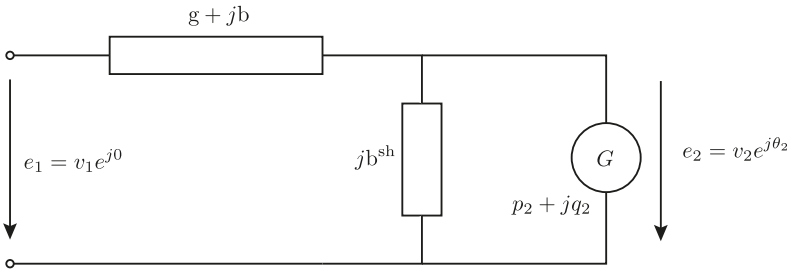

Figure 14.1: 2-bus system with shunt capacitor

Example 14.1 (2-bus grid / Nose Curve). Consider the 2-bus system in Figure 14.1 which is inspired by [251]. Namely, we consider a generator at Bus 2 connected to the rest of the grid (represented by Bus 1) through a lossy transmission line (without line shunt impedance). Bus 1 serves as the angle reference and exhibits a constant voltage magnitude of $v_1 = 1.0$. At the generator bus, a fixed shunt capacitor is installed.

Unlike the classical 2-bus system studied in [252, Ch 2], we consider a lossy transmission line and include a shunt capacitance at the terminal bus. Furthermore, for consistency with the sign convention, instead of a load connected to the terminal bus, we consider a generator. In particular, a positive value of p_2 refers to real power injected into the grid.

The power flow equations for the generator bus are given by

$$\begin{aligned} p_2 &= -v_1 v_2 (\mathrm{g} \cos\theta + \mathrm{b} \sin\theta) + \mathrm{g} v_2^2 \\ q_2 &= -v_1 v_2 (\mathrm{g} \sin\theta - \mathrm{b} \cos\theta) - \underbrace{(\mathrm{b} + \mathrm{b}^{\mathrm{sh}})}_{\tilde{\mathrm{b}}} v_2^2 \,. \end{aligned} \tag{14.2}$$

We can solve these equations for v_2. For this, we rearrange terms and square both sides of each equation, which yields

$$\begin{aligned} (p_2 - \mathrm{g} v_2^2)^2 &= v_1^2 v_2^2 (\mathrm{g}^2 \cos^2\theta + \mathrm{b}^2 \sin^2\theta + 2\mathrm{gb} \cos\theta \sin\theta) \\ (q_2 + \tilde{\mathrm{b}} v_2^2)^2 &= v_1^2 v_2^2 (\mathrm{g}^2 \sin^2\theta + \mathrm{b}^2 \cos^2\theta - 2\mathrm{gb} \cos\theta \sin\theta) \,. \end{aligned}$$

Adding up both equations and using $\sin^2\theta + \cos^2\theta = 1$, we get

$$p_2^2 - 2\mathrm{g}p_2v_2^2 + \mathrm{g}^2(v_2^2)^2 + q_2^2 + 2\tilde{\mathrm{b}}q_2v_2^2 + \tilde{\mathrm{b}}^2(v_2^2)^2 = v_1^2v_2^2(\mathrm{g}^2 + \mathrm{b}^2)$$

which is a quadratic function in v_2^2. After reorganizing the terms, we get

$$(\mathrm{g}^2 + \tilde{\mathrm{b}}^2)(v_2^2)^2 - (2\mathrm{g}p_2 - 2\tilde{\mathrm{b}}q_2 + v_1^2(\mathrm{g}^2 + \mathrm{b}^2))v_2^2 + p_2^2 + q_2^2 = 0\,,$$

and we see that (14.2) admits a solution if and only if

$$(2\mathrm{g}p_2 - 2\tilde{\mathrm{b}}q_2 + v_1^2(\mathrm{g}^2 + \mathrm{b}^2))^2 \geq 4(\mathrm{g}^2 + \tilde{\mathrm{b}}^2)(p_2^2 + q_2^2) \tag{14.3}$$

holds. More precisely, the points (p_2, q_2) which satisfy (14.3) with equality define the impasse line of the grid. Whenever, (14.3) is strictly satisfied, (14.2) admits two solutions corresponding to the two voltage surfaces separated by the impasse line. In this case, v_2 can be explicitly computed as

$$v_2^2 = \tfrac{1}{2(\mathrm{g}^2+\tilde{\mathrm{b}}^2)}\Big((2\mathrm{g}p_2 - 2\tilde{\mathrm{b}}q_2 + v_1^2(\mathrm{g}^2 + \mathrm{b}^2)) \pm \sqrt{(2\mathrm{g}p_2 - 2\tilde{\mathrm{b}}q_2 + v_1^2(\mathrm{g}^2 + \mathrm{b}^2))^2 - 4(\mathrm{g}^2 + \tilde{\mathrm{b}}^2)(p_2^2 + q_2^2)}\Big)\,.$$

In Figure 14.2, this parametrization of v_2 as a function of (p_2, q_2) is illustrated. In particular, one can clearly identify the two voltage surfaces separated by the impasse line. Each voltage set is an open submanifold of the PFM. The observation that low-voltage manifold exhibits an "kink" at the point $(p_2, q_2, v_2) = (0, 0, 0)$ does not contradict this fact because Figure 14.2 does not show an embedding of the PFM but merely a parametrization of the voltage magnitude. ■

Remark 14.2. The assignment of input and output variables according to Table 14.1 is very basic, yet it captures the main source of algebraic singularities.

First, since the variables x_d are uncontrollable (and assumed to be fixed) we may as well consider a *reduced* input-output map $y = \hat{h}(x_c)$ mapping x_c to the outputs. In other words, we have $F(\hat{h}(x_c), x_c, x_d) = 0$.

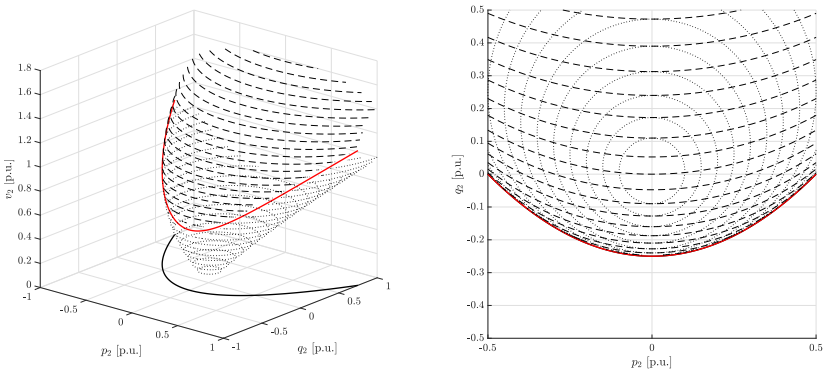

Figure 14.2: Voltage Surfaces of 2-bus example with impasse surface (red) and upper voltage surface (dashed voltage level curves) and lower voltage surface (dotted voltage level curves).

The existence of $\hat{h}$ depends on the non-singularity of $\nabla_y F(y, x_c, x_d)$ as before, and the Jacobian of $\hat{h}$ is given by

$$\nabla \hat{h}(x_c) := -\left(\nabla_y F(y, x_c, x_d)\right)^{-1} \nabla_{x_c} F(y, x_c, x_d)\,,$$

and therefore the domain of $\hat{h}$ is the same as for h.

Second, one can define an augmented output $\hat{y}$ composed of y and quantities that can be explicitly computed from u and y. One such example are the line currents $i^{\mathrm{f}} = \mathbf{Y}^{\mathrm{f}} e$ and $i^{\mathrm{t}} = \mathbf{Y}^{\mathrm{t}} e$. Defining new output variables in this way, does not affect the existence of the input-output map. ■

14.5 Projected Gradient Flows on the PFM

We now discuss why and under which conditions projected gradient flows on the PFM are well-defined, thus establishing the connection to the results in Parts I and II. At this point, we are not concerned about the feedback control design, but only about whether the desired closed-loop dynamics are well-posed.

Since we are interested in steering the physical grid to an optimal steady state instead of directly solving (13.2), we face a *reachability problem*: Assume that we actuate slowly enough not to excite any underlying dynamics, and, consequently, the grid is fully characterized by the steady-state ACPF equations. Further, assume that bus types are fixed and that inputs and outputs are separated according to Table 14.1. If the initial grid state is the flat-voltage solution, we cannot hope to steer the system outside of $\mathcal{M}^{\text{high}}$. Trying to do so would require passing through an algebraic singularity at which the steady-state input-to-output map is not well-defined.

For this reason, instead of the ACOPF problem (13.2), we consider the optimal steady-state problem

$$\begin{aligned} \text{minimize} \quad & \Phi(y,u) \\ \text{subject to} \quad & y = h(u) \\ & u \in \mathcal{U} \\ & (y,u) \in \mathcal{X} \\ & u \in \mathcal{U}^{\text{high}}\,, \end{aligned} \tag{14.4}$$

where $\mathcal{U}$ collects all the constraints on the input variables, e.g., the real and reactive power injection at PQ buses, and $\mathcal{X} \subset \mathbb{R}^{4N}$ combines the remaining (engineering) constraints which may depend on a combination of the input and output variable such as line flow limits.

More precisely, departing from (13.2), we generally have the structure

$$\mathcal{U} := \{u \,|\, \underline{u} \leq u \leq \overline{u}\}\,, \tag{14.5}$$

where $\underline{u}, \overline{u}$ combine components from $\underline{p}, \underline{q}, \underline{v}, \overline{p}, \overline{q}, \overline{v}$, respectively, and incorporate the angle reference constraint $0 \leq \theta_1 \leq 0$. The engineering constraints $\mathcal{X}$ take the form

$$\mathcal{X} := \left\{(y,u) \,|\, \underline{y} \leq y \leq \overline{y},\, i^{\text{t}}(y,u) \leq \overline{i},\, i^{\text{f}}(y,u) \leq \overline{i}\right\}\,, \tag{14.6}$$

where $\underline{y}, \overline{y}$ combine the remaining components from $\underline{p}, \underline{q}, \underline{v}, \overline{p}, \overline{q}, \overline{v}$, respectively. The notation $i^{\text{t}}(y,u)$ emphasizes the fact that the branch currents are a function of u and y.

Furthermore, in (14.4) the set

$$\mathcal{U}^{\text{high}} := \{u \in \mathbb{R}^{2N} \mid \exists y : \ (y, u) \in \mathcal{M}^{\text{high}}\}$$

defines the open set on which h is well-defined.

Next, we can reduce (14.4) by replacing y with $h(u)$, which results in

$$\begin{aligned} \text{minimize} \quad & \tilde{\Phi}(u) := \Phi(h(u), u) \\ \text{subject to} \quad & u \in \mathcal{U} \cap h^{-1}(\mathcal{X}) \cap \mathcal{U}^{\text{high}} , \end{aligned} \tag{14.7}$$

where $h^{-1}(\mathcal{X}) := \{u \in \mathcal{U}^{\text{high}} \mid (h(u), u) \in \mathcal{X}\}$.

Note that every (local) optimizer of the ACOPF problem (13.2) restricted to $\mathcal{M}^{\text{high}}$ is a (local) optimizer of (14.4) (and (14.7)) and vice versa. However, because $\mathcal{U}^{\text{high}}$ is an open set, (14.4) is not generally well-posed: It can happen (as we will see in the forthcoming illustrative examples) that a minimum of Φ is reached on the boundary of $\mathcal{U}^{\text{high}}$, i.e., at a voltage collapse point. More precisely, although Φ has an infimum over the feasible set of (14.4), it might not have a minimum.

For this reason, it is convenient, if not necessary, to introduce the following no-collapse assumptions:

Assumption 14.1 (strong no-collapse condition)**.** *The problem* (14.4) *satisfies* $\mathcal{U} \subset \mathcal{U}^{\text{high}}$.

Assumption 14.2 (weak no-collapse condition)**.** *The problem* (14.4) *satisfies* $(\mathcal{U} \cap h^{-1}(\mathcal{X})) \subset \mathcal{U}^{\text{high}}$ *and* $(\mathcal{U} \cap h^{-1}(\mathcal{X}))$ *is closed.*

In other words, under the strong no-collapse assumption, for every admissible input $u \in \mathcal{U}$, the steady-state map h exists and the feasible set of (14.4) (and (14.7)) is closed. Under the weak no-collapse assumption, the same is true for any feasible state (y, u) (i.e., $u \in \mathcal{U}$ and $(h(u), u) \in \mathcal{X}$). This distinction between weak and strong no-collapse conditions is relevant when discussing control designs: If one can guarantee that no transient constraint violations of $\mathcal{X}$ and $\mathcal{U}$ can occur, then the weak no-collapse assumption is sufficient to ensure well-posedness, i.e., that trajectories do not leave $\mathcal{M}^{\text{high}}$. However, if one can only guarantee that

the input constraints $\mathcal{U}$ are enforced at all times, but the remaining engineering constraints may be violated temporarily, the strong no-collapse assumption is required.

Under normal operating conditions, a well-designed power grid with effective engineering constraints will satisfy (at least) the weak no-collapse condition. In particular, constraints on the bus voltage magnitudes guarantee that voltage collapse cannot occur. However, under critical conditions, e.g. after line failures, these no-collapse assumptions might not be satisfied anymore [251, 252].

At this point, we can introduce two types of optimization dynamics based on projected gradient flows on $\mathcal{M}^{\text{high}}$. These dynamics define the closed-loop behavior that we implement in the next chapter with appropriate feedback controllers. The purpose of this section is to summarize the fact that these dynamics are mathematically well-posed and exhibit desirable convergence and stability properties.

A first model will approximately enforce the engineering constraint set $\mathcal{X}$ with a penalty function while the second approach incorporates both $\mathcal{U}$ and $\mathcal{X}$ directly into the projected gradient flow.

14.5.1 Penalty-Relaxed Projected Gradient Flows

Let the strong no-collapse condition (Assumption 14.1) hold. Then, (14.7) is equivalent to

$$\begin{aligned} &\text{minimize} && \tilde{\Phi}(u) \\ &\text{subject to} && u \in \mathcal{C} := \mathcal{U} \cap h^{-1}(\mathcal{X})\,. \end{aligned} \tag{14.8}$$

Further, let $\Psi_{\mathcal{X}} : \mathcal{M}^{\text{high}} \to \mathbb{R}$ be an appropriate *penalty function* for the engineering constraints $\mathcal{X}$. We will assume that $\Psi_{\mathcal{X}}$ is a quadratic penalty on constraint violations, i.e., in view of (14.6), we define

$$\begin{aligned} \Psi_{\mathcal{X}}(y,u) := &\,\|\max\{0, y - \overline{y}\}\|^2 + \|\max\{0, \underline{y} - y\}\|^2 \\ &+ \|\max\{0, i^{\text{t}}(y,u) - \overline{i}\}\|^2 + \|\max\{0, i^{\text{f}}(y,u) - \overline{i}\}\|^2\,. \end{aligned} \tag{14.9}$$

This allows us to approximate (14.8) by

$$\text{minimize} \quad \tilde{\Phi}(u) + \rho\tilde{\Psi}_{\mathcal{X}}(u) \qquad \text{subject to} \quad u \in \mathcal{U}\,, \tag{14.10}$$

where $\rho > 0$ is a penalty parameter and $\tilde{\Psi}_{\mathcal{X}}(u) := \Psi_{\mathcal{X}}(h(u), u)$. Given a continuous metric $\mathcal{U} \to \mathbb{S}^{2N}_{+}$, we consider the projected gradient flow

$$\dot{u} = \Pi^{G}_{\mathcal{U}}\left[-\operatorname{grad}_G(\tilde{\Phi} + \rho\tilde{\Psi}_{\mathcal{X}})\right](u)\,. \tag{14.11}$$

Note that any solution $u : [0, T] \to \mathcal{C}$ with $T > 0$ of (14.11) defines a trajectory on the PFM by the fact the outputs at time t are given by $y(t) = h(u(t))$ and consequently $(y(t), u(t)) = (v(t), \theta(t), p(t), q(t)) \in \mathcal{M}$.

For the system (14.11), we can apply the results from previous chapters and summarize their consequences as follows:

Statement 1. *Consider the projected gradient flow (14.11) derived from the ACOPF problem (13.2). If Assumption 14.1 (strong no-collapse condition) is satisfied the following hold:*

(i) The system (14.11) admits a complete solution for every initial condition $u(0) \in \mathcal{U}$.

Reason: *Because Corollaries 4.2 and 4.3 apply. In particular, $\mathcal{U}$ is defined by box constraints according to (14.5), hence it is compact and (trivially) Clarke regular.*

(ii) Every trajectory of (14.11) converges to the KKT points of (14.10), but only minimizers can be stable, and asymptotic stability of a limit point implies strict minimality.

Reason: *Because Proposition 8.1 and Theorem 8.1 apply.*

(iii) Assume that the power system (and low-level control) dynamics $\dot{\zeta} = f(\zeta, u)$ are exponentially stable for fixed u. Further, let the input-output map $y = h(u)$ result from combination the steady-state map of $\dot{\zeta} = f(\zeta, u)$ and an output map $y = h'(\zeta, u)$.

Then, trajectories of the feedback interconnection of (14.11) *and* $\dot{\zeta} = f(\zeta, u)$ *with a scalar control gain* ϵ *converge to the KKT points of the ACOPF problem* (14.10) *for all* ϵ *smaller than some threshold* $\epsilon^\star$. *However, only minimizers can be stable, and asymptotic stability implies strict minimality.*

Reason: *Because Theorem 11.1 applies.*

(iv) The system (14.11) *(as well as the interconnection with a dynamical plant* $\dot{\zeta} = f(\zeta, u)$*) can be approximated in terms of uniform convergence and semiglobal practical asymptotic stability with an* anti-windup implementation *as in Chapter 5.*

Reason: *Because Corollary 5.2 and Theorem 5.2 apply.*

(v) Independently of the anti-windup gain K*, any equilibrium point of the anti-windup implementation of* (14.11) *corresponds the power grid being at a KKT point of* (14.10). *Moreover, if* $G \equiv \mathbb{I}$ *is the identity matrix*[a]*, then trajectories of the anti-windup implementation converge to the set of these equilibrium points for any* $K > 0$.

Reason: *Because Proposition 9.1 and Theorem 9.3 apply.*

[a] Presumably, the assumption that $G \equiv \mathbb{I}$ can be relaxed in Theorem 9.3.

Statement 1 can be strengthened in several ways. For instance, under pertinent Lipschitz assumptions, solutions of (14.11) are unique.

Further, the constraint structure of $\mathcal{U}$ can be generalized as long as $\mathcal{U}$ remains compact, Clarke regular, prox-regular, or convex, depending on the specific statement. For example, the box constraints on power injection can be replaced with more general capability curves incorporating field and stator current limits of synchronous machines, or power factor and apparent power constraints of inverters.

14.5.2 Fully Projected Gradient Flow

As a second desirable closed-loop behavior, we incorporate the engineering constraints $\mathcal{X}$ into the projected gradient dynamics. This choice

allows us to reach slightly different conclusions.

First, we weaken the no-collapse assumption. Namely, note that (14.8) is equivalent to

$$\begin{aligned} &\text{minimize} && \tilde{\Phi}(u) \\ &\text{subject to} && u \in \mathcal{C} := \mathcal{U} \cap h^{-1}(\mathcal{X}) \end{aligned} \tag{14.12}$$

under the weak no-collapse condition (instead of the strong no-collapse condition).

Given a continuous metric $G : \mathcal{C} \to \mathbb{S}^{2N}_{+}$, we can therefore consider the projected gradient flow

$$\dot{u} = \Pi^{G}_{\mathcal{C}} \left[-\operatorname{grad}_G \tilde{\Phi} \right](u) \tag{14.13}$$

to solve (14.12) instead of the penalty-relaxed problem (14.10).

Statement 2. *Consider the projected gradient flow (14.13) derived from the ACOPF problem (13.2). If the weak no-collapse assumption is satisfied the following hold:*

(i) The system (14.13) generically admits a complete solution for every initial condition $u(0) \in \mathcal{C}$.

Reason: *Because Corollaries 4.2 and 4.3 apply, as before, if $\mathcal{C}$ is compact and Clarke regular. Compactness is guaranteed because $\mathcal{C} \subset \mathcal{U}$ and $\mathcal{U}$ is compact. Clarke regularity is guaranteed if $\mathcal{C}$ satisfies LICQ (as seen in Example 2.2). LICQ is, however, generically satisfied for ACOPF problems of the form (13.2) as shown in [Ha11].*

(ii) Every trajectory of (14.13) starting in $\mathcal{U}$ converges to the KKT points of (13.2) which are "high-voltage" solutions. Only minimizers can be stable, and asymptotic stability of a limit point implies strict minimality.

Reason: *Because Proposition 8.1 and Theorem 8.1 apply.*

Under the weak additional assumption that the "linearized output

sets" satisfy LICQ (cf. Assumption 6.3), the following hold:

(iii) The system (14.13) can be approximated in terms of uniform convergence and semiglobal practical asymptotic stability with an LOP discretization *as proposed in Chapter 6.*

Reason: *Because Theorems 6.1 and 6.2 apply.*

(iv) Moreover, for all step-sizes α smaller than some threshold $\alpha^\star$, the solutions of the LOP discretization converge to critical points of (13.2) on $\mathcal{M}^{\text{high}}$ and asymptotic stability of a limit point implies strict minimality.

Reason: *Because Theorem 10.1 applies.*

14.6 Riemannian Metrics on the PFM

We have seen that the projected gradient flows (14.11) and (14.13) exhibit the properties in Statements 1 and 2 for any continuous metric G on $\mathcal{U}$ and $\mathcal{C}$, respectively. It remains to discuss viable choices for G.

In the following, we present three different metrics that have been considered in the literature (although they have not been identified as such). With a small example on the 2-bus test case, we then illustrate the metrics' different behavior with respect to algebraic singularities.

This investigation is motivated by the fact that feedback-based optimization schemes have been proposed to manage curative actions in critical operating conditions after contingencies [144, 174]. Therefore, it is natural to ask whether feedback-based algorithms exhibit a robust behavior against voltage collapse in these challenging situations. Even though our insights are very preliminary, they indicate that the choice of metric is crucial for this purpose.

Explicit Euclidean Metric

A straightforward choice for the metric G is the identity matrix $\mathbb{I}_{2N}$, which refer to as the *explicit (Euclidean) metric* G^{E}. Clearly, since

$G^{\mathrm{E}} \equiv \mathbb{I}_{2N}$ is constant, Statements 1 and 2 hold for the penalty-relaxed partially projected gradient flow (14.11) as well as the fully projected gradient flow (14.13), respectively.

By using the explicit metric, we fully accept the viewpoint that we can formulate the ACOPF problem in the space of inputs $u \in \mathbb{R}^{2N}$. The penalty-relaxed partially projected gradient flow (14.11) can consequently be discretized as a standard projected gradient descent scheme (see Remark 6.1) of the form

$$u^+ = P_{\mathcal{U}}\left(u - \alpha \nabla \tilde{\Phi}(u)\right) = P_{\mathcal{U}}\left(u - \alpha H(u)^T \nabla \Phi(h(u), u)^T\right) ,$$

where $\alpha > 0$ is a step-size and $H(u)^T = \begin{bmatrix} \nabla h(u)^T & \mathbb{I}_{2N} \end{bmatrix}$. This approach has been pursued, for example, in [111] with a barrier instead of a penalty function.

Apart from its simplicity, one particular advantage of this metric is that it preserves sparsity and is therefore well-suited for distributed implementations. Although, for this to work, other requirements have to be met. E.g., $\mathcal{U}$ and $\tilde{\Phi}$ need to decomposable and $\nabla h(u)$ needs to be approximated with a sparse matrix. In particular, if $\mathcal{U} = \mathcal{U}_1 \times \ldots \times \mathcal{U}_N$ decomposes into N sets, each associated with a single bus, the projection onto $\mathcal{U}$ can be performed locally at each bus.

The "explicit" viewpoint, i.e., considering a problem purely in the input space, has been adopted not only for projected gradient schemes, but also for constrained saddle-point approaches [29, 57, 71] and dual methods [43, 42].

Newton Metric

Alternatively to the explicit Euclidean metric, if Φ is strongly convex and twice continuously differentiable, one can choose

$$G^N(u) := \nabla^2_{uu} \tilde{\Phi}(u) = H(u)^T \nabla^2_{xx} \Phi(h(u), u) H(u) \succ 0$$

resulting in a projected Newton flow in the input variables u. Hence we refer to G^{N} as the *Newton metric* on the PFM. At least in the unconstrained case, this metric yields superlinear convergence when discretized.

If a penalty function is used to enforce the engineering constraints $\mathcal{X}$, the Hessian of the penalty-augmented cost function $\tilde{\Phi} + \rho\tilde{\Psi}$ should be used. However, since a penalty function like (14.9) is not twice continuously differentiable, additional theoretical considerations have to be made to guarantee well-posedness.

This approach has been taken in [241], where a time-varying setup was considered, and a quasi-Newton method was used to approximate the Hessian of a penalty-augmented cost function and yield a distributed implementation.

One drawback of the Newton metric is the fact that it puts additional requirements on the cost function and does, in general, not allow for constant cost function terms, e.g., renewable generation with zero marginal cost of generation.

Implicit Metric

As a final possibility, we consider the *implicit Euclidean metric* defined by

$$G^{\mathrm{I}}(u) := H(u)^T H(u) \succ 0\,.$$

This metric has been used in [Ha14, Ha15] and results from projecting the ambient gradient $\nabla\Phi(v,\theta,p,q) \in \mathbb{R}^{4N}$ onto the tangent space $T_x\mathcal{M} \subset \mathbb{R}^{4N}$, i.e., $\operatorname{grad}_{G^{\mathrm{I}}}\Phi(u)$ is given by the solution of

$$\underset{w}{\text{minimize}} \quad \tfrac{1}{2}\|\nabla\Phi(x)^T - w\|^2 \qquad \text{subject to} \quad \nabla F(x)w = 0\,. \tag{14.14}$$

To motivate the expression for G^{I}, note that w can be decomposed into w_u and w_y according the input-output assignments. Consequently, we may write $\nabla F(x)w = 0$ as

$$w_y = -\left(\nabla_y F(x)\right)^{-1} \nabla_u F(x) w_u = \nabla h(u) w_u\,.$$

Hence, using $w = H(u)w_u$, we can rewrite (14.14) as

$$\underset{w_u}{\text{minimize}} \quad \tfrac{1}{2}\|\nabla\Phi(x)^T - H(u)w_u\|^2\,, \tag{14.15}$$

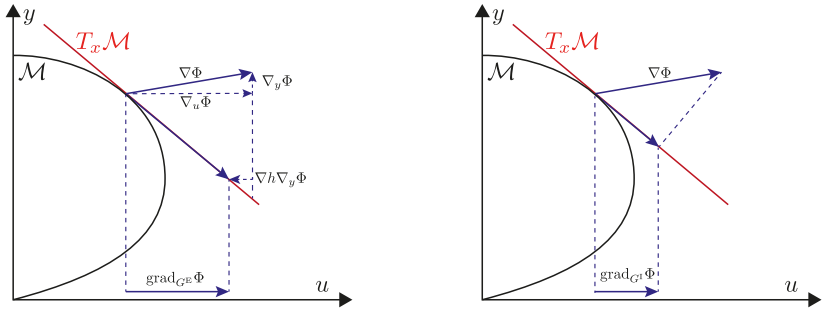

Figure 14.3: Construction of explicit (left) and implicit (right) gradient from the ambient Euclidean gradient $\nabla\Phi(x)$. The explicit gradient is given by $\text{grad}_{G^{\text{E}}}\,\Phi = \nabla_u\Phi + \nabla h\nabla_y h\Phi$ (omitting arguments) whereas the implicit gradient is obtained from a projection of $\nabla\Phi$ onto the tangent space $T_x\mathcal{M}$.

and we conclude that $w_u^\star = \text{grad}_{G^{\text{I}}}\,\Phi(x)$ satisfies

$$\nabla\Phi(x)^T = H(u)\,\text{grad}_{G^{\text{I}}}\,\Phi(u)\,.$$

Multiplying both sides with $(H(u)^T H(u))^{-1}H(u)^T$ yields the fact that G^{I} has to be defined as $G^{\text{I}}(u) = H(u)^T H(u)$.

Although it is more involved in its (algebraic) construction, the implicit metric can be shown to be well-defined the entire PFM $\mathcal{M}$ rather than just on $\mathcal{M}^{\text{high}}$. In particular, trajectories of the projected gradient flow models (14.11) and (14.13) are independent of the choice of input and output variables.

This fact is illustrated in Figure 14.3, where the definition of the implicit metric is compared to the construction of the explicit gradient. In particular, we can see that the implicit gradient is well-defined for any linearization point x on $\mathcal{M}$ and its geometric construction does not rely on the existence of h. The explicit gradient, on the other hand, requires $\nabla h(x)$ to map the gradient component $\nabla_y\Phi$ into the input space. Namely, we have $\text{grad}_{G^{\text{E}}}(x)\Phi = \nabla_x\Phi(x)^T + \nabla h(x)^T\nabla_y\Phi(x)^T$.

Similarly to the implicit metric, the Newton metric also enjoys the

property of being well-defined on all of $\mathcal{M}$ independently of the input-output assignment. However, for the Newton metric to be well-defined, the additional requirement that Φ is twice continuously differentiable and strongly convex (in $\mathbb{R}^{4N}$) needs to be satisfied.

With the implicit or Newton metric used in the full input-output projection model (14.13) it is therefore possible, in theory, to swap of input and output variables during runtime. However, we will see in the next section that the Newton metric, nevertheless, has shortcomings with respect to avoiding algebraic singularities.

14.7 Numerical Illustration: 2-Bus Case

To illustrate the effect of different metrics, we study unconstrained gradient descent on the power flow manifold of the 2-bus test case introduced in Example 14.1.

We ignore all constraints such as voltage limits, line ratings, etc. except for the power flow equations. Namely, we consider unconstrained gradient flows on $\mathcal{M}^{\text{high}}$ in the input coordinates the form

$$\begin{aligned} \dot{u} = -\operatorname{grad}_G \tilde{\Phi}(u) &= -G(u)^{-1} \nabla \tilde{\Phi}(u)^T \\ &= -G(u)^{-1} H(u)^T \nabla \Phi(h(u), u)^T , \end{aligned} \tag{14.16}$$

where $H(u)^T = \begin{bmatrix} \nabla h(u)^T & \mathbb{I}_{2N} \end{bmatrix}$ holds with h being the steady-state map on $\mathcal{U}^{\text{high}}$, as before. Clearly, (14.16) is well-defined only on $\mathcal{U}^{\text{high}}$, but neither the strong nor the weak no-collapse assumption hold.

In order to simulate (14.16), we perform a simple forward Euler integration where we use the standard ACPF solver shipped with MATPOWER to evaluate $h(u)$. The step-size in the following simulations has been chosen for each metric separately, and results are not representative of convergence speed.

The cost is a convex quadratic function in (p_2, q_2) with a minimizer outside of $\mathcal{U}^{\text{high}}$. The minimizer subject to $\mathcal{M}$ lies on the impasse surface. More concretely, the cost function promotes power delivery at Bus 2 and has a minimum for a larger power output than the maximum power that the line can deliver.

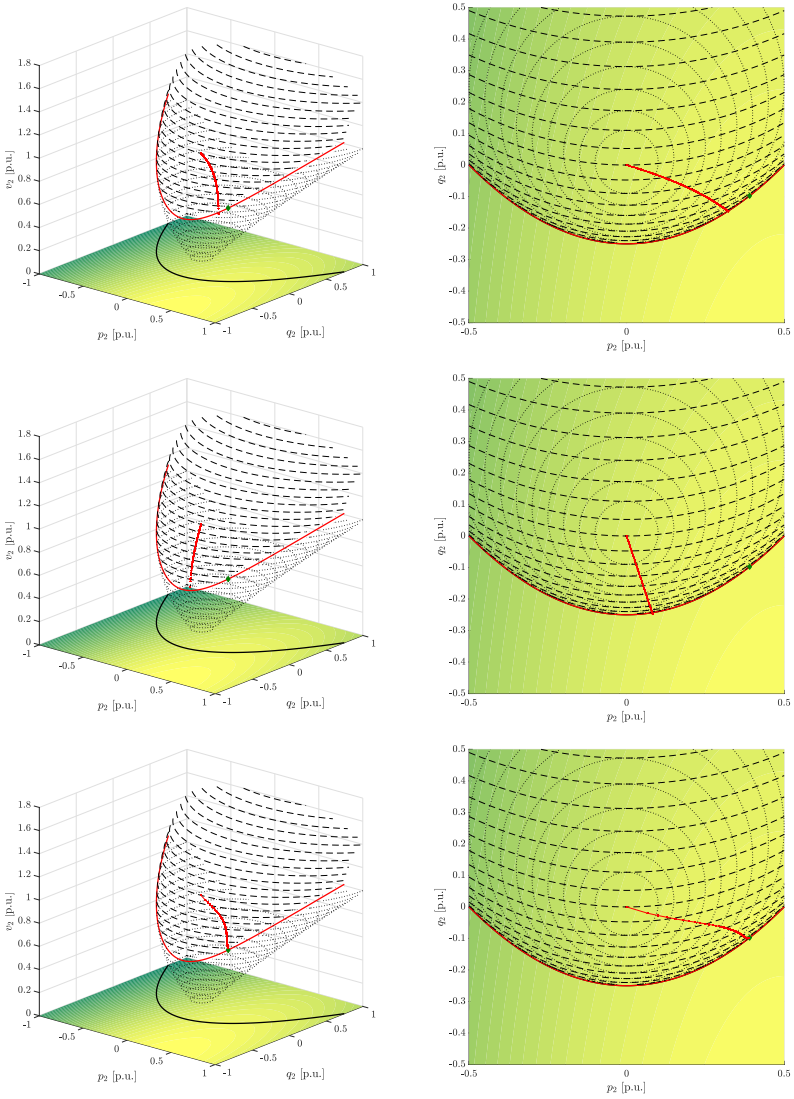

Figure 14.4: Trajectories for explicit (top), Newton (middle), and implicit gradient descent (bottom) on the 2-bus "nose curve"; The green diamond represents the constrained minimizer.

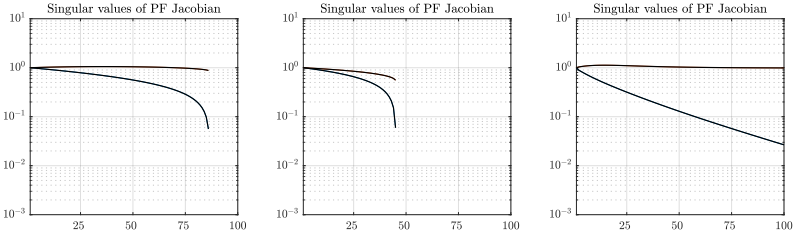

Figure 14.5: Conditioning of power flow Jacobian for explicit (left), Newton (middle), and implicit gradient descent (right)

Figure 14.4 illustrates the behavior of a gradient descent on the PFM for each of the three proposed metrics for the 2-bus grid from Example 14.1.

The explicit and Newton gradient descents ignore the nonlinearity of the power flow equations. This becomes particularly apparent in the $p-q$-plane (right-hand panels in Figure 14.4): The trajectory of the explicit gradient descent crosses level curves in a normal direction (which is a classical feature of gradient flows in $\mathbb{R}^n$). The Newton method approaches the unconstrained optimizer in the $p-q$-plane in a straight line since the cost is quadratic (in p_2 and q_2). These methods reach and cross the impasse line after few iterations upon which the numerical integration fails because no valid power flow solution exists anymore.

The trajectory based on the implicit metric, on the other hand, also approaches the impasse surface but slows down in its vicinity since ∇h is almost singular.

The difference between the three algorithm variations becomes particularly apparent when comparing the evolution of the singular values of the power flow Jacobian in Figure 14.5. While the trajectories based on the explicit and Newton metrics experience a sharp drop in the lower singular value right before reaching the impasse line, the trajectory for the implicit metric experiences a slower decrease and does not cross the impasse line within a reasonable number of iterations.

The fact that this phenomenon persists when decreasing the step-

size indicates that this type of collapse is not a numerical artifact but an inherent property of these iterative schemes for the setup under consideration.

Also, note that only the implicit metric produces a trajectory that converges to the constrained minimizer on the PFM, indicated by the green diamond shape in Figure 14.4. The trajectories based on the explicit and the Newton metric, on the other hand, move towards the unconstrained optimizer, essentially ignoring the fact the constrained optimizer has to lie on $\mathcal{M}^{\text{high}}$.

14.8 Notes & Comments

Dynamical systems subject to algebraic constraints are common in power systems analysis. These *differential-algebraic equations* appear naturally due to timescale separation and singular perturbation considerations, where fast dynamics are approximated by their steady-state relation. In our case, these algebraic constraints are given by the ACPF equations and lower-level controllers.

Interpreting these constraints as a differentiable manifold has allowed us to consider our optimal steady-state control problem from a different perspective: Rather than designing a feedback controller directly, given an input-output map of a system, we first define the desired closed-loop dynamics in the form of a dynamical system on the PFM.

By doing so, we have been able to identify families of possible optimization dynamics by their constraint enforcement mechanism and the (Riemannian) metric defined on the manifold.

Importantly, this geometric viewpoint has also allowed us to conceive dynamics independent of the particular choice of input and output variables. This formulation opens the door to study "switched" controllers, which dynamically change the set controllable variable, e.g., switching between voltage-controlled and reactive-power-controlled generation modes.

Our investigation of three possible metrics on the PFM remains very

preliminary. From a theoretical viewpoint, it is clear that the implicit and Newton metrics are well-defined on the entire PFM, independently of the input-output assignments. In contrast, the explicit metric is defined only on the high-voltage components of the PFM induced by a particular configuration of bus types.

The choice of metric appears to significantly affect algorithm performance in situations where the power flow Jacobian is ill-conditioned, e.g., close to algebraic singularities. Besides the simulations in Section 14.7, this behavior is also supported by the fact that, for the larger grids studied in the forthcoming chapter, algorithms using the explicit metric are much more difficult, if not impossible, to tune.

CHAPTER 15

Numerical Simulations

We can now illustrate the different behaviors of three feedback-based optimization schemes in a more realistic power system setup. We consider the following control designs:

(i) A projected Euler discretization of the penalty-relaxed partially projected gradient flow (14.11), which we refer to as *penalty scheme*,

(ii) An extension of the penalty scheme with an integrator for the constraint violations. We refer to this design as *primal-dual scheme*, since it is technically based on a saddle-point flow.

(iii) A LOP discretization of the fully projected gradient flow (14.13) as studied in Chapters 6 and 10.

For all three algorithms we consider only the implicit Euclidean metric since, for the test case under consideration, the explicit and Newton metrics lead to satisfactory performance only after additional regularization and extensive tuning.

In terms of computational complexity, each of the algorithms requires the solution of a QP at each iteration. Moreover, these QPs scale linearly with the number of buses and nodes in the network.

Recall that solving such a convex (and often sparse) QP is computationally very efficient since specialized solvers with strong theoretical performance guarantees are available.

In contrast, a "feedforward" scheme that attempts to solve the non-convex ACOPF problem directly relies on more sophisticated (and, arguably, more fragile) solvers which do not come with the same level of theoretical guarantees and often require the solution of QPs as a subroutine at every iteration. On top of that, feedforward schemes require more data than just an estimate of the input-output sensitivities and measurements. They may produce solutions that "jump" between sampling instants. In other words, feedforward methods do not come with the same robustness guarantees as developed for feedback-based schemes in the previous parts of this thesis.

For the simulations, we consider a challenging scenario using a modified version of the IEEE 30-bus system with large solar and wind plants. In particular, we simulate a time horizon of 24 h in which power consumption and the availability of renewable energy varies. The objective is to autonomously track the solution of a time-varying ACOPF and, in particular, to enforce the various constraints on voltage, line flows, and generator output.

Besides their capability to track the solution of the underlying ACOPF and to enforce constraints, we also investigate the effects of time delays and the impact of inexact input-output sensitivities on the three schemes.

15.1 30-Bus Test System

Our forthcoming simulations are based on a modified version of the IEEE 30-bus system [268, 205] illustrated in Figure 15.1. This modified setup has previously been used in [Ha14].

The grid topology, line parameters, and shunt capacitors correspond to the original IEEE 30-bus model. The transformer tap ratios, as well as the placement of the synchronous condensers, were taken from [205]. The placement of the generators at Bus 1 and 2 is original, but their capacity has been adjusted to better fit the new context. Similarly, the load distribution has been slightly changed. Voltage limits are slightly relaxed from [268] to [0.94, 1.06] p.u.. For easier visualization, line ratings are ignored except for the two lines connecting Bus 6 and 8 and

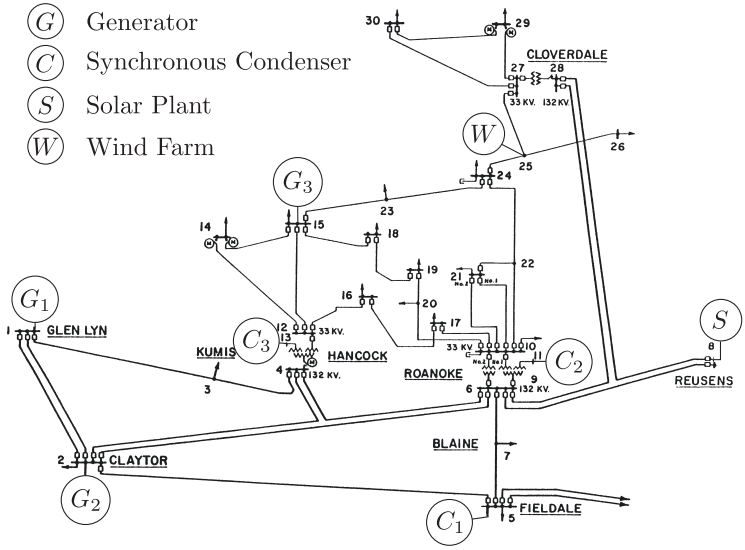

Figure 15.1: modified IEEE 30-bus test system

Bus 24 and 25, respectively.

Two large renewable plants have been added to the system: On one hand, a large solar plant is connected to Bus 8, and, on the other hand, a wind farm feeds power into the grid at Bus 25. Furthermore, an additional conventional generator has been added to Bus 15 to demonstrate the effects of a generator outage.

Unless noted otherwise, the objective function is a sum of quadratic cost terms associated to each generator i of the form $C_i(p_i) = a_i p_i^2 + b_i p_i$. The cost parameters and other generator data are given in Table 15.1. Voltage limits at all nodes are at 0.94 p.u. and 1.06 p.u.. The objective is loosely based on the economic cost of each generator. For instance, the power produced by the wind and solar plants has close to zero marginal cost of production and should be used whenever available.

We will consider the steady-state power flow in this 30-bus grid over a 24-h period sampled in **1-minute intervals**, and we try to track

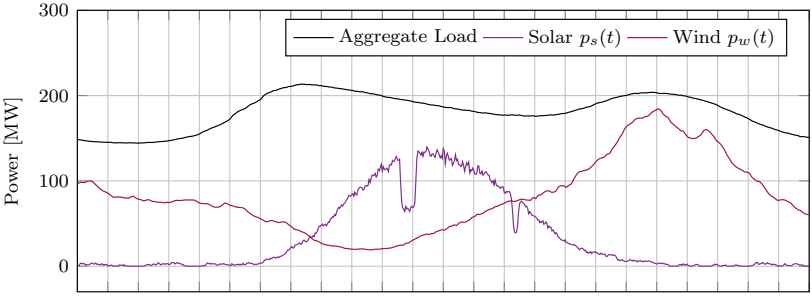

Figure 15.2: Load, solar, and wind profiles for 30-bus test case

the solution of the ACOPF problem (13.2). The main sources of time-variability are power consumption (which is assumed to be inflexible) and the availability of wind and solar energy (which can be curtailed). The corresponding profiles are illustrated in Figure 15.2. Besides time-varying loads and renewable generation, we also consider an outage of Generator 3 at 4:00, which puts the grid into a stressed state for the remainder of the simulation horizon.

	Unit	Type	Node #	a_i	b_i	$\bar{p}_i^{\mathrm{G}}$	$\underline{p}_i^{\mathrm{G}}$	$\bar{q}_i^{\mathrm{G}}$	$\underline{q}_i^{\mathrm{G}}$
G_1	Gen.1	slack	1	0.1	0.9	100	0	50	-50
G_2	Gen.2	PV	2	0.04	0.5	120	0	50	-50
G_3	Gen.3	PQ	15	0.02	0.2	100	0	50	-50
C_1	Cond.1	PV	5	0	0	0	0	50	-50
C_2	Cond.2	PV	11	0	0	0	0	50	-50
C_3	Cond.3	PV	13	0	0	0	0	50	-50
S	Solar	PQ	8	0	0	$\bar{p}_{\mathrm{s}}(t)$	0	50	-50
W	Wind	PQ	25	0	0	$\bar{p}_{\mathrm{w}}(t)$	0	50	-50

Table 15.1: Cost coefficients a and b in [\$/MW^{2}h] and [\$/MWh], respectively. Active power generation limits in [MW] and reactive power generation limits in [MVAr]. The system base power is fixed to 100MVA.

The optimal cost and power flow solution of an offline ACOPF computation for the entire time horizon are given in Figures 15.3 and 15.4, respectively. Note that this solution does not constitute a "feedforward"

strategy but rather a sequence of ACOPF solutions. Individual solutions might not be connected by a continuous transition (as it would be the case for a physical system). Furthermore, this offline solution does not incorporate any delay between measurement, computation, and actuation. Hence, Figures 15.3 and 15.4 represent the *ground truth* optimal solution rather than a feedforward control.

Further, note the following points about Figures 15.3 and 15.4:

- The ACOPF problem (13.2) is feasible for the entire simulation horizon.
- The outage of Generator 3 is immediately compensated mostly by Generator 2, which is cheaper to operate than Generator 1.
- Solar and wind generation have to be curtailed between 10:00-13:00 and 17:00-21:30 in order not to overload the line between Buses 6-8 and 24-25, respectively.
- The optimal response to the dip in solar generation around 11:00 is a coordinated action, adjusting both real and reactive power generation to satisfy voltage limits and meet the demand.
- Around 13:00, the offline solution is volatile and changes significantly between sampling instances, particularly, with respect to reactive power and voltage magnitude. This volatility is due to the absence of any regularizing cost component on reactive power generation. As a consequence the objective function is very "flat" in the q-v-dimension and thus sensitive to perturbations.

15.2 Feedback-Based Online Optimization Schemes

To track the ACOPF solution of the 30-bus test setup as illustrated in Figure 15.4, we study the three feedback-based optimization schemes presented in the following

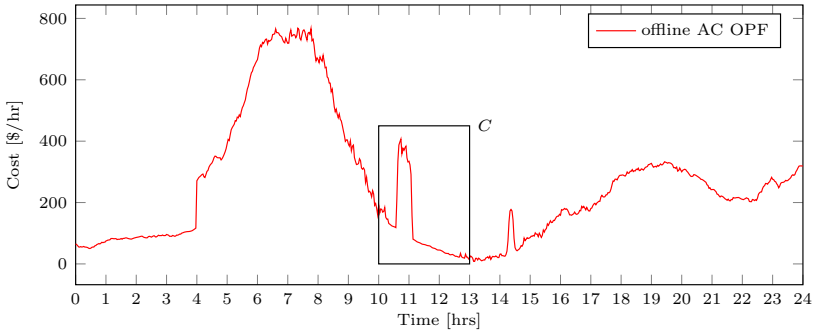

Figure 15.3: A-posteriori offline ACOPF optimal cost for 30-bus case

15.2.1 Penalty Scheme

As a first algorithm we consider a projected forward Euler discretization of (14.11) featured in [Ha14] and given by

$$\begin{aligned} u^{k+1} &= P_{\mathcal{U}^k}^{G^{\mathrm{I}}(u^k)} \left(u^k - \alpha G^{\mathrm{I}}(u^k)^{-1} \tilde{H}(y^k, u^k) \nabla \left(\Phi + \rho \Psi_{\mathcal{X}} \right) (y^k, u^k)^T \right) \\ y^{k+1} &= h \left(P_{\mathcal{U}^{k+1}}(u^{k+1}) \right) , \end{aligned} \tag{15.1}$$

where $\Psi_{\mathcal{X}}$ is the penalty function from (14.9), $\alpha > 0$ is a step-size, and

$$P_{\mathcal{U}^k}^{G^{\mathrm{I}}(u^k)}(w) := \arg\min_{v \in \mathcal{U}^k} \|w - v\|^2_{G^{\mathrm{I}}(u^k)}$$

is the projection onto the set of admissible inputs $\mathcal{U}^k$ at iteration k with respect to the norm induced by $G^{\mathrm{I}}(u^k) \in \mathbb{S}^{2N}_+$. Note that $\mathcal{X}$ is stationary since it does not express any time-varying constraints because solar and wind power generation constraints are exclusively part of $\mathcal{U}$.

Furthermore, $\tilde{H}(y^k, u^k)$ replaces $H(u)$ from the nominal scheme, because the time-varying loads perturb the power flow equations (13.1) and therefore, the input-output sensitivities ∇h are not solely a function of u. However, by using the measurement y^k, we know the full state, and we can evaluate the linearization of the power flow equations according to (14.1). By doing so, we do not need to explicitly estimate the loads.

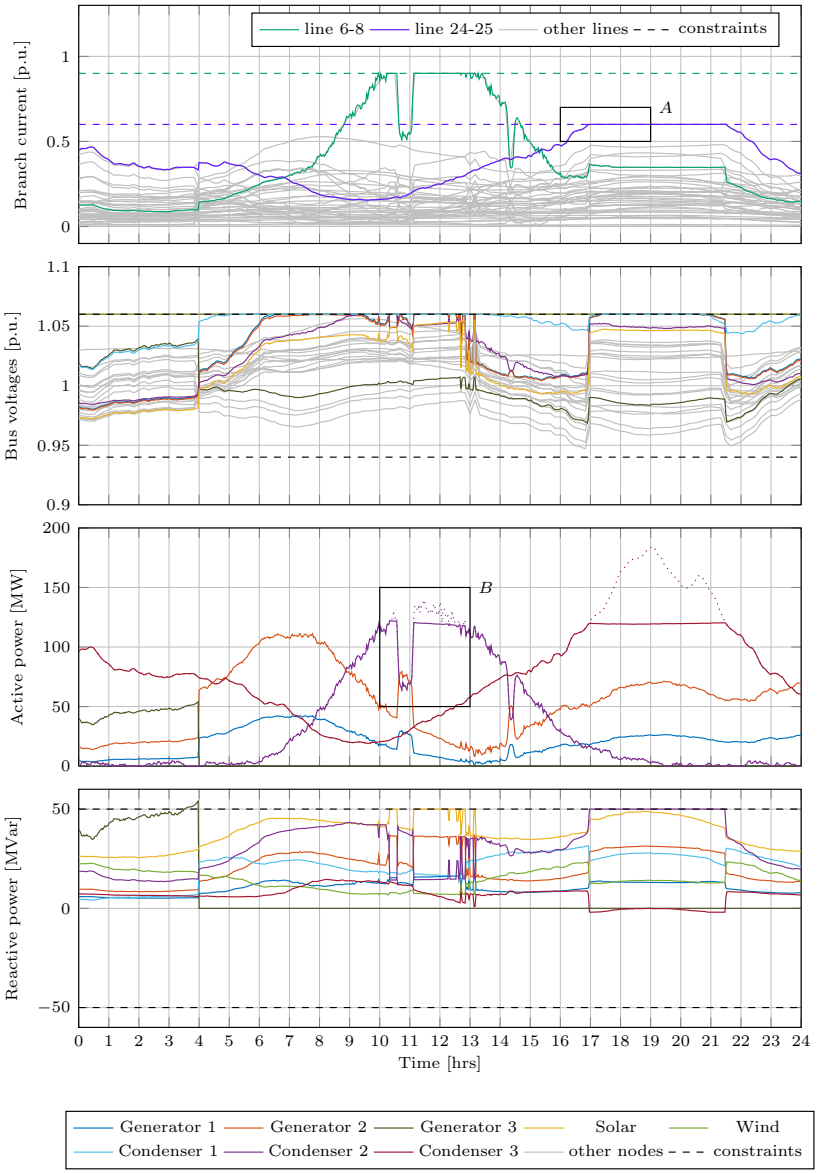

Figure 15.4: A-posteriori offline ACOPF computation for 30-bus case

Remark 15.1. Notice that the set $\mathcal{U}^k$ varies over time due to the time-varying availability of the solar and wind plants. The computed control input u^{k+1} lies in $\mathcal{U}^k$. However, the physical system does not simply "evaluate" $h(u^{k+1})$, but also performs the input saturation onto the true feasible set $\mathcal{U}^{k+1}$ at time $k+1$. The underlying assumption is that the control algorithm at time k has access only to $\mathcal{U}^k$ and uses this set as the best estimate for $\mathcal{U}^{k+1}$. The fact that all three proposed schemes perform well despite this "double projection" (once algorithmically on $\mathcal{U}^k$ and once physically onto $\mathcal{U}^{k+1}$) is another indication of the robustness of feedback-based optimization.

At least from a theoretical perspective, the solution to this kind of problem are AWAs which were studied in Chapters 5 and 9. A simple discretization of an AWA of (15.1) is given by

$$\begin{aligned} u^{k+1} &= u^k - \alpha G^{\mathrm{I}}(u^k)^{-1}\Big(\tilde{H}(y^k,u^k)^T\nabla\left(\Phi+\rho\Psi_{\mathcal{X}}\right)(y^k,u^k)^T \\ &\qquad\qquad - \tfrac{1}{K}(u^k + P_{\mathcal{U}^k}(u^k))\Big) \\ y^{k+1} &= h\left(P_{\mathcal{U}^{k+1}}(u^{k+1})\right), \end{aligned}$$

where $K>0$ is the inverse anti-windup gain and $P_{\mathcal{U}^k}(u^k)$ is the (measured) saturated control input at iteration k. In particular, notice that the control algorithm does not require an explicit model for $\mathcal{U}^k$.

Although this type of implementation is theoretically well-motivated, at the time of writing, the AWA methods have not been tested on the power system examples. ■

15.2.2 Primal-Dual Scheme

In order to achieve asymptotic constraint satisfaction for the engineering constraints $\mathcal{X}$ (e.g., line current limits), we amend the penalty-relaxed partially projected gradient descent by integrating constraint violations

over time. This is achieved by the system

$$\begin{aligned} u^{k+1} &= P_{\mathcal{U}^k}^{G^{\mathrm{I}}(u^k)}\left(u^k - \alpha G^{\mathrm{I}}(u^k)^{-1}\tilde{H}(y^k,u^k)^T\Delta_k\right) \\ &\quad \text{where} \quad \Delta_k := \nabla\left(\Phi + \rho\Psi_{\mathcal{X}}\right)(y^k,u^k)^T + \mu^k\nabla g_{\mathcal{X}}(y^k,u^k) \\ \mu^{k+1} &= P_{\mathbb{R}^r_{\geq 0}}\left(\mu^k + \eta k(y^k,u^k)\right) \\ y^k &= h(u^k)\,, \end{aligned} \tag{15.2}$$

where $\eta > 0$ is a dual integration gain and $g_{\mathcal{X}} : \mathbb{R}^{2N} \times \mathbb{R}^r \to \mathbb{R}^r$ defines the engineering constraint set $\mathcal{X}$, i.e., $\mathcal{X} = \{(y,u) \,|\, g_{\mathcal{X}}(y,u) \leq 0\}$.

Hence, technically, (15.2) is a discretization of a projected saddle-point flow for the partial Lagrangian $L : \mathcal{U} \times \mathbb{R}^r_{\geq 0} \to \mathbb{R}$ given by

$$L(u,\mu) := \Phi(h(u),u) + \rho\tilde{\Psi}_{\mathcal{X}}(u) + \mu^T g_{\mathcal{X}}(h(u),u)\,.$$

More precisely, (15.2) is a discretization of

$$\begin{aligned} \dot{u} &= \Pi_{\mathcal{U}}^{G^{\mathrm{I}}}\left[-\alpha\,\mathrm{grad}_{G^{\mathrm{I}}}\,L(\cdot,\mu)\right](u) \\ \dot{\mu} &= \Pi_{\mathbb{R}^n_{\geq 0}}\left[\eta\,\mathrm{grad}_{\mathbb{I}}\,L(u,\cdot)\right](\mu)\,. \end{aligned} \tag{15.3}$$

Strictly speaking, no known results guarantee the convergence of (15.3) to a primal-dual solution pair of the underlying optimization problem. This lack of theoretical guarantees is primarily due to the non-convexity of h and the non-Euclidean metric G^{I}. For this reason, [29, 71, 244] use an additional regularization of the dual variables which modifies the equilibria, but allows for convergence (and tracking) guarantees even for non-convex problems.

We do not this type of dual regularization. Nevertheless, in our simulation context, and for small $\eta > 0$, we observe the desired convergence to local optimizer of the ACOPF problem, including asymptotic constraint satisfaction (which is not guaranteed by the penalty scheme).

15.2.3 Linearized Output Projection Scheme

As a final scheme, we consider the LOP discretization of the fully projected gradient flow (14.13). The implementation follows (6.5) without

any modification other than accounting for the time-varying nature of the problem. Namely, we implement

$$\begin{aligned} u^{k+1} &= u^k + \alpha\, \Sigma^{G^{\mathrm{I}}}_{\mathcal{C}^k}[-\operatorname{grad}_{G^{\mathrm{I}}}(\Phi + \rho\Psi_{\mathcal{X}})](u^k, \alpha) \\ y^{k+1} &= h\left(P_{\mathcal{U}^{k+1}}(u^{k+1})\right) , \end{aligned}$$

where we define

$$\begin{aligned} \Sigma^{G}_{\mathcal{C}^k}[f](u^k, \alpha) := \arg\min_{w \in \mathbb{R}^n} \quad & \left\| -G^{\mathrm{I}}(u^k)^{-1} \tilde{H}(y^k, u^k)^T \Delta_k - w \right\|^2_{G^{\mathrm{I}}(u^k)} \\ \text{subject to} \quad & u^k + \alpha w \in \mathcal{U}^k \\ & g_{\mathcal{X}}(y^k, u^k) + \alpha \tilde{H}(y^k, u^k)^T \nabla g(y^k, u^k) w \leq 0 \end{aligned}$$

with $\Delta_k := \nabla(\Phi + \rho\Psi_{\mathcal{X}})(y^k, u^k)^T$.

15.3 Simulation Results

Figures 15.5, 15.6, and 15.7 show the evolution of the various state variables for each of the three feedback-based optimization schemes applied to the 30-bus test setup.

The simulations were performed by computing the control value u^{k+1} for the respective scheme and evaluating $h\left(P_{\mathcal{U}^{k+1}}(u^{k+1})\right)$ using the AC power flow solver of the MATPOWER software package.

It is worth noting that parameter values are uniform across all three algorithms. In particular, the step-size α is the same for all three methods, and the penalty parameter ρ is the same for the penalty and the primal-dual scheme.

Remark 15.2. Interestingly, this type iterative simulations can also be interpreted in the framework of *optimization on manifolds* [4, 3] in the sense that we optimize over the manifold $\mathcal{M}^{\text{high}}$ embedded in a high dimensional vector. At every iteration, we compute a descent direction on the tangent space of $\mathcal{M}^{\text{high}}$. A finite step in this direction will generally result in a point that is not on $\mathcal{M}^{\text{high}}$. However, the computation of an AC power flow solution acts as a so-called *retraction*, which "projects" the point back onto $\mathcal{M}^{\text{high}}$.

Our setup, however, is more demanding than in [3] because we incorporate complex constraints (which is, so far, not within the scope of this field of research), and $\mathcal{M}^{\text{high}}$ is not a closed manifold. ■

15.3.1 Tracking Performance

We first consider how well each of the three feedback-based optimization schemes can track the solution of the ACOPF problem (13.2).

Although we cannot provide explicit theoretical tracking guarantees, we can see from Figure 15.8 that all of the three schemes manage to track the optimum in terms of cost almost perfectly. The only exception is the penalty approach, which achieves a lower-than-optimal cost from 17:00 to 21:30 by slightly violating a line flow and voltage constraint. This behavior is expected since the use of penalty functions implies that constraints can be violated both during run-time and upon convergence.

Apart from the economic cost and real power generation, the trajectories for the voltages and reactive power in-feed in Figures 15.5, 15.6, and 15.7 differ significantly from the optimal trajectory in Figure 15.4. This deviation is mainly because no cost is associated with reactive power generation, making different voltage/reactive power configurations almost equally costly.

This phenomenon should not be considered as a flaw of feedback-based schemes. Instead, comparing the voltage profiles at about 13:00, we can see the (offline) optimal solution is very volatile. In contrast, the all of the feedback schemes achieve relatively smooth profiles without sacrificing cost optimality. The latter behavior is preferable in practice.

Finally, note that the generator outage at 4:00 triggers a spike in output of Generator 1, which serves as the slack generator. A more realistic modeling framework including this type of contingency would lead to a response by multiple generators and would be better modeled by distributed slack generation. The critical thing to note is that all three feedback-based schemes experience this compensation by the slack bus. This unilateral response contrasts with the offline solution in Figure 15.4, which does not exhibit this spike in real power generation at the slack

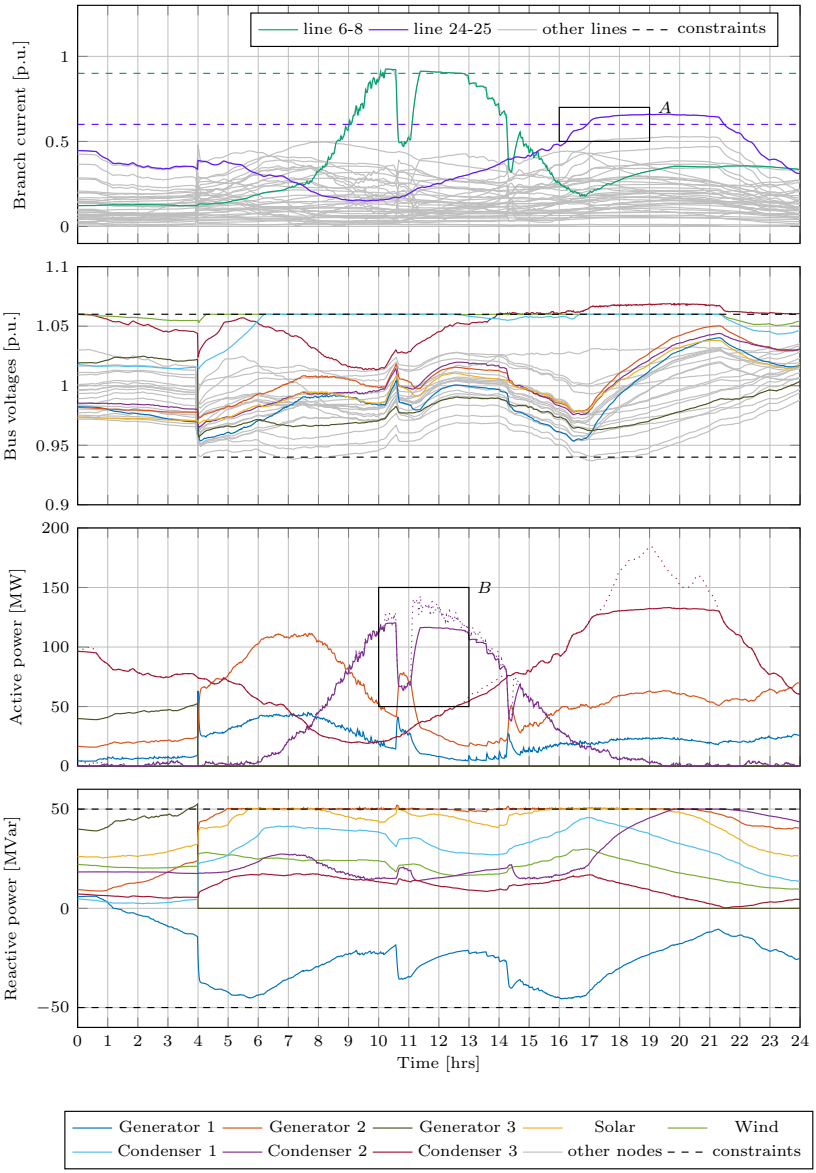

Figure 15.5: Penalty scheme applied to the 30-bus test case

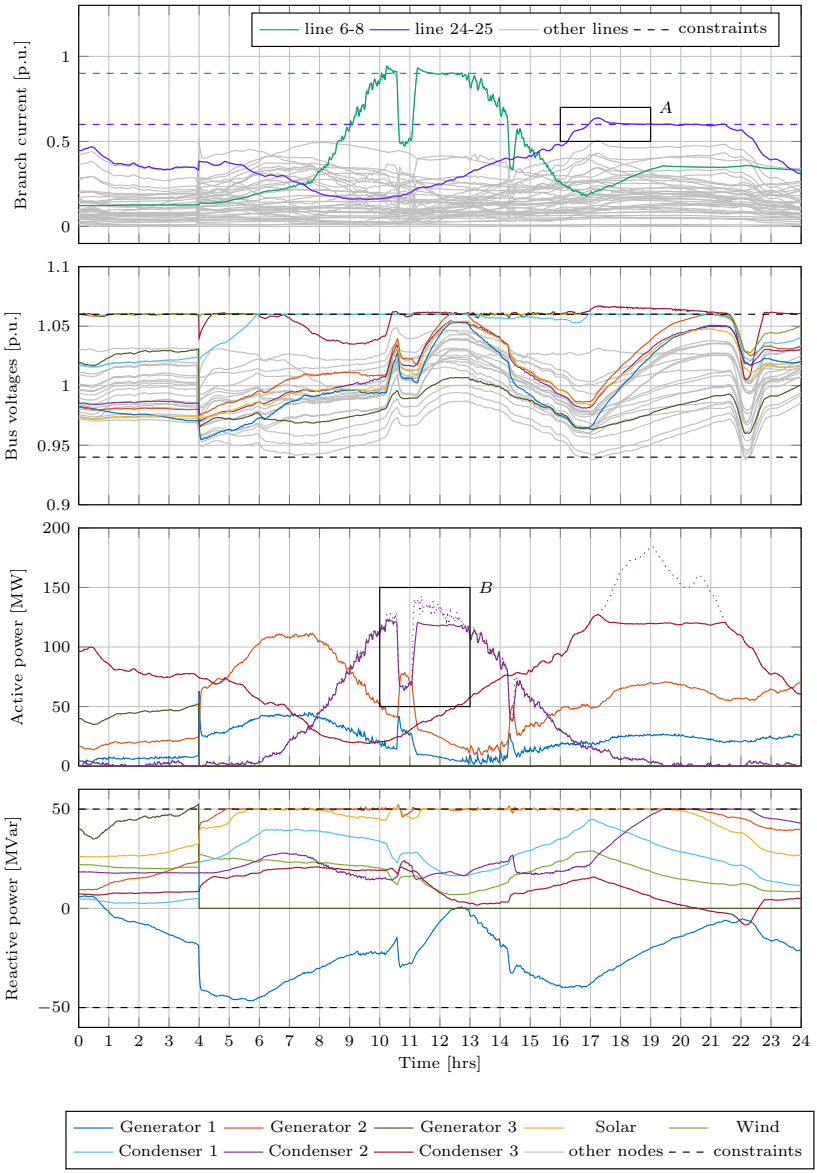

Figure 15.6: Primal-dual scheme applied to the 30-bus test case

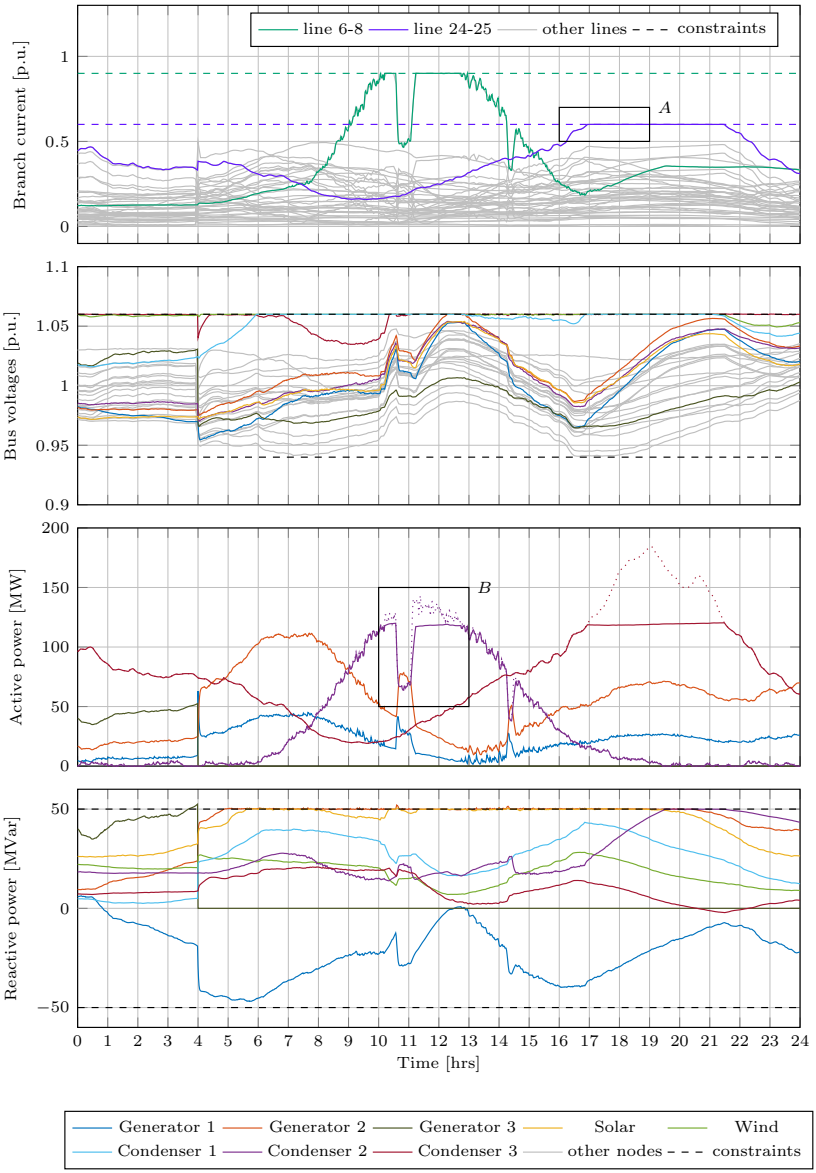

Figure 15.7: LOP scheme applied to the 30-bus test case

bus. This behavior is not a shortcoming of the feedback approach, but only due to the fact that the feedback schemes do not anticipate the generator outage.

15.3.2 Constraint Enforcement

The main difference between the proposed feedback-based optimization schemes lies in enforcing the engineering constraints $\mathcal{X}$. Figure 15.9 illustrates the characteristic behavior of each algorithm with respect to the line flow constraint for the line connecting Buses 24 and 25. This line is quickly at full capacity once the wind generation at Bus 25 ramps up around 17:00. Namely, Figure 15.9 shows Detail A[1] in Figures 15.5, 15.6, and 15.7.

As expected, the penalty scheme allows for a small but significant constraint violation up to the point where the effect of the penalty is strong enough to compensate for the cost and other penalty terms, e.g., from voltage constraints which are also active (i.e., violated) in the given time interval.

The primal-dual scheme improves this behavior by integrating the constraint violation and eventually driving it to zero, resulting in an asymptotically feasible state (in the time-invariant case). However, the tuning of the dual gain η is challenging and different types of constraints may require other gains. A control gain that is too small will not drive the constraint violation back to zero quickly enough, whereas a high gain leads to oscillations around the constraint surface.

Finally, the LOP scheme performs exceedingly well with almost no discernible constraint violations. This behavior is due to the fact that the LOP scheme anticipates the violation of output constraints based on the linearization around the current operating point and curtails the set-point update accordingly to avoid (or minimize) constraint violations in advance. Technically, constraint violations can happen, but the results from Chapter 6 indicate that they are bounded and vanish upon

[1] Details B and C are not discussed here because they are less interesting from a constraint enforcement perspective.

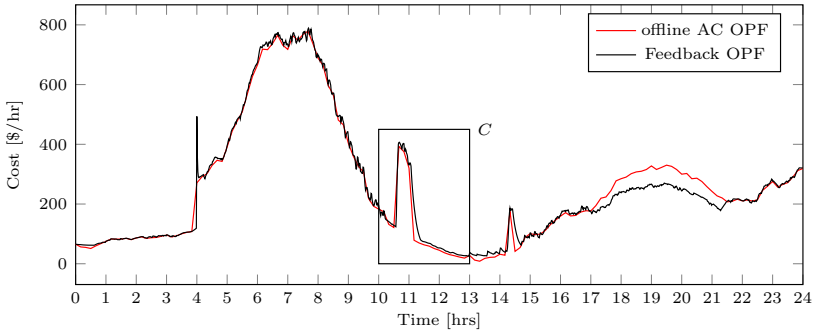

(a) Penalty scheme

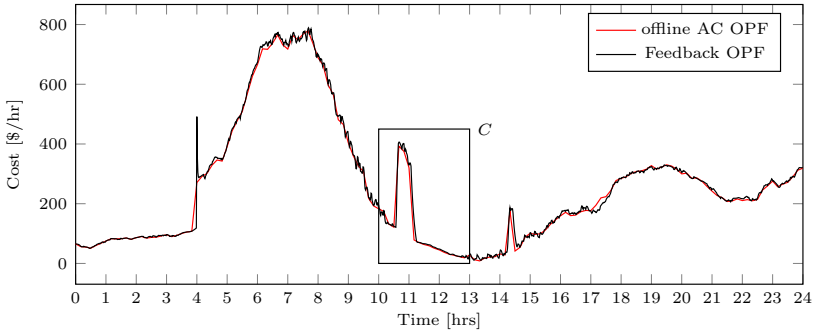

(b) Primal-Dual scheme

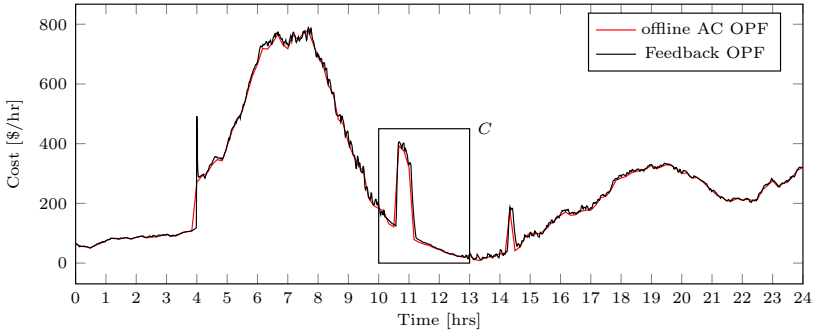

(c) LOP scheme

Figure 15.8: Cost value compared to a-posteriori offline solution

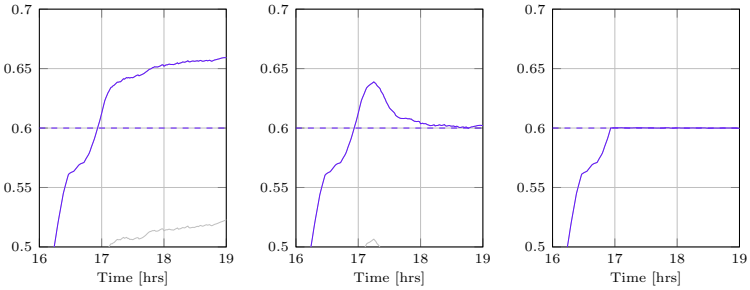

(a) Penalty-Relaxed (b) Dual Integration (c) LOP

Figure 15.9: Enforcement of line flow constraint for the feedback-based optimization schemes (Detail A in Figures 15.5, 15.6, and 15.7)

convergence.

15.3.3 Robustness against Time Delays

Next, we consider the effect of time delays in the feedback loop. For this purpose, we adopt a very simple model which applies the computed set-points with a delay of d_k iterations (i.e., d_k minutes). This delay is applied on top of the one iteration delay between measurement and implementing new set-points, which we have implicitly applied so far.

More precisely, to compute set-points for iteration $k+1$, each of the three proposed schemes has access only to $\mathcal{U}^{k-d_k}$, u^{k-d_k}, and y^{k-d_k}. For the primal-dual scheme, we assume that μ, as an internal variable, is not affected by the time delay, and μ^{k-1} is available. This assumption is motivated by the fact that μ is generally an internal variable of the controller, easily computed, and not subject to communication delays.

The simulation results for the three proposed schemes for different values of d_k are illustrated in Figures 15.10, 15.11, and 15.12 which show Details A, B, and C as defined before. The complete simulations can be found in Appendix B.

To assess the effect on constraint enforcement, we consider Detail A,

i.e., the line current around 17:00 from Bus 24 to Bus 25. Comparing the left columns of Figures 15.10, 15.11, and 15.12, we observe that the constraint enforcement of all three schemes is not significantly affected by the feedback delay.

The effect of time delays on the tracking in terms of the cost is best seen through Details B and C, which show how solar generation is recovering after a drop in availability (e.g., due to cloud cover). In this situation, all three schemes experience a slowed recovery of solar in-feed.

However, it should be stressed that the investigated model for time delays is very crude and does not reflect more complex situations encountered in practice such as heterogeneous delays. Nevertheless, the simulations suggest that the mechanisms for enforcing constraints suffer relatively little from time delays. In contrast, cost optimality seems to be more directly impacted.

15.3.4 Robustness against Model Uncertainty

Finally, in practical situations, it is not always possible to get an accurate estimate of the input-output sensitivities ∇h, which are required to compute the control action. Therefore, we study the effect of a linearization that does not correspond to the current state of the system.

Although an inexact linearization can have many different reasons (such as modeling uncertainty, noisy measurements, etc.) that produce an inexact state estimate, we limit ourselves to the use case where the linearization is simply not updated at every iteration. Instead, we will assume that the sensitivities are computed every 60 iterates (i.e., every 60 minutes) at every full hour and kept constant for the next 59 iterates.

The resulting behavior for the three proposed schemes is summarized in Figures 15.13, 15.14, and 15.15 with the complete simulations in Appendix B. Again, we are interested in the effect on both constraint enforcement and optimality.

As before, we observe that the deferred linearization does not impact the capacity to enforce output constraints for any of the three schemes in Detail A. All three methods display roughly the same behavior as

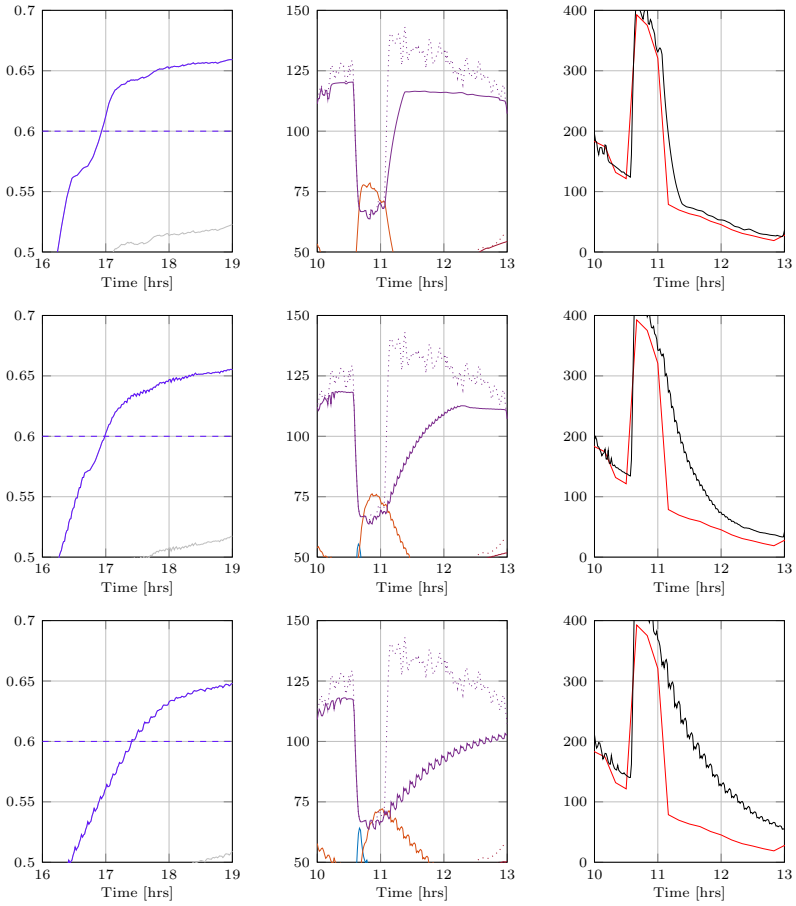

(a) Detail A (branch currents [p.u.])

(b) Detail B (active power [MW])

(c) Detail C (generation cost [$/hr])

Figure 15.10: Impact of communication delays on penalty scheme; from top to bottom: $d_k = \{0, 2, 5\}$ min

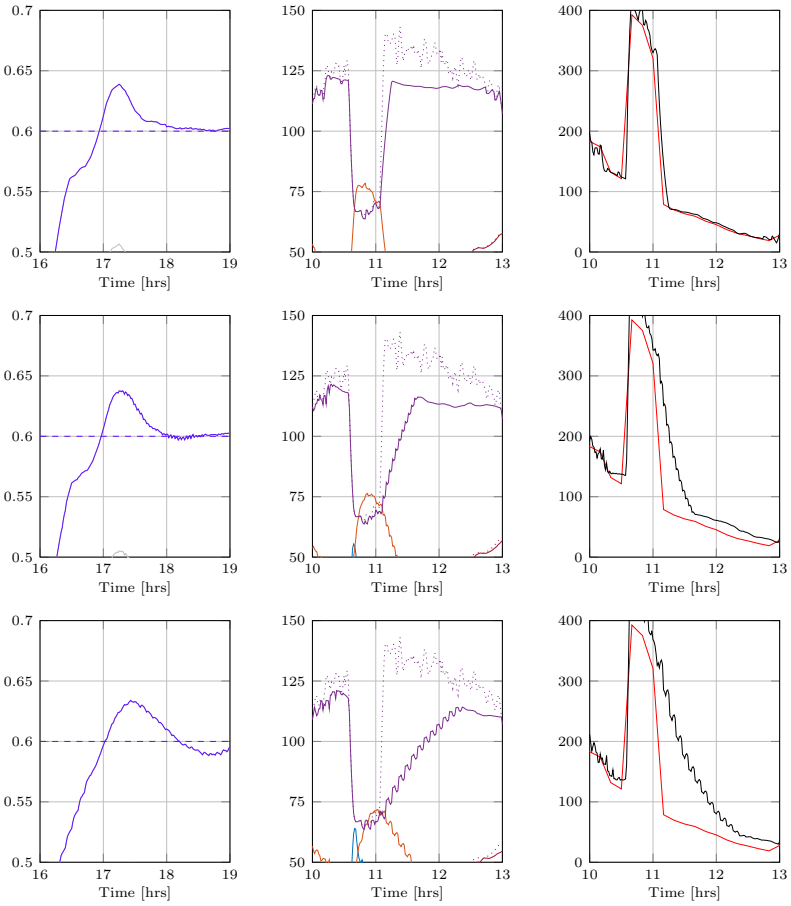

(a) Detail A (branch currents [p.u.])

(b) Detail B (active power [MW])

(c) Detail C (generation cost [$/hr])

Figure 15.11: Impact of communication delays on primal-dual scheme; from top to bottom: $d_k = \{0, 2, 5\}$ min

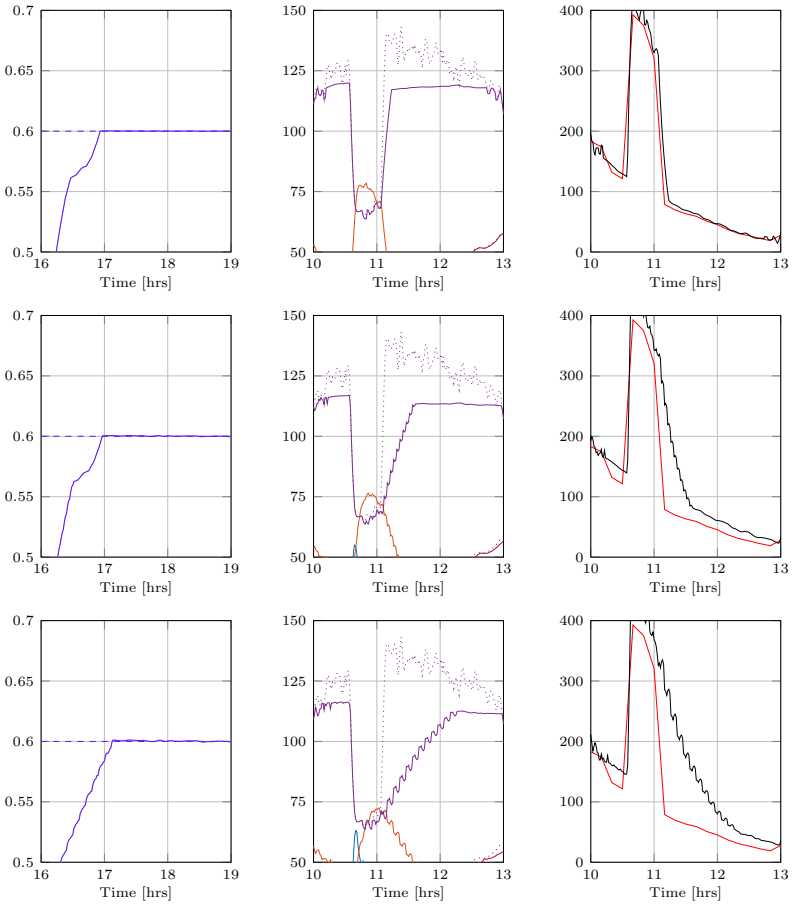

(a) Detail A (branch currents [p.u.])

(b) Detail B (active power [MW])

(c) Detail C (generation cost [$/hr])

Figure 15.12: Impact of communication delays on LOP scheme; from top to bottom: $d_k = \{0, 2, 5\}$ min

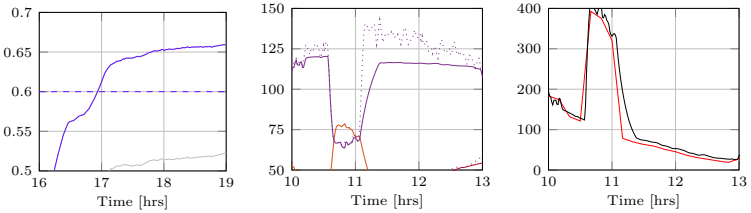

(a) Update of input-output sensitivities every 1min

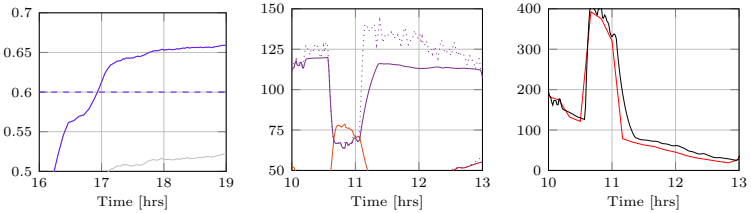

(b) Update of input-output sensitivities every 60min

Figure 15.13: Details of "lazy linearization" for penalty scheme

with the exact linearization. Although, as can be seen in Figure 15.16, the LOP scheme does suffer from temporary line flow limit violations between 11:00 and 13:00.

Concerning optimality, we observe that the penalty scheme is almost unaffected by the inexact linearization in response to the dip in solar power around 11:00. Solar power generation recovers quickly after the drop in availability.

However, the primal-dual and the LOP schemes are both impacted by an inexact linearization. Detail B provides a particularly illustrative example: Between 11:00 and 12:00, the optimal solution calls for an increase in the real power generation of the solar plant. However, both the primal-dual and the LOP scheme decrease solar in-feed again after an initial increase, thus remaining suboptimal (right-hand panels in Figures 15.14 and 15.15).

After an update of the input-output sensitivities at 12:00, both algorithms start increasing solar generation immediately. A look at the

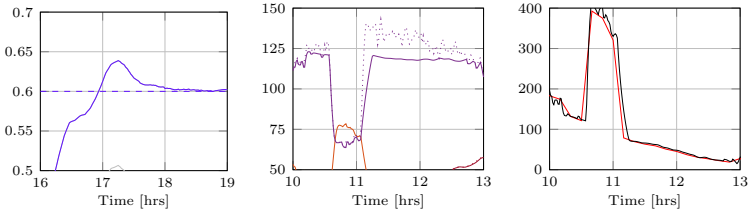

(a) Update of input-output sensitivities every 1min

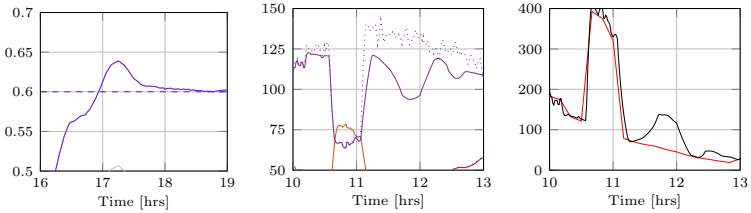

(b) Update of input-output sensitivities every 60min

Figure 15.14: Details of"lazy linearization" for primal-dual scheme

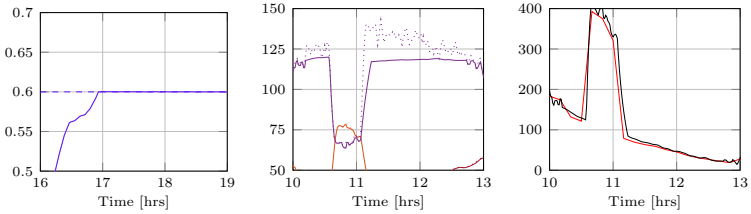

(a) Update of input-output sensitivities every 1 min

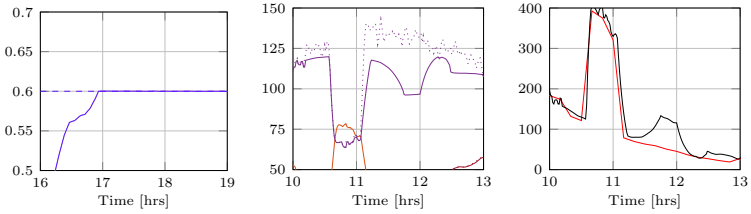

(b) Update of input-output sensitivities every 60 min

Figure 15.15: Details of "lazy linearization" for LOP scheme

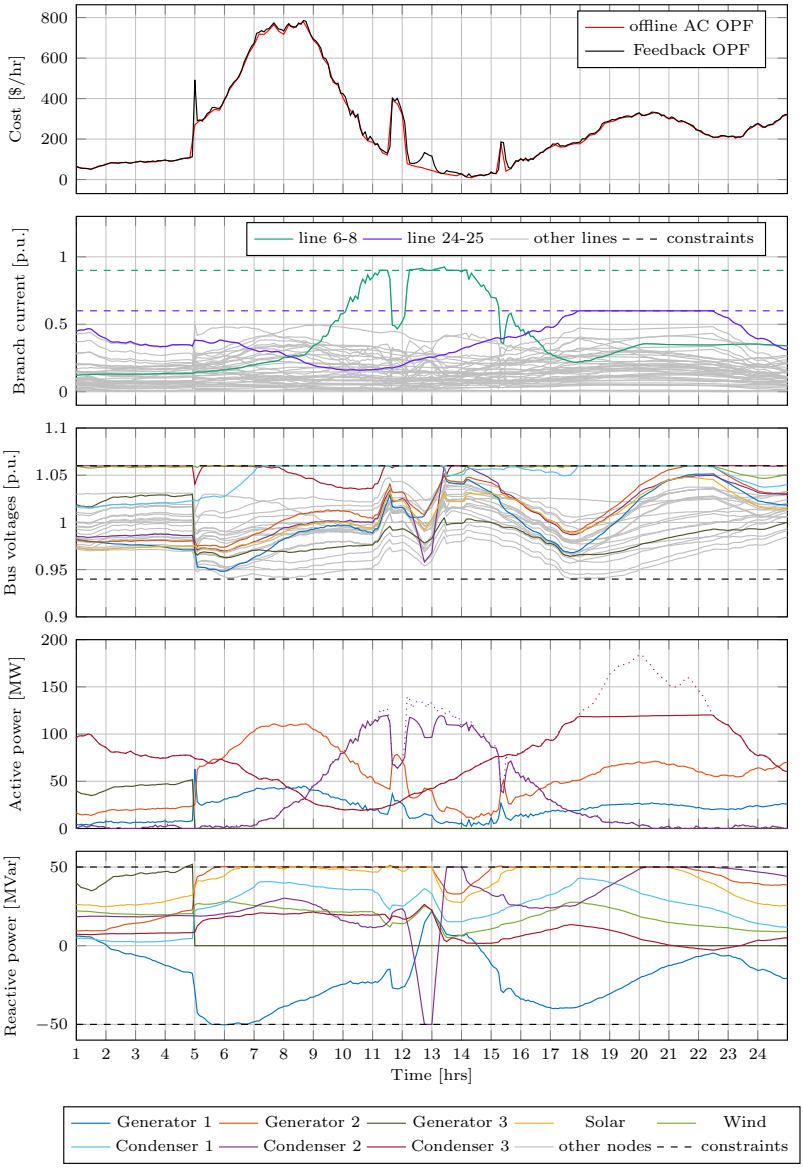

Figure 15.16: "lazy linearization" of LOP scheme (update every 60min)

reactive power in-feed and voltage in Figure 15.16 also reveals a complete reversal in reactive power generation at the solar plant right after 12:00.

The new linearization is critical for the system to leave a suboptimal equilibrium point and impacts the input-output sensitivities to the extent that the new equilibrium point far from the previous (inexact) equilibrium state.

15.4 Summary

Although the simulations presented in this chapter rely on a simplistic model for power systems, they illustrate the robustness of feedback-based optimization schemes based on projected gradient flows, irrespective of the particular discretization or implementation.

Besides the ability to track time-varying optimizers, these schemes also do well under time delays and inexact linearizations of the steady-state map. Although the tracking performance with respect to the a-posteriori optimal solution may suffer under these imperfections, the capacity to enforce constraints is retained by all of the proposed schemes to an extensive degree and under challenging conditions.

CHAPTER 16

Conclusions

In this thesis, we have studied novel, non-smooth multi-input-multi-output controllers to track solutions of constrained, non-convex optimization problems. We have backed these control designs with rigorous theoretical guarantees by starting from fundamental mathematical models for discontinuous dynamical systems. Their capability to enforce complicated constraints with little model information makes these new feedback-based optimization methods a promising approach for the real-time optimal operation of future power grids.

The control designs presented in this thesis are based on projected gradient flows due to their generality, robustness, and various implementation possibilities. To emphasize this fact, note the following:

- Projected gradient flows **work out-of-the-box**: Their trajectories converge to local minimizers for non-convex objectives and non-convex feasible sets under minimal technical assumptions, without regularization, and with little tuning (Chapters 4 and 8).
- Projected gradient flows, despite their non-smooth nature, can be **implemented as feedback controllers**. The concepts of anti-windup approximations and linearized output projection discretizations, introduced in this thesis, are two possibilities towards this end (Chapters 5, 6, 9, and 10).
- Projected gradient flows have many of the **same properties as**

unconstrained gradient flows. For example, when interconnected with a dynamical system, the same stability threshold applies to projected and unconstrained gradient flows (Chapter 11).

- Projected gradient flows are part of a set of methods for constrained optimization that can be **easily "mixed and matched"**, while still retaining many theoretical guarantees: We have seen three methods to enforce constraints (by penalty function, by dualization, and by projection). We have also identified the freedom to choose a metric, and we have studied anti-windup approximations and LOP discretizations as a means of implementing general projected dynamical systems (Chapter 14).

Nevertheless, many open problems and possibilities remain unexplored. Hence, in the following, we formulate open theoretical problems, possible power system modeling extensions, and potential use cases.

16.1 Open Theoretical Questions

Parts I and II provide extensive results on projected dynamical systems, their implementation as feedback controllers, and applications to nonlinear optimization. However, various cross-connections remain to be explored.

For instance, the combination of AWA and LOP schemes has not yet been studied but appears to be a plausible possibility since AWAs are used to enforce input constraints whereas LOP discretizations excel at enforcing output constraints. Furthermore, both can be analyzed with the same type of analysis based on perturbed differential inclusions.

Also, stability conditions analogous to Chapter 11 quantifying the need for timescale separation have yet to be developed for AWAs and LOP discretizations of projected gradient flows.

Extensions to time-varying problems, as presented in Chapter 12, are still in their early stage. This is mainly because perturbed sweeping processes have not yet reached the same level of generality as projected dynamical systems in terms of a coordinate-free formulation.

Furthermore, the robustness against time delays and model uncertainty, illustrated in Chapter 15, has yet to be analyzed theoretically.

Another aspect, very relevant in the power systems context, is incorporating discrete decisions through a hybrid systems formalism.

Finally, in this thesis, we have not touched upon stochastic aspects, e.g., the influence of process and measurement noise. But also, the design of stochastic control designs is a potential future research direction.

16.2 Power Systems Modeling Possibilities

Because the methods presented in this thesis apply to general problems of the form

$$\begin{aligned} \text{minimize} \quad & \Phi(u, y) \\ \text{subject to} \quad & y = h(u) \\ & u \in \mathcal{U} \\ & (u, y) \in \mathcal{X}\,, \end{aligned} \tag{16.1}$$

the simple power systems model used in Part III can be significantly generalized to incorporate more realistic objectives, constraints and actuation capabilities.

One relatively straightforward possibility is to consider **more realistic constraints** on power generation instead of simple box constraints. For instance, one could easily consider field and stator current limits for synchronous machines, whereas, for inverters, one can impose constraints on the power factor and apparent power output. Most of our technical assumptions on the constraint sets (such as Clarke regularity) would be satisfied for these models.

Another easy extension is the use of **more complicated load models**. For instance, loads $p^{\mathrm{L}}(v, \theta) + jq^{\mathrm{L}}(v, \theta)$ may be voltage-dependent, have a constant current component, or follow an exponential model. As long as ∇p^{L} and ∇q^{L} can be estimated, this knowledge can be incorporated into the control design.

Yet another possibility is the **control of FACTS devices**, such as

Static Var Compensators (SVC) or Thyristor Controlled Series Compensators (TCSC). These can be modeled as a variable shunt susceptance or line reactance, respectively, subject to engineering constraints. In both cases, it is possible to compute the sensitivities with respect to these variable grid parameters and incorporate them as inputs into the control design. Similarly, our framework also allows us to **model HVDC-links**. Thus, it becomes possible to use these devices to dynamically optimize the power flow in a transmission grid in real time.

Next, since we have not made any strong assumptions on the objective function, we can replace the generic cost of generation considered in Chapter 15 with more sophisticated objectives. Choices may include cost terms that directly **promote voltage stability**, e.g., by penalizing the distance to voltage collapse. As another example, in the restoration phase after a contingency, it might be necessary to align voltage angles to reconnect a line between two buses. This temporary objective can also be easily included in our problem setup.

All in all, the generality of (16.1) paves the way to consider many power system operation challenges and make more efficient use of modern devices to improve the reliability and optimality of power transmission and delivery.

16.3 Potential Power System Use Cases

Throughout Part III, we have adopted a simple, yet general, load flow model which can be used to model both transmission as well as distribution grids. However, the ideas in this thesis are presumably more interesting for transmission-level applications because transmission grids exhibit more challenging grid topologies, more complicated actuation capabilities, and, at the same time, reliable monitoring and state estimation is available. As of today, the active management of distribution systems, on the other hand, faces different challenges. Apart from economic considerations, a lack of monitoring infrastructure, privacy issues, and other aspects, the limited number of controllable devices does not, generally, motivate the deployment of a completely integrated control

infrastructure.

On the transmission level, a feedback-based optimization approach is well-suited to overhaul secondary and tertiary frequency control and include more complex functionality such as congestion-aware redispatch, thus **rethinking many classical ancillary services** [63]. Work in this direction is already well on its way (see the relevant references in [179]). However, methods developed in this thesis due to their generality and modeling possibilities (discussed in Section 16.2 above), are potentially instrumental to study even more complicated scenarios that have not been previously considered.

Another possible use case for closed-loop optimization is the **fast, coordinated voltage control in transmission systems**. Namely, the control of reactive power and voltages at the transmission level has been identified as a critical aspect of future power systems operations [61, 195]. Although several transmission system operators already perform this type of *secondary voltage control* [197], the concept is still not widely adopted or based on relatively simple models. The tools presented in this thesis are very well suited to provide a structured approach for designing such voltage control schemes and supporting them with theoretical statements.

Moreover, the **adequate response to contingencies and emergencies** is expected to become ever more crucial as the complexity of power systems grows [62]. In this context, a lack of situational awareness and subsequent poor decisions have been identified as a critical factor in recent emergencies and blackouts [64]. Therefore, it is crucial to provide operational strategies or even fully automated emergency systems that can address this challenge by enabling well-informed actions in the face of extreme events to sustain a highly reliable electricity supply. Again, feedback-based optimization approaches such as those presented in this thesis provide new but well-motivated possibilities for designing such strategies and systems.

Part IV

Appendices

APPENDIX A

Quantitative Sensitivity Bounds for Nonlinear Optimization Problems

In this chapter, we study how the sensitivities of optimizers for a general class of nonlinear optimization problems can be quantified. These types of results are particularly interesting for establishing tracking guarantees for time-varying optimization schemes as in Chapter 12.

In the following, using the classical sensitivity results [99, 139, 98, 210], we establish sufficient conditions under which the optimizer and its dual variables of parametrized nonlinear optimization problems are continuous functions of a perturbation (not necessarily time), and, more importantly we derive quantitative bounds on their Lipschitz constants. Moreover, we provide and discuss special cases of constrained parametrized optimization problems for which the sensitivities of optimizers and dual variables take a simple and easy-to-interpret form.

The results in this chapter build directly upon the definitions in Chapter 2 and, in particular, extend Section 2.2.1. In the remainder of the chapter, we particularly refer to paremtrized nonlinear optimization

This chapter is based on the first part of [Ha10]. The main technical results are Proposition A.1, Theorem A.2, and Corollary A.2. This work is a collaboration with I. Subotić.

problems of the form (2.3), i.e.,

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \Phi(x,\xi), \\ \text{subject to} \quad & h(x,\xi) = 0 \\ & g(x,\xi) \leq 0\,, \end{aligned} \tag{A.1}$$

where $\xi \in \Xi$ is a perturbation parameter and $\Xi \subset \mathbb{R}^o$. Furthermore, $\Phi : \mathbb{R}^n \times \Xi \to \mathbb{R}$, $h : \mathbb{R}^n \times \Xi \to \mathbb{R}^m$, and $g : \mathbb{R}^n \times \Xi \to \mathbb{R}^p$ are continuously differentiable in x.

Moreover, our results in this chapter will extend Theorem 2.3 [138, Thm. 2.3.2 & 2.3.3], which we repeat hereafter for convenience.

Theorem A.1. *Consider (A.1) and assume that Φ and g are twice continuously differentiable in x, and that Φ, h, g, $\nabla_x\Phi, \nabla_x h, \nabla_x g, \nabla^2_{xx}\Phi$, $\nabla^2_{xx}h$, and $\nabla^2_{xx}g$ are continuous in ξ. Further, let $\bar{x}^\star$ be a regular minimizer of (2.2) for $\bar{\xi}$ with multipliers $(\bar{\lambda}^\star, \bar{\mu}^\star)$. Then, on a neighborhood $\mathcal{N} \subset \Xi$ of $\bar{\xi}$, there exist continuous maps $x^\star : \mathcal{N} \to \mathbb{R}^n$, $\lambda^\star : \mathcal{N} \to \mathbb{R}^m$ and $\mu^\star : \mathcal{N} \to \mathbb{R}^p$ such that*

(i) $x^\star(\bar{\xi}) = \bar{x}^\star$, $\lambda^\star(\bar{\xi}) = \bar{\lambda}^\star$, and $\mu^\star(\bar{\xi}) = \bar{\mu}^\star$,

(ii) for all $\xi \in \mathcal{N}$, LICQ is satisfied at $x^\star(\xi)$, and $(x^\star(\xi), \lambda^\star(\xi), \mu^\star(\xi))$ is a KKT point for which SSOSC holds,

(iii) for all $\xi \in \mathcal{N}$, $x^\star(\xi)$ is a local minimizer for (2.2) and $(\lambda^\star(\xi), \mu^\star(\xi))$ are the corresponding Lagrange multipliers.

If, in addition, Φ, h and g are twice continuously differentiable in (x,ξ),

(iv) $x^\star, \lambda^\star$ and $\mu^\star$ are locally Lipschitz at $\bar{\xi}$.

The chapter is structured as follows: Appendix A.1 deals with quantifying the local differentiability and Lipschitz continuity of the solution maps for nonlinear programs like (A.1). Then, in Appendix A.2, we introduce and discuss assumptions that allow for global statements about solution maps and their sensitivities. Finally, these findings are discussed in Appendix A.3, using illustrative special cases.

A.1 Quantifying the Lipschitz Continuity of Solution Maps

Theorem A.1 and similar results in [139, 98, 210] guarantee the Lipschitz continuity of the solution map $x^\star$, however, they do not quantify the Lipschitz constants of $x^\star$ and $\mu^\star$.

Hence, in this section, we refine these results and give quantitative bounds on the rate of change of $x^\star$. The key insight for this purpose, also used in [139, 98], is that a KKT point $(x^\star, \mu^\star) \in \mathcal{X}(\xi) \times \mathbb{R}^m_{\geq 0}$ solves the system

$$F(x^\star, \mu^\star, \xi) := \begin{bmatrix} \nabla_x L(x^\star, \mu^\star, \xi)^T \\ h(x^\star, \xi) \\ \operatorname{diag}(\mu^\star) g(x^\star, \xi) \end{bmatrix} = \mathbb{0} \in \mathbb{R}^{n+p+m}. \tag{A.2}$$

Under additional assumptions, the implicit function theorem can be applied to (A.2), to express $(x^\star, \mu^\star)$ as a function of ξ.

Throughout the remainder of this section, we consider the same setup as in Theorem A.1. Hence, existence and local Lipschitz continuity of $x^\star : \mathcal{N} \to \mathbb{R}^n$ and $\mu^\star : \mathcal{N} \to \mathbb{R}^m$ for some neighborhood $\mathcal{N} \subset \Xi$ around $\bar{\xi}$, for which $(\bar{x}^\star, \bar{\mu}^\star)$ is a regular optimizer, are guaranteed. Further, we use the shorthand notations $\mathbf{I}^\star(\xi) := \mathbf{I}^\xi_{x^\star(\xi)}$, $\bar{\mathbf{I}}^\star(\xi) := \bar{\mathbf{I}}^\xi_{x^\star(\xi)}$, and, for all $\xi \in \mathcal{N}$, we define

$$\begin{aligned}
A(\xi) &:= \nabla^2_{xx} L(x^\star(\xi), \lambda^\star(\xi), \mu^\star(\xi), \xi)) \in \mathbb{R}^{n \times n}, \\
B(\xi) &:= \begin{bmatrix} \nabla_x h(x^\star(\xi), \xi) \\ \nabla_x g_{\mathbf{I}^\star(\xi)}(x^\star(\xi), \xi) \end{bmatrix} \in \mathbb{R}^{(p+|\mathbf{I}^\star|) \times n}, \\
L^\star(\xi) &:= \nabla^2_{\xi x} L(x^\star(\xi), \lambda^\star(\xi), \mu^\star(\xi), \xi) \in \mathbb{R}^{n \times r}, \\
G^\star(\xi) &:= \begin{bmatrix} \nabla_\xi h(x^\star(\xi), \xi) \\ \nabla_\xi g_{\mathbf{I}^\star(\xi)}(x^\star(\xi), \xi) \end{bmatrix} \in \mathbb{R}^{(p+|\mathbf{I}^\star|) \times r}.
\end{aligned}$$

If (2.2) does not have any equality constraints and $\mathbf{I}^\star(\xi) = \emptyset$ we follow the convention that $B(\xi) = 0$ and $G^\star(\xi) = 0$.

In the following, unless there is any ambiguity, we drop the argument from any map whose sole argument is ξ.

Finally, we make one simplifying assumption that can, presumably, be slightly relaxed in future work:

Assumption A.1. *In the context of Theorem A.1, the matrix $A(\xi)$ is positive definite for all $\xi \in \Xi$.*

In particular, Assumption A.1 is satisfied for convex optimization problems with strongly convex objective – though one can perceive also scenarios with weaker regularity assumptions. Furthermore, the assumption is common in multi-parametric programming, when the KKT system needs to be solved explicitly (e.g., [249]). This assumption will be made in forthcoming sections to provide global bounds on the Lipschitz constants of $x^\star$ and $(\lambda^\star, \mu^\star)$.

To quantify the Lipschitz constant of the solution map we have to distinguish between two cases related to constraints becoming active or inactive when varying ξ. Namely, if every active constraint is associated with a positive multiplier (which we refer to as *strict complementarity* below), these constraints remain active for small variations in ξ. In this case, the solution maps $x^\star$ and $\mu^\star$ are differentiable. However, when a constraint is active but its multiplier is zero, then differentiability is in general not guaranteed. In this case, we need to make a careful case distinction considering the constraint as being either active or inactive.

A.1.1 Sensitivity under Strict Complementarity

Given a KKT-point $(x^\star, \lambda^\star, \mu^\star)$ of (2.2) for a given $\xi \in \Xi$, we define the sets of *strongly* active inequality constraints as

$$\mathbf{I}^{\mathbf{s}}(x^\star, \mu^\star, \xi) := \{i \in \mathbf{I}^{\xi}_{x^\star} \mid \mu_i^\star > 0\}$$

and we say that $(x^\star, \lambda^\star, \mu^\star)$ satisfies *strict complementary slackness (SCS)* if $\mathbf{I}^{\mathbf{s}}(x^\star, \mu^\star, \xi) = \mathbf{I}^{\xi}_{x^\star}$, i.e., if all active inequality constraints are strongly active.

Under strict complementary slackness, the solution maps $x^\star$ and $(\lambda^\star, \mu^\star)$ are continuously differentiable in a neighborhood of $\bar{\xi}$. The following result gives an explicit expression for the respective derivatives.

Proposition A.1. *Consider the same setup as in Theorem A.1 and let SCS and Assumption A.1 hold at* $(\bar{x}^\star, \bar{\lambda}^\star, \bar{\mu}^\star)$*. Then, in a neighborhood* $\mathcal{N} \subset \Xi$ *of* $\bar{\xi}$*, the maps* $x^\star : \mathcal{N} \to \mathbb{R}^n$ *and* $(\lambda^\star, \mu^\star) : \mathcal{N} \to \mathbb{R}^p \times \mathbb{R}^m_{\geq 0}$ *are continuously differentiable with*

$$\begin{aligned} \nabla_\xi x^\star &= -\Pi A^{-1} L^\star - \Sigma B^\dagger G^\star \in \mathbb{R}^{n \times r}, \\ \nabla_\xi \begin{bmatrix} \lambda^\star \\ \mu^\star_{\mathbf{I}^\star} \end{bmatrix} &= {B^\dagger}^T A \Sigma \left(A^{-1} L^\star - B^\dagger G^\star \right) \in \mathbb{R}^{(p+|\mathbf{I}^\star|) \times r} \end{aligned} \tag{A.3}$$

(and $\nabla_\xi \mu^\star_{\bar{\mathbf{I}}^\star} = 0$*), where* $B^\dagger := B^T (B^T)^{-1} \in \mathbb{R}^{n \times (p+|\mathbf{I}^\star|)}$ *is the pseudoinverse of* B*, and* $\Sigma, \Pi \in \mathbb{R}^{n \times n}$ *are given by*

$$\Sigma := A^{-1} B^T (B A^{-1} B^T)^{-1} B \quad \text{and} \quad \Pi := \mathbb{I} - \Sigma .$$

If (2.2) has only inequality constraints and none of them are active at $\bar{\xi}$*, then we have* $B^\dagger G^\star = 0$*.*

For the proof we require the following matrix inversion lemma [30, Ch. 2.17]:

Lemma A.1. *Let* $A \in \mathbb{R}^{n \times n}$*,* $B, C^T \in \mathbb{R}^{n \times m}$ *and* $D \in \mathbb{R}^{m \times m}$*. If* A *is non-singular, then the inverse*

$$M^{-1} = \begin{bmatrix} M_1 & M_2 \\ M_3 & M_4 \end{bmatrix} \quad \text{of} \quad M := \begin{bmatrix} A & C \\ B & D \end{bmatrix}$$

exists if and only if $D - BA^{-1}C$ *is invertible. Then, we have*

$$\begin{aligned} M_1 &= A^{-1} + A^{-1} C (D - BA^{-1}C)^{-1} B A^{-1}, \\ M_2 &= -A^{-1} C (D - BA^{-1}C)^{-1}, \\ M_3 &= -(D - BA^{-1}C)^{-1} B A^{-1}, \\ M_4 &= (D - BA^{-1}C)^{-1}. \end{aligned}$$

Proof of Proposition A.1. We follow the same procedure as in [138, 98]. However, Assumption A.1, the matrix inversion Lemma A.1 in the appendix, and some attention to details allow us to derive the explicit expressions (A.3) for the derivative of the solution map.

For the moment let $z^\star = (x^\star, \lambda^\star, \mu^\star)$. The implicit function theorem [213, Ch. 9] applied to (A.2) states that $\nabla_\xi z^\star = -\left(\nabla_z F\right)^{-1} \nabla_\xi F$.

Hence, without loss of generality, assume that the first $|\mathbf{I}^\star|$ constraints are (strongly) active and, by SCS, the remaining constraints are inactive. Then, ∇F can be written as

$$\nabla_z F := \begin{bmatrix} A & B^T & \bar{B}^T \\ D_\mu B & \mathbb{0} & \mathbb{0} \\ \mathbb{0} & \mathbb{0} & D_{\bar{g}} \end{bmatrix}, \quad \nabla_\xi F := \begin{bmatrix} L^\star \\ D_\mu G^\star \\ \mathbb{0} \end{bmatrix},$$

where we partitioned and regrouped the matrices according to the ordering $(x^\star, [\lambda^\star, \mu^\star_{\bar{\mathbf{I}}^\star}], \mu^\star_{\bar{\mathbf{I}}^\star})$, and used $\bar{B}(\xi) := \nabla_x g_{\bar{\mathbf{I}}^\star(\xi)}(x^\star(\xi), \xi)$. Further, we have used $D_\mu := \operatorname{diag}([\mathbb{1}_p^T \ \mu^{\star T}_{\bar{\mathbf{I}}^\star}])$ and $D_{\bar{g}} := \operatorname{diag}(g_{\bar{\mathbf{I}}^\star})$. Denote the upper left part of $\nabla_z F$ by

$$M := \begin{bmatrix} A & B^T \\ D_\mu B & \mathbb{0} \end{bmatrix}.$$

Then, using Lemma A.1 (and assuming for the moment that M is invertible), we have that

$$\left(\nabla_{(x,\mu)} F\right)^{-1} = \begin{bmatrix} M^{-1} & \star \\ \mathbb{0} & \star \end{bmatrix},$$

where $\star$ denotes non-zero, but irrelevant components.

Hence, we can already conclude from the expression for $\nabla_\xi z^\star$ that $\nabla_\xi \mu^\star_{\bar{\mathbf{I}}^\star} = 0$ for the inactive constraints.

Note that A is invertible by Assumption A.1, and the Schur complement of M given by $-D_\mu B A^{-1} B^T$ is invertible because D_μ in invertible by SCS. Hence, M is nonsingular, and its inverse can be computed according to Lemma A.1 as

$$M^{-1} = \begin{bmatrix} M_1 & M_2 \\ M_3 & M_4 \end{bmatrix},$$

where

$$\begin{aligned}
M_1 &= \left(\mathbb{I} - A^{-1}B^T(D_\mu BA^{-1}B^T)^{-1}D_\mu B\right)A^{-1} \\
&= (\mathbb{I} - A^{-1}B^T(BA^{-1}B^T)^{-1}B)A^{-1} = \Pi A^{-1}\,, \\
M_2 &= A^{-1}B^T(D_\mu BA^{-1}B^T)^{-1} = \\
&= A^{-1}B^T(BA^{-1}B^T)^{-1}D_\mu^{-1} \\
&= A^{-1}B^T(BA^{-1}B^T)^{-1}D_\mu^{-1}\underbrace{D_\mu B^\dagger D_\mu^{-1}}_{=\mathbb{I}} = \Sigma B^\dagger D_\mu^{-1}\,, \\
M_3 &= (D_\mu BA^{-1}B^T)^{-1}D_\mu BA^{-1} \\
&= \underbrace{B^{\dagger T}A^{-1}B^T}_{\mathbb{I}}(BA^{-1}B^T)^{-1}BA^{-1} = B^{\dagger T}A\Sigma A^{-1}\,, \\
M_4 &= -(D_\mu BA^{-1}B^T)^{-1} = -(BA^{-1}B^T)^{-1}D_\mu^{-1} \\
&= -B^{\dagger T}A^{-1}B^T(BA^{-1}B^T)^{-1}D_\mu^{-1}D_\mu B^\dagger D_\mu^{-1} \\
&= -B^{\dagger T}A\Sigma B^\dagger D_\mu^{-1}.
\end{aligned}$$

Finally, recalling $\nabla_\xi z^\star = -\left(\nabla_z F\right)^{-1}\nabla_\xi F$, we get

$$\begin{bmatrix} \nabla_\xi x^\star \\ \nabla_\xi \lambda^\star \\ \nabla_\xi \mu^\star_{\mathbf{I}^\star} \end{bmatrix} = -\begin{bmatrix} M^{-1} & \star \end{bmatrix}\nabla_\xi F$$

which yields the desired expressions. □

We defer an in-depth discussion of (A.3) to Appendix A.3, where we explore special cases. One special feature of (A.3) that we require though is the fact that Π and Σ can be interpreted as (oblique) projections onto the kernel of B and its orthogonal complement, respectively. Consequently, exploiting non-expansivity, their norm can be bounded independently of B:

Lemma A.2. *Consider the setup of Proposition A.1. For all $\xi \in \mathcal{N}$, it holds that*

$$\max\{\|\Pi(\xi)\|, \|\Sigma(\xi)\|\} \leq \sqrt{\lambda^{\max}_{A(\xi)}/\lambda^{\min}_{A(\xi)}}\,.$$

Proof. As before, we drop the argument ξ. Consider

$$\underset{w}{\text{minimize}}\ \tfrac{1}{2}(v-w)^T A(v-w)\ \text{ subject to } Bw = 0\,. \tag{A.4}$$

Since A is positive definite for all $\xi \in \Xi$ by Assumption A.1, (A.4) can be rewritten as

$$\underset{\tilde{w}}{\text{minimize}}\quad \tfrac{1}{2}\|Qv - \tilde{w}\|^2 \quad \text{subject to} \quad \tilde{B}\tilde{w} = 0 \tag{A.5}$$

where $Q := A^{1/2}$ is the unique symmetric positive definite matrix such that $A = Q^2$, $\tilde{w} := Qw$, and $\tilde{B} := BQ^{-1}$.

Since (A.5) is a Euclidean projection onto the kernel of $\tilde{B}$, its solution can be stated explicitly as [175, eq. 5.13.3]

$$\begin{aligned}\tilde{w}^\star &= Qv - \tilde{B}^T(\tilde{B}\tilde{B}^T)^{-1}\tilde{B}Qv \\ &= Q(\mathbb{I} - A^{-1}B^T(BA^{-1}B^T)^{-1}B)v = Q\Pi v\,.\end{aligned}$$

By non-expansivity of an orthogonal projection onto a subspace, we have $\|\tilde{w}^\star\| \le \|Qv\|$, and therefore we can write

$$\lambda_Q^{\min}\|\Pi v\| \le \|Q\Pi v\| = \|\tilde{w}^\star\| \le \|Qv\| \le \lambda_Q^{\max}\|v\|\,.$$

Rearranging the leftmost and rightmost terms together with the fact that $\lambda_A^{\max} = (\lambda_Q^{\max})^2$ and $\lambda_A^{\min} = (\lambda_Q^{\min})^2$ yields the desired bound for $\|\Pi\|$.

The projection on the orthogonal complement of $\ker \tilde{B}$ is given by $\Sigma := \mathbb{I} - \Pi$ [175, eq. 5.13.6]. The same bound as for $\|\Pi\|$ holds for $\|\Sigma\|$ since Σ is also non-expansive. □

Lemma A.2 allows us to bound (A.3) and thereby establish bounds for the local Lipschitz constants of $x^\star$ and $\mu^\star$ at $\bar{\xi}$.

Corollary A.1. *Consider the setup of Proposition A.1. Then,*

$$\left\|\nabla_\xi x^\star(\bar{\xi})\right\| \le \ell_{x^\star} \quad \textit{and} \quad \left\|\nabla_\xi \begin{bmatrix}\lambda^\star \\ \mu^\star\end{bmatrix}\right\| \le \ell_{(\lambda^\star,\mu^\star)}$$

holds for all $\xi \in \mathcal{N}$ *where*

$$\begin{aligned} \ell_{x^\star} &:= \sqrt{\frac{\lambda^{\max}_{A(\bar{\xi})}}{\lambda^{\min}_{A(\bar{\xi})}}} \left(\frac{\|L^\star(\bar{\xi})\|}{\lambda^{\min}_{A(\bar{\xi})}} + \frac{\|G^\star(\bar{\xi})\|}{\sigma^{\min}_{B^T(\bar{\xi})}} \right), \\ \ell_{(\lambda^\star,\mu^\star)} &:= \frac{{\lambda^{\max}_{A(\bar{\xi})}}^{3/2}}{\sigma^{\min}_{B^T(\bar{\xi})} {\lambda^{\min}_{A(\bar{\xi})}}^{1/2}} \left(\frac{\|L^\star(\bar{\xi})\|}{\lambda^{\min}_{A(\bar{\xi})}} + \frac{\|G^\star(\bar{\xi})\|}{\sigma^{\min}_{B^T(\bar{\xi})}} \right). \end{aligned} \tag{A.6}$$

Proof. Using singular value decomposition $B^T = V^T \Lambda^T U$, where B^T (i.e., Λ^T) has a full rank since the columns of B are linearly independent. Next we note that $B^\dagger = V^T \Lambda^\dagger U$ also has a full rank and corresponds to the right pseudoinverse of B, where $\Lambda^\dagger$ is the right pseudoinverse of Λ. Comparing the expressions for B^T and $B^\dagger$ we have that for the minimal and maximal singular values $\sigma^{\max}_{B^\dagger} = 1/\sigma^{\min}_{B^T}$ and $\sigma^{\min}_{B^\dagger} = 1/\sigma^{\max}_{B^T}$ hold. Consequently, $\|B^\dagger\| = 1/\sigma^{\min}_{B^T(\bar{\xi})}$ holds. Finally, the bounds in (A.6) follows from applying the triangle inequality, Cauchy-Schwarz, and Lemma A.2 to (A.3). □

A.1.2 Sensitivity of General Regular Optimizers

If the SCS assumption in Proposition A.1 does not hold, then at least one constraint i is *weakly active*, i.e., $g_i(x^\star(\xi)) = 0$ but $\mu_i^\star(\xi) = 0$. Weakly active constraints occur, for example, when the unconstrained minimizer of f lies exactly on a constraint surface. In this case, the constraint is active, but is not "pushing" the optimizer inwards.

Whenever we vary ξ and a constraint changes from being strongly active to inactive (or vice versa), the constraint is momentarily weakly active for some ξ. This is a consequence of the continuity of $(x^\star, \mu^\star)$.

For values of ξ, for which some constraints are weakly active, the map $(x^\star, \lambda^\star, \mu^\star)$ might not be differentiable, i.e., (A.3) might not apply.[1] Consequently, the best we can hope for in the general case where SCS does not hold, is a bound on the Lipschitz constant of $(x^\star, \lambda^\star, \mu^\star)$.

[1] However, it is generally possible to compute directional derivatives of $(x^\star, \lambda^\star, \mu^\star)$ at any point in any direction $\Delta\xi$.

To deal with weakly active constraints, we follow the approach in [138]. For this purpose, consider the same setup as in Proposition A.1 and let $(\bar{x}^\star, \bar{\lambda}^\star, \bar{\mu}^\star)$ be a regular optimizer for (2.2) for $\bar{\xi} \in \Xi$, satisfying LICQ at $\bar{x}^\star$, KKT, and SSOSC but not necessarily SCS. Further, let $\mathbf{R}$ be any index set satisfying

$$\mathbf{I}^{\mathbf{s}}(\bar{x}^\star, \bar{\mu}^\star, \bar{\xi}) \subset \mathbf{R} \subset \mathbf{I}^\star(\bar{\xi})\,, \tag{A.7}$$

that is, $\mathbf{R}$ includes all strongly active constraints and an arbitrary subset of weakly active constraints at $(\bar{x}^\star, \bar{\lambda}^\star, \bar{\mu}^\star, \bar{\xi})$. Next, consider the equality-constrained problem

$$\begin{aligned} \text{minimize} \quad & f(x, \xi) \\ \text{subject to} \quad & h(x, \xi) = 0 \\ & g_{\mathbf{R}}(x, \xi) = 0\,. \end{aligned} \tag{A.8}$$

By considering (A.8) we make a choice for every constraint that is weakly active at $(\bar{x}^\star, \bar{\lambda}^\star, \bar{\mu}^\star)$ to consider it like a strongly active constraint (by including it in $\mathbf{R}$) or to ignore it.

Note that Proposition A.1 and Corollary A.1 can be applied to (A.8). In particular, the SCS condition is vacuous because (A.8) does not have any inequality constraints. Hence, let $x^\star_{\mathbf{R}} : \mathcal{N}_{\mathbf{R}} \to \mathbb{R}^n$ and $(\lambda^\star_{\mathbf{R}}, \mu^\star_{\mathbf{R}}) : \mathcal{N}_{\mathbf{R}} \to \mathbb{R}^{(p+|\mathbf{R}|)}$ denote the primal and dual solution map of (A.8) in a neighborhood $\mathcal{N}_{\mathbf{R}}$ of $\bar{\xi}$.

In the proof of [138, Thm 2.3.2] it was established[2] that for every ξ in a neighborhood $\mathcal{N} \subset \bigcap_{\mathbf{R}} \mathcal{N}_{\mathbf{R}}$ of $\bar{\xi}$ the following holds: If the point $(x^\star(\xi), \lambda^\star(\xi), \mu^\star(\xi))$ is a regular optimizer of (2.2), then there exists an index set $\mathbf{R}$ satisfying (A.7) such that $(x^\star(\xi), \lambda^\star(\xi), \mu^\star_{\mathbf{R}}(\xi))$ is a regular optimizer of (A.8). In particular, we have $x^\star(\xi) = x^\star_{\mathbf{R}}(\xi)$, $\lambda^\star(\xi) = \lambda^\star_{\mathbf{R}}(\xi)$, and $\mu^\star_i(\xi) = 0$ for all $i \notin \mathbf{R}$.

Note that the set $\mathbf{R}$ for which this equivalence of solutions between (2.2) and (A.8) holds, depends on ξ and there does in general not exist a single set $\mathbf{R}$ that works for the entire neighborhood $\mathcal{N}$ of $\bar{\xi}$.

[2] To establish this equivalence the author passes through an exact penalty reformulation, which is beyond the scope of this chapter.

This key insight can be used to establish bounds on the Lipschitz constants of $x^\star$ and $\mu^\star$ at $\bar{\xi}$ in the absence of SCS:

Theorem A.2. *Consider the same setup as in Theorem A.1 and let Assumption A.1 hold at $(\bar{x}^\star, \bar{\lambda}^\star, \bar{\mu}^\star)$. Then, $x^\star : \mathcal{N} \to \mathbb{R}^n$ and $(\lambda^\star, \mu^\star) : \mathcal{N} \to \mathbb{R}^{p+m}$ are locally Lipschitz at $\bar{\xi}$ with bounds for the Lipschitz constants given by $\ell_{x^\star}$ and $\ell_{(\lambda^\star,\mu^\star)}$ as in (A.6).*

The difference in (A.6) for Corollary A.1 and Theorem A.2 lies in the fact that for Corollary A.1 the set $\mathbf{I}^\star$ (required for $A, B, L^\star$ and $G^\star$) contains only strictly active constraints, whereas for Theorem A.2 it may also contain weakly active constraints.

Proof. Given $\mathbf{R}$ satisfying (A.7), let $x^\star_{\mathbf{R}} : \mathcal{N} \to \mathbb{R}^n$, $\lambda^\star_{\mathbf{R}} : \mathcal{N} \to \mathbb{R}^m$, and $\mu^\star_{\mathbf{R}} : \mathcal{N} \to \mathbb{R}^{|\mathbf{R}|}$ the solution maps of (A.8) in a neighborhood $\mathcal{N}$ of $\bar{\xi}$. Consequently, we apply Corollary A.1 to (A.8) for every $\mathbf{R}$ satisfying (A.7), and we let $\ell^{\mathbf{R}}_{x^\star}$ and $\ell^{\mathbf{R}}_{(\lambda^\star,\mu^\star)}$ denote the respective bounds on the Lipschitz constant of $x^\star_{\mathbf{R}}$ and $(\lambda^\star_{\mathbf{R}}, \mu^\star_{\mathbf{R}})$ at $\bar{\xi}$ according to (A.6). Similarly, let $A_{\mathbf{R}}(\bar{\xi})$, $B_{\mathbf{R}}(\bar{\xi})$, $L^\star_{\mathbf{R}}(\bar{\xi})$ and $G^\star_{\mathbf{R}}(\bar{\xi})$ denote the corresponding quantities for (A.8), evaluated at $\bar{\xi}$.

Since, for every $\xi \in \Xi$ the solution $(x^\star(\xi), \mu^\star(\xi))$ is equal to the point $(x^\star_{\mathbf{R}}(\xi), \mu^\star(\xi))$ for some $\mathbf{R}$ satisfying (A.7), we can upper bound the local Lipschitz constants of $x^\star$ and $(\lambda^\star, \mu^\star)$ at $\bar{\xi}$ by maximizing over all (finite) possibilities of $\mathbf{R}$ satisfying (A.7). Namely, for $x^\star$ we have

$$\sup_{\xi \to \bar{\xi}} \tfrac{\|x^\star(\xi) - x^\star(\bar{\xi})\|}{\|\xi - \bar{\xi}\|} \leq \max_{\mathbf{R}} \sup_{\xi \to \bar{\xi}} \tfrac{\|x^\star_{\mathbf{R}}(\xi) - x^\star(\bar{\xi})\|}{\|\xi - \bar{\xi}\|} \leq \max_{\mathbf{R}} \ell^{\mathbf{R}}_{x^\star} ,$$

and the case for $(\lambda^\star, \mu^\star)$ follows analogously.

In particular, we claim that this maximum is achieved exactly for the choice $\mathbf{R} = \mathbf{I}^\star(\bar{\xi})$ where all strongly and weakly active constraints are considered.

To see this, note the following: Since lifting a vector to higher dimensions by adding components can only increase its 2-norm, we have $\|G^\star_{\mathbf{R}}(\bar{\xi})\| \leq \|G^\star_{\mathbf{I}^\star(\bar{\xi})}(\bar{\xi})\|$ for any $\mathbf{R}$ satisfying (A.7). Analogously, we have $\sigma^{\min}_{B^T_{\mathbf{R}}(\bar{\xi})} \geq \sigma^{\min}_{B^T_{\mathbf{I}^\star(\bar{\xi})}}$ since adding columns to B^T can only reduce its minimum singular value.

Next, at $\bar{\xi}$ we that $\mu^\star(\bar{\xi}) = \mu^\star_{\mathbf{R}}(\bar{\xi})$ for any $\mathbf{R}$ satisfying (A.7). Further, since $\mu^\star_i(\bar{\xi}) = 0$, for all weakly active constraints with index i, we have $L^\star(\xi) = L^\star_{\mathbf{R}}(\xi) = L^\star_{\mathbf{I}^\star(\bar{\xi})}(\xi)$ and $A(\bar{\xi}) = A_{\mathbf{R}}(\bar{\xi}) = A_{\mathbf{I}^\star(\bar{\xi})}(\bar{\xi})$. In other words, at $\bar{\xi}$ the weakly active constraints do not affect the quantities A and $L^\star$ which are derivatives of the Lagrangian.

Applying these facts to (A.6), it follows that, for all $\mathbf{R}$ satisfying (A.7), it holds that $\ell^{\mathbf{R}}_{x^\star} \leq \ell^{\mathbf{I}^\star(\bar{\xi})}_{x^\star}$ and $\ell^{\mathbf{R}}_{(\lambda^\star,\mu^\star)} \leq \ell^{\mathbf{I}^\star(\bar{\xi})}_{(\lambda^\star,\mu^\star)}$ and thus we have $\ell_{x^\star} = \max_{\mathbf{R}} \ell^{\mathbf{R}}_{x^\star} = \ell^{\mathbf{I}^\star(\bar{\xi})}_{x^\star}$ and $\ell_{(\lambda^\star,\mu^\star)} = \max_{\mathbf{R}} \ell^{\mathbf{R}}_{(\lambda^\star,\mu^\star)} = \ell^{\mathbf{I}^\star(\bar{\xi})}_{(\lambda^\star,\mu^\star)}$ which completes the proof. □

A.2 Conditions for Global Solution Maps

Theorem A.2 provides bounds on the local Lipschitz constants for $x^\star$ and $\mu^\star$ at $\bar{\xi}$. Next, we provide a set of assumptions to give global bounds on the sensitivity. Naturally, some of these assumptions are restrictive, but they provide important intuition and use cases for the general result in Theorem A.2.

For simplicity, instead of (2.2), we focus on problems with only inequality constraints which require a more careful investigation than equality constraints as shown in the previous analysis, i.e., we consider

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f(x,\xi) \\ \text{subject to} \quad & g(x,\xi) \leq 0\,, \end{aligned} \tag{A.9}$$

for the remainder of this section. All of the following statements and assumptions can be generalized.

First, it is necessary that (A.9) admits a solution for every ξ:

Assumption A.2. *The problem* (A.9) *is feasible for all* $\xi \in \Xi$.

Next, in order to guarantee that the solution map $x^\star : \Xi \to \mathbb{R}^n$ is single-valued for all $\xi \in \Xi$, it is convenient, if not necessary, to assume (strong) convexity:

Assumption A.3. *For (A.9), let f be α-strongly convex and have a β-Lipschitz gradient $\nabla_x f$ in x for all $\xi \in \Xi$. Further, for all $i = 1, \dots, m$, let g_i be convex and $\nabla_x g$ be ℓ_i-Lipschitz in x for all $\xi \in \Xi$.*

Assumption A.3 guarantees not only that (2.2) admits a unique optimizer for every $\xi \in \Xi$, it also implies that any KKT point of (A.9) (trivially) satisfies the SSOSC since $\nabla^2_{xx} L$ is positive definite for all $(x, \mu, \xi) \in \mathbb{R}^n \times \mathbb{R}^m_{\geq 0} \times \Xi$. Assumption A.3 also implies the lower bound $\lambda^{\min}(A(\xi)) \geq \alpha$ for all $\xi \in \Xi$ and thus replaces Assumption A.1. Lipschitz continuity of $\nabla_x f$ and $\nabla_x g$ is required to upper bound $\lambda^{\max}(A)$ as discussed below. Moreover, we generally require the active constraints to have uniformly full rank:

Assumption A.4. *Given (A.9), there exists $\omega > 0$ such that, for all $\xi \in \Xi$ and all $x \in \mathcal{X}(\xi)$ such that $\mathbf{I}^\xi_x \neq \emptyset$, we have*

$$\omega^2 \mathbb{I} \preceq \nabla_x g_{\mathbf{I}^\xi_x}(x, \xi) \nabla_x g_{\mathbf{I}^\xi_x}(x, \xi)^T .$$

In particular, Assumption A.4 guarantees that LICQ is satisfied at all $x \in \mathcal{X}(\xi)$, for all $\xi \in \Xi$, and that $\sigma^{\min}_{B^T(\xi)} \geq \omega$ in (A.6) is lower bounded away from zero whenever at least one constraint is active. Note that Assumption A.4 is independent of the choice of the cost function.

Further note that, under Assumptions A.3 and A.4, (A.9) has a unique regular optimizer for every $\xi \in \Xi$. Hence, by Theorem A.1, continuous solution (and multiplier) maps exist around every $\xi \in \Xi$. This implies that the maps $x^\star : \Xi \to \mathbb{R}^n$ and $\mu^\star : \Xi \to \mathbb{R}^m$ exist globally on all of Ξ and are continuous.

Finally, because $A = \nabla^2_{xx} f + \sum_i \mu_i \nabla_{xx} g$ and $L^\star = \nabla_{\xi x} f + \sum_i \mu_i \nabla^2_{\xi x} g$ are weighted sums over μ, we generally require an upper bound on the dual multiplier $\mu^\star$.

Assumption A.5. *There exist $\zeta_i > 0$ such that, for all $\xi \in \Xi$ and every KKT point $(x^\star(\xi), \mu^\star(\xi))$ of (A.9), we have $\mu^\star_i(\xi) \leq \zeta_i$ for all $i = 1, \dots, m$.*

Finding (tight) upper bounds on the dual multipliers is tricky and depends on the problem structure.

If $\|\nabla_x f(x,\xi)\| \le B_f$ and $\|\nabla_x g_i(x,\xi)\| \le B_{g_i}$ are uniformly bounded for all $\xi \in \Xi$, all $x \in \mathcal{X}(\xi)$ and all $i = 1,\dots,m$, then it is easy to see that $\mu_i^\star \le \zeta_i := B_f / B_{g_i}$.

Another possibility, documented in [33, p.647], applies specifically to convex optimization problems (i.e., under Assumption A.3). Namely, if for all $\xi \in \Xi$ a strict lower bound $\tilde{f}_\xi^\star$ (i.e., $\tilde{f}_\xi^\star < f(x,\xi)$ for all $x \in \mathcal{X}(\xi)$) and a strictly feasible point $\tilde{x}_\xi$ (i.e., $g(\tilde{x}_\xi,\xi) < 0$) are known, then we can choose

$$\zeta_i = \sup_{\xi\in\Xi} \frac{f(\tilde{x}_\xi,\xi) - \tilde{f}_\xi^\star}{-g_i(\tilde{x}_\xi)}\,.$$

Combining Assumptions A.3 and A.5 we can guarantee that

$$A = \nabla_{xx}^2 f + \sum\nolimits_{i=1}^m \mu_i^\star \nabla_{xx}^2 g \preceq \left(\beta + \sum\nolimits_i \zeta_i \ell_i\right) \mathbb{I}\,.$$

In other words, we have $\lambda^{\max}(A(\xi)) \le \beta + \sum_i \zeta_i \ell_i$.

The assumptions made so far can be summarized in the following statement:

Corollary A.2. *Under Assumptions A.2, A.3, A.4, and A.5 the primal and dual solution maps $x^\star : \Xi \to \mathbb{R}^n$ and $\mu^\star : \Xi \to \mathbb{R}^m$ of (A.9) are single-valued and Lipschitz continuous (on Ξ) with respective Lipschitz constants*

$$\begin{aligned} \ell_{x^\star} &:= \sqrt{\frac{\beta + \sum_i \zeta_i \ell_i}{\alpha}} \left(\frac{\bar{L}^\star}{\alpha} + \frac{\bar{G}^\star}{\omega}\right), \\ \ell_{\mu^\star} &:= \frac{(\beta + \sum_i \zeta_i \ell_i)^{3/2}}{\alpha^{1/2}\omega} \left(\frac{\bar{L}^\star}{\alpha} + \frac{\bar{G}^\star}{\omega}\right), \end{aligned} \tag{A.10}$$

where $\bar{L}^\star := \sup_{\xi\in\Xi} \|L^\star(\xi)\|$ and $\bar{G}^\star := \sup_{\xi\in\Xi} \|G^\star(\xi)\|$ are upper bounds on $L^\star$ and $G^\star$ over Ξ, respectively.

It remains to establish bounds on $L^\star$ and $G^\star$. Such bounds, however, are highly problem dependent, and therefore we consider only special cases in the next section.

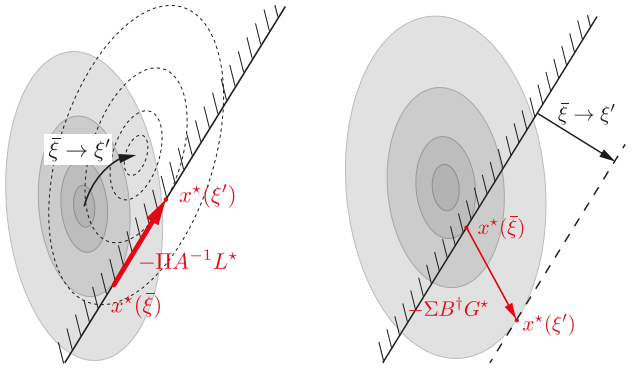

Figure A.1: Evolution of the optimizer for (a) stationary constraints and perturbed cost function; (b) fixed objective and linear constraint with righthand side perturbation.

A.3 Discussion of Sensitivity Bounds

We now discuss special cases for the established sensitivity bounds and how to relax the differentiability of f and g.

A.3.1 Special Cases

Stationary Constraints

We first assume that the constraints are stationary and only the cost function is perturbed, i.e., we consider the problem

$$\text{minimize} \;\; f(x, \xi) \quad \text{subject to} \;\; g(x) \leq 0$$

and assume that Assumptions A.2, A.3, A.4, and A.5 are satisfied. Let $(\bar{x}^\star, \bar{\mu}^\star)$ be a regular optimizer for $\bar{\xi}$ for which SCS holds, and therefore Proposition A.1 is applicable.

In this case, $G^\star(\xi)$ is vacuous and we have $L^\star = \nabla^2_{x\xi} L = \nabla^2_{x\xi} f$. Hence, the evolution of $x^\star$ and $\mu^\star$ around $\bar{\xi}$ is governed by the first terms in (A.3).

In particular, we have $\nabla_\xi x^\star = -\Pi A^{-1} L^\star$, which is independent of $B^\dagger$ (and hence of $\sigma_{B^T}^{\min}$). In other words, the conditioning of the constraints (as quantified by Assumption A.4) does not affect the sensitivity of $x^\star$ (however, it does affect the sensitivity of $\mu^\star$).

In this context, we can also make sense of the operator Π which projects the quantity $A^{-1}L^\star$ onto $\ker B^T$ which is the tangent space of the surface of (strongly) active constraints. In other words, Π guarantees that $x^\star$ can change only along the surface spanned by the (strongly) active constraints, which is plausible since the constraints are stationary. Figure A.1a illustrates this effect.

If we further assume that no constraints are active at $\bar{x}^\star$, we find ourselves in a situation of unconstrained optimization. In this case, (A.3) reduces to $\nabla_\xi x^\star = -\left(\nabla_{xx}^2 f(\xi)\right)^{-1} \nabla_{\xi x}^2 f$. This results can also be derived by a direct application of the implicit function theorem, for which $\ell_{x^\star}$ reduces to

$$\ell_{x^\star} := \frac{\ell_f}{\alpha}, \quad \text{where} \quad \ell_f := \sup_{\xi' \in \Xi} \|\nabla_{\xi x}^2 f(x^\star(\xi'), \xi')\| \,. \tag{A.11}$$

In particular, for unconstrained time-varying optimization, the assumption that $\nabla_x f$ is β-Lipschitz is not required.

Translational Perturbation of Objective

A particular class of objective functions is when the perturbation is in the form of a translation, i.e., we consider

$$\text{minimize } \hat{f}(x - c(\xi)) \quad \text{subject to } \; g(x) \leq 0 \,, \tag{A.12}$$

where $\hat{f} : \mathbb{R}^n \to \mathbb{R}$ is α-strongly convex and β-smooth and $c : \mathbb{R}^r \to \mathbb{R}^n$ is continuously differentiable and ℓ_c-Lipschitz. Otherwise, let Assumptions A.2, A.3, A.4, and A.5 hold. Considering $f(x, \xi) = \hat{f}(x - c(\xi))$, the structure of (A.12) implies that $\nabla_x f(x, \xi) = \nabla \hat{f}(x - c(\xi))$, and therefore $\nabla_{\xi x}^2 f(x, \xi) = -\nabla_{xx}^2 \hat{f}(x - c(\xi)) \nabla c(\xi)$, and therefore $\|L^\star\| \leq \beta \ell_c$ (since $\nabla_{\xi x} g = 0$). In fact, c does not need to be continuously differentiable, but only Lipschitz as discussed in the forthcoming Appendix A.3.2.

Right-hand Constraint Perturbations

Next, we consider a stationary cost function and a constraint parametrization that takes the form of a right-hand side perturbation, i.e.,

$$\text{minimize } f(x) \quad \text{subject to } u(x) \leq v(\xi)\,, \tag{A.13}$$

where $u : \mathbb{R}^n \to \mathbb{R}^m$ and $v : \Xi \to \mathbb{R}^m$ are twice continuously differentiable, and Assumptions A.2, A.3, A.4, and A.5 hold.

Let $(\bar{x}^\star, \bar{\mu}^\star)$ be a regular optimizer for $\bar{\xi}$ and let SCS hold, such that $x^\star$ and $\mu^\star$ are differentiable at $\bar{\xi}$.

Exploiting the structure of (A.13), we have $\nabla^2_{\xi x} g \equiv 0$, and $\nabla^2_{\xi x} f \equiv 0$, and thus $L^\star \equiv 0$. Consequently, according to Proposition A.1 it holds that $\nabla_\xi x^\star = -\Sigma B^\dagger G^\star$.

Recall from Lemma A.2 and its proof, that Σ denotes an orthogonal projection onto the space spanned by $Q^{-1}\nabla g_i(\xi)$ for all $i \in \mathbf{I}^\star(\xi)$ where $Q := A^{1/2}$. Roughly speaking, as the constraints vary according to (A.13), the optimizer can move only Q^{-1}-orthogonally to the surface spanned by the active constraints. This behavior is illustrated in Figure A.1b.

Linear Constraints

Yet, a more special case is to assume that $\mathcal{X}(\xi)$ is a polyhedron, that is, to consider

$$\text{minimize } f(x) \quad \text{subject to } Ux \leq v(\xi)\,, \tag{A.14}$$

where $U \in \mathbb{R}^{m \times n}$ is assumed to have full row rank (thus satisfying Assumption A.4). The fact that in this case we have that $\nabla^2_{xx} g = 0$ (and that $L^\star = \nabla^2_{\xi x} f$) obviates the need for Assumption A.5, i.e., there is no need to estimate the upper bound of the dual multipliers. Hence, the bounds $\ell_{x^\star}$ and $\ell_{\mu^\star}$, in Corollary A.2 simplify to

$$\ell_{x^\star} := \sqrt{\frac{\beta}{\alpha}}\frac{\ell_v}{\omega} \quad \text{and} \quad \ell_{\mu^\star} := \frac{\beta^{3/2}}{\alpha^{1/2}}\frac{\ell_v}{\omega^2}\,, \tag{A.15}$$

where $\ell_v := \sup_{\xi \in \Xi} \|\nabla v(\xi)\|$ is the Lipschitz constant of v.

A.3.2 Relaxation of Differentiability Assumption

All results thus far require f and g, from (A.9), to be twice continuously differentiable in (x, ξ). Whether it is possible to directly relax this (possibly restrictive) assumption without losing Lipschitz continuity is still an open question.

However, this issue sometimes can be circumvented by a perturbation mapping that isolates the non-smoothness. Namely, assume that there exist a map $\tilde{\xi} : \mathbb{R}^p \to \mathbb{R}^r$ that is $\ell_{\tilde{\xi}}$-Lipschitz. Then, based on (A.9), we consider the problem

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f(x, \tilde{\xi}(\psi)), \\ \text{subject to} \quad & g(x, \tilde{\xi}(\psi)) \leq 0 \end{aligned} \tag{A.16}$$

parametrized in ψ. If (A.9) admits a $\ell_{x^\star}$-Lipschitz solution map with respect to ξ, then (A.16) admits a $\ell_{x^\star}\ell_{\tilde{\xi}}$-Lipschitz solution map $\tilde{x}^\star$, in ψ. This holds since the composition of Lipschitz maps is also Lipschitz.

Based on the special problem classes discussed in the previous subsection, we can also derive the following result which incorporates non-differentiable perturbations with tighter bound on the Lipschitz constant than derived from (A.16):

Corollary A.3. *Consider the problem*

$$\begin{aligned} \text{minimize} \quad & \hat{f}(x - c(\xi)) \\ \text{subject to} \quad & Ux \leq v(\xi) \end{aligned} \tag{A.17}$$

where $\xi \in \Xi \subset \mathbb{R}^r$. *Define* $\hat{\mathcal{X}}(\xi) := \{x \,|\, Ux \leq v(\xi)\}$ *and let*

(i) $\hat{f} : \mathbb{R}^n \to \mathbb{R}$ *be twice continuously differentiable,* α*-strongly convex, with* β*-Lipschitz* $\nabla \hat{f}$,

(ii) $c : \Xi \to \mathbb{R}^n$ *and* $v : \Xi \to \mathbb{R}^m$ *be* ℓ_c*- and* ℓ_v*-Lipschitz, respectively, and*

(iii) $U \in \mathbb{R}^{n \times m}$ *and* $\omega > 0$ *such that for every* $\xi \in \Xi$ *and ever* $x \in \hat{\mathcal{X}}(\xi)$ *one has* $\omega^2 \mathbb{I} \preceq U_{\mathbf{I}_x^\xi} U_{\mathbf{I}_x^\xi}^T$.

If $\hat{\mathcal{X}}(t) \neq \emptyset$ for every $\xi \in \Xi$, then the primal and dual solution maps $x^\star : \Xi \to \mathbb{R}^n$ and $\mu^\star : \Xi \to \mathbb{R}^m$ of (A.17) are Lipschitz with respective bounds on Lipschitz constants

$$\ell_{x^\star} = \sqrt{\frac{\beta}{\alpha}} \left(\frac{\beta \ell_c}{\alpha} + \frac{\ell_v}{\omega} \right) \text{ and } \ell_{\mu^\star} = \frac{\beta^{3/2}}{\alpha^{1/2}\omega} \left(\frac{\beta \ell_c}{\alpha} + \frac{\ell_v}{\omega} \right) .$$

Proof. The key point about Corollary A.3 is the fact that c and v are Lipschitz continuous, but not twice differentiable. To address this issue, consider, instead of (A.17), the problem

$$\text{minimize} \quad \hat{f}(x - c) \qquad \text{subject to} \quad Ux \leq v \tag{A.18}$$

parametrized in (c, v). Under the stated assumptions, the parametrized objective $(x, c) \mapsto \hat{f}(x - c)$ and constraint function $(x, v) \mapsto Ux - v$ are twice continuously differentiable in (x, c) and (x, v), respectively. Therefore, the solution map $(c, v) \mapsto \hat{x}^\star(c, v)$ of (A.18) is Lipschitz by Corollary A.2. In fact, by considering the reasoning for translational perturbations of the objective and righthand perturbations of polyhedra in the previous subsection, we know that $\hat{x}^\star$ is Lipschitz in c and v with Lipschitz constants $(\frac{\beta}{\alpha})^{3/2}$ and $\sqrt{\frac{\beta}{\alpha}}\frac{1}{\omega}$, respectively.

Moreover, replacing c with $c(\xi)$ and using ℓ_c-Lipschitz continuity of c we have, for all $c', c \in \mathbb{R}^n$ and all $v \in \mathbb{R}^m$,

$$\|\hat{x}^\star(c(\xi'), v) - \hat{x}^\star(c(\xi), v)\| \leq \sqrt{\tfrac{\beta}{\alpha}}\tfrac{\beta}{\alpha}\ell_c\|\xi' - \xi\| \tag{A.19}$$

Similarly, replacing v with $v(\xi)$ we have

$$\|\hat{x}^\star(c, v(\xi')) - \hat{x}^\star(c, v(\xi))\| \leq \sqrt{\tfrac{\beta}{\alpha}}\tfrac{1}{\omega}\ell_v\|\xi' - \xi\| \tag{A.20}$$

for all $c \in \mathbb{R}^n$ and all $v', v \in \mathbb{R}^m$.

The mutual uniformity of the these Lipschitz constants, i.e., the fact that (A.19) and (A.20) are independent of v and c, respectively, allows us to write

$$\|x^\star(\xi') - x^\star(\xi)\| \leq \sqrt{\tfrac{\beta}{\alpha}} \left(\tfrac{\beta \ell_c}{\alpha} + \tfrac{\ell_v}{\omega} \right) \|\xi' - \xi\| ,$$

where $x^\star(\xi) = \hat{x}^\star(c(\xi), v(\xi))$ and thus the solution map $x^\star : \Xi \to \mathbb{R}^n$ of (A.17) is Lipschitz with the desired Lipschitz constant. The Lipschitz constant for $\mu^\star$ follows analogously. □

APPENDIX B

Complete Simulation Results

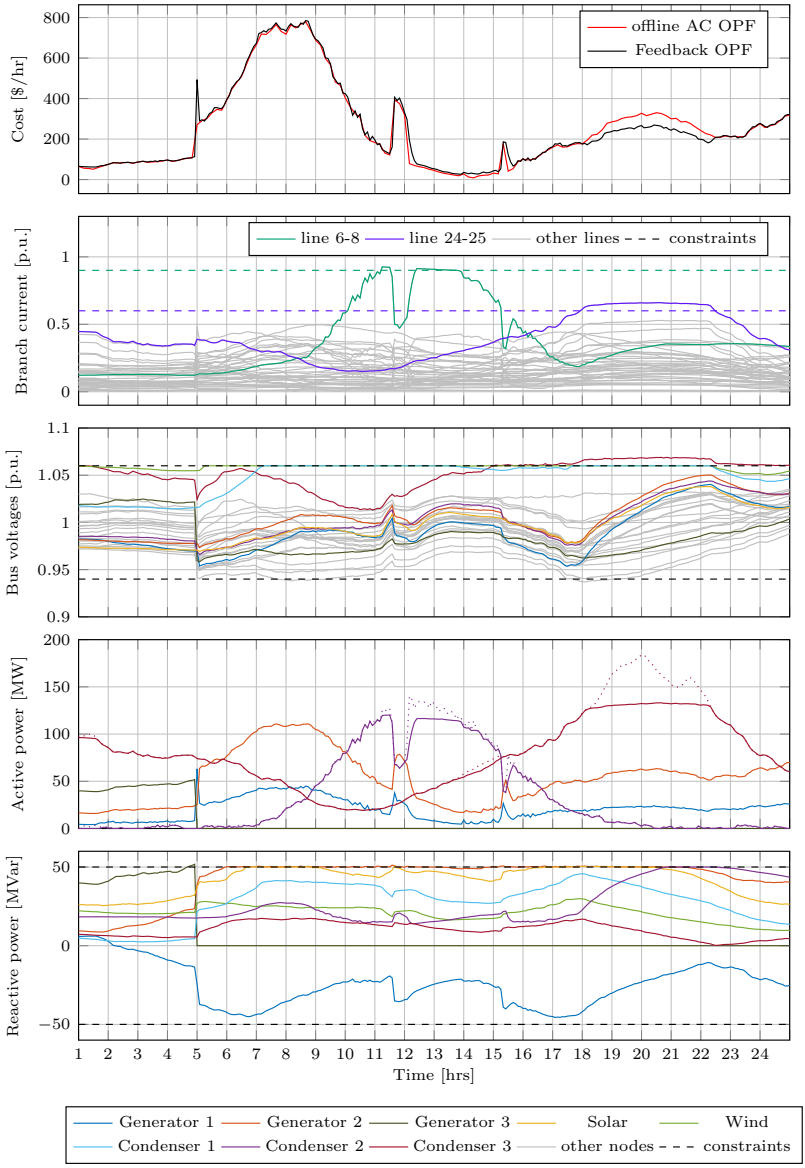

Figure B.1: Penalty scheme under nominal conditions

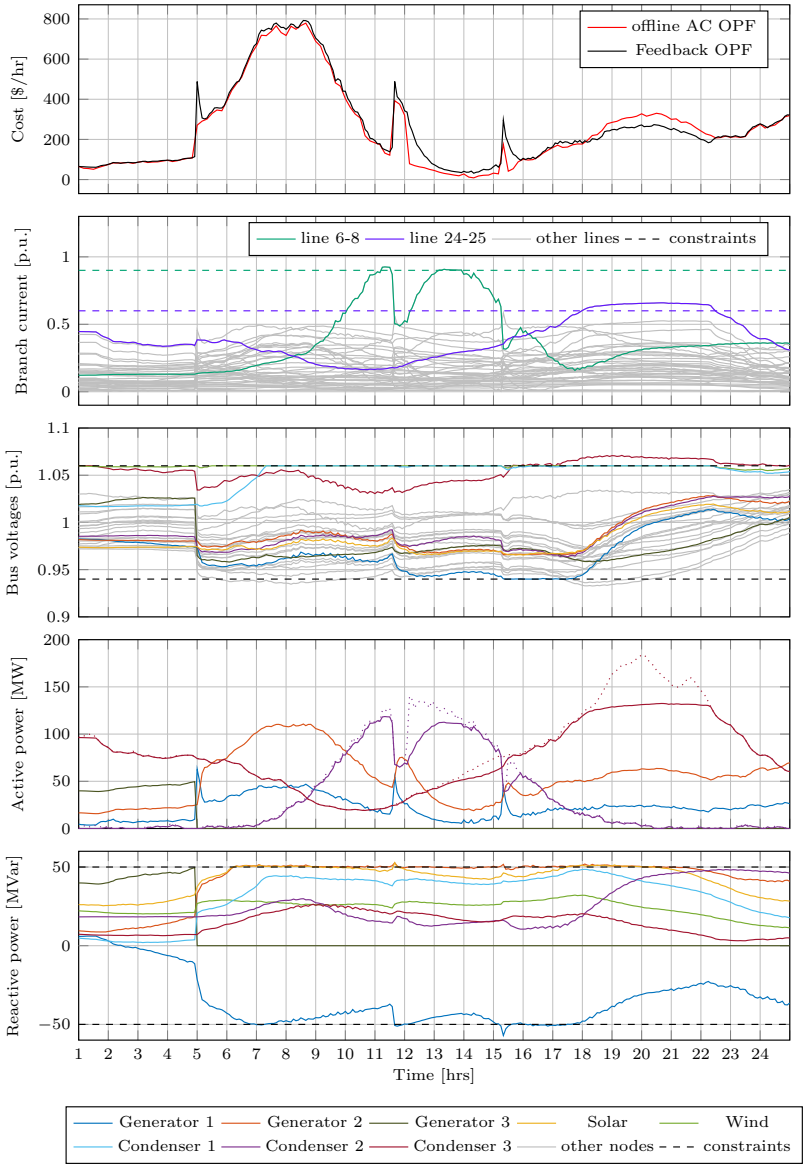

Figure B.2: Penalty scheme with delay of 2min

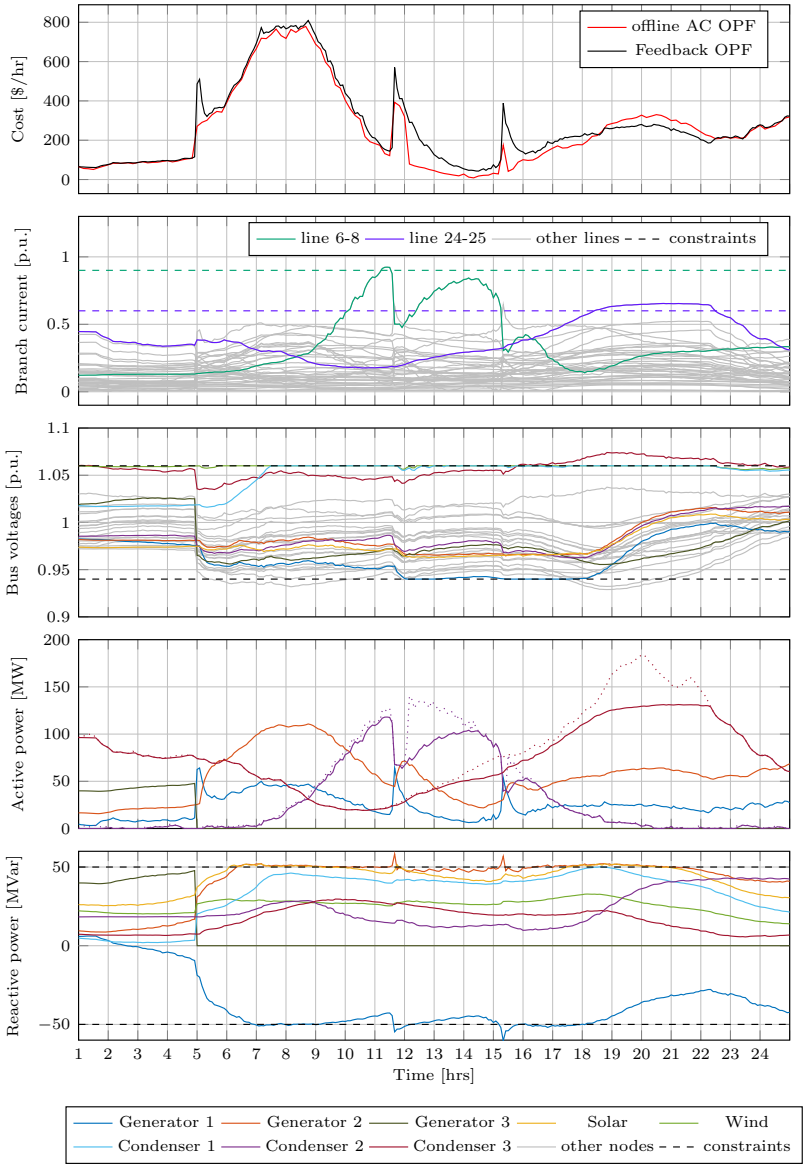

Figure B.3: Penalty scheme with delay of 5min

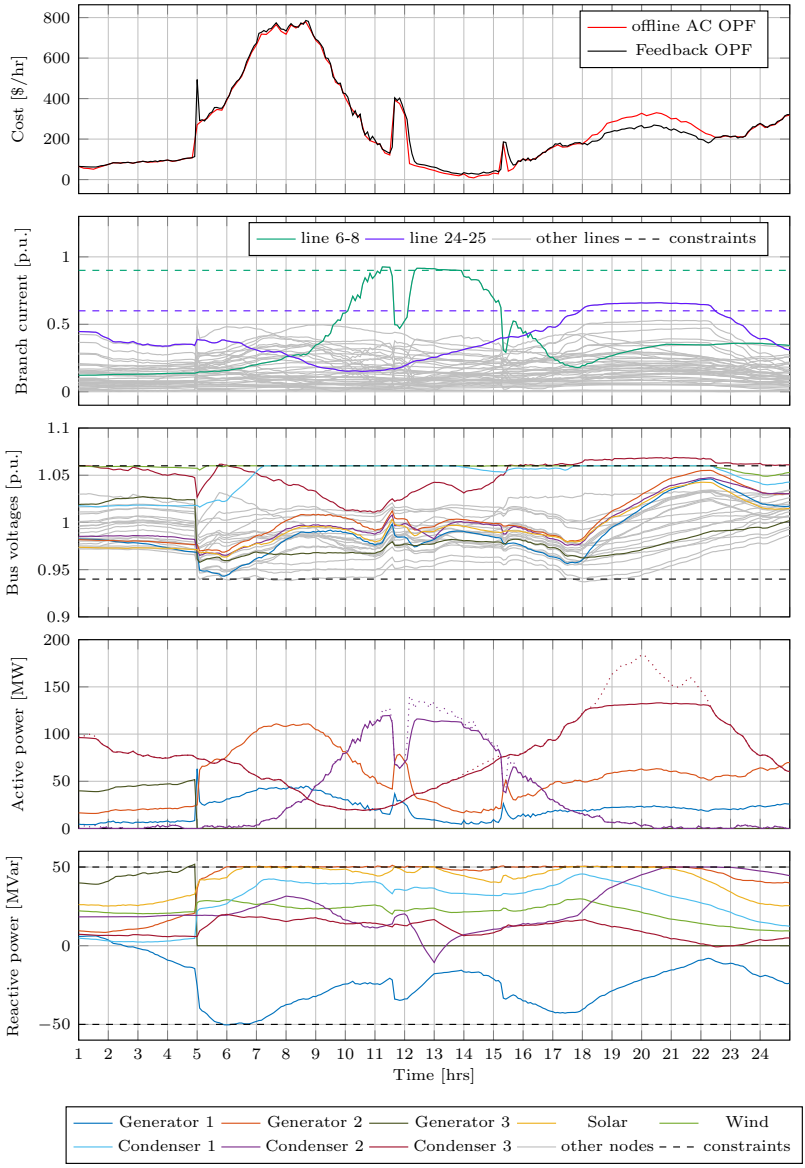

Figure B.4: Penalty scheme with linearization update every 60min

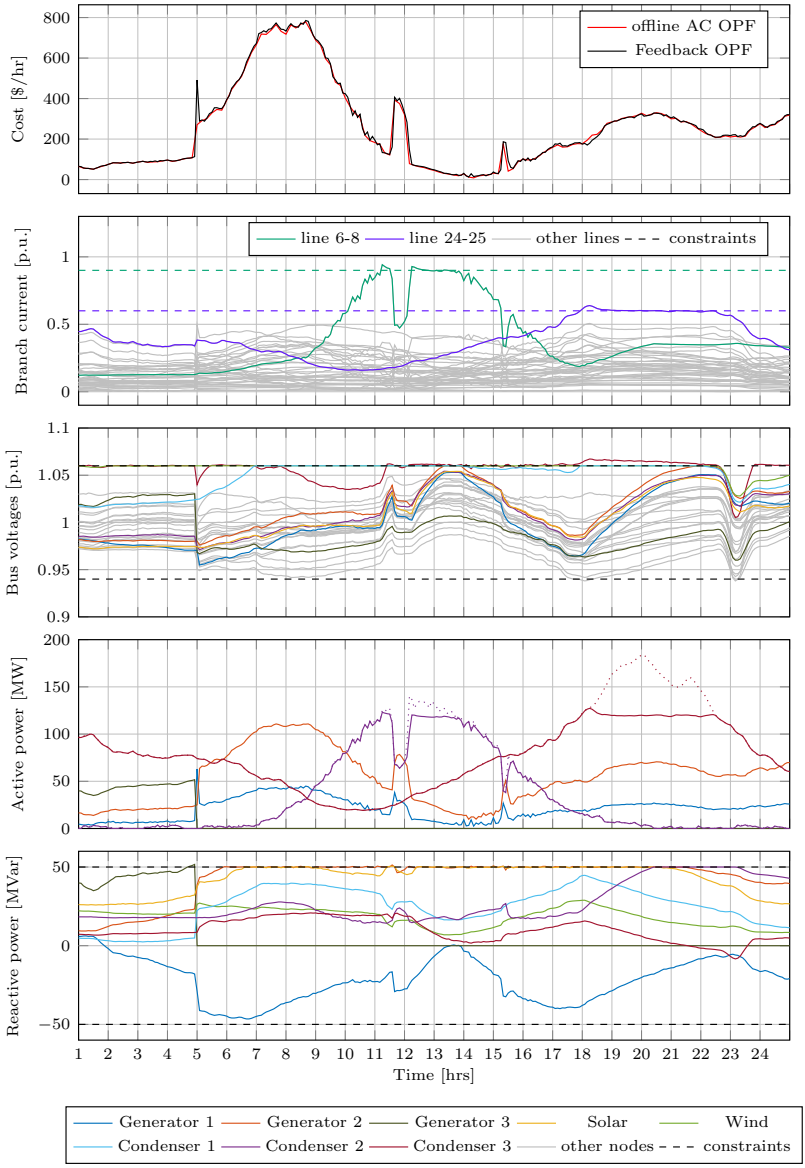

Figure B.5: Primal-dual scheme under nominal conditions

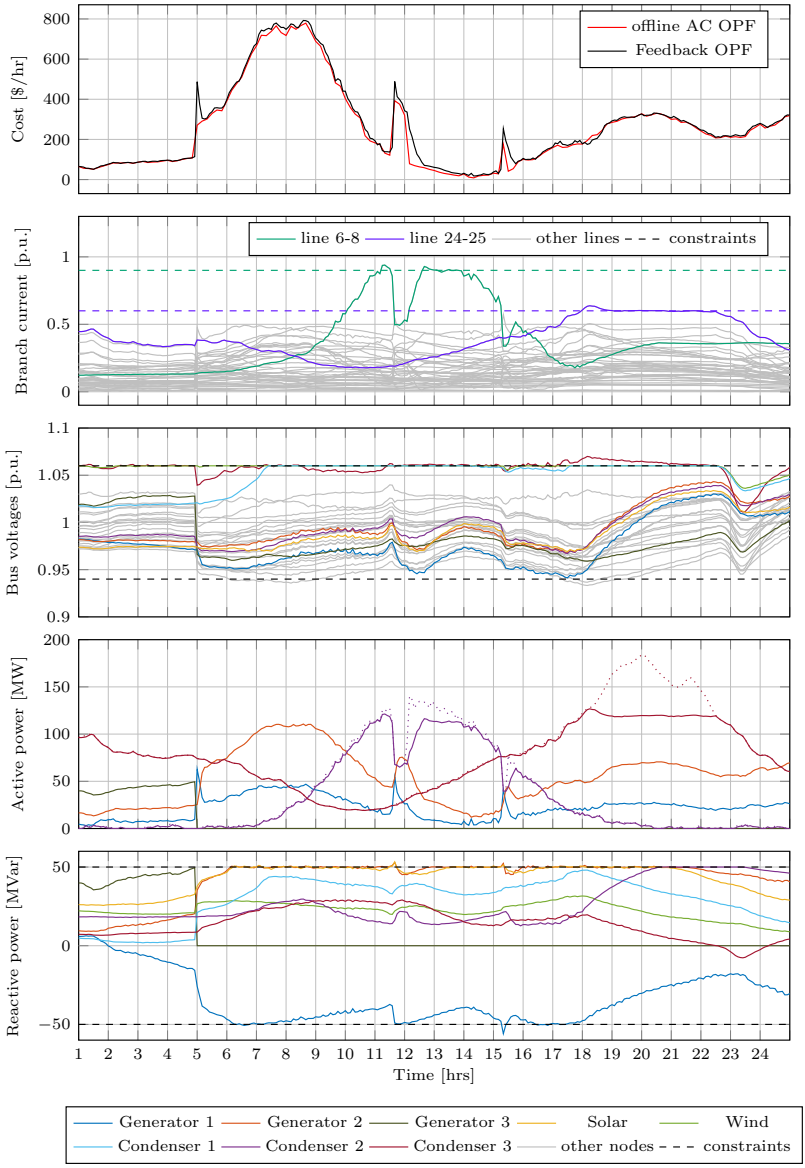

Figure B.6: Primal-dual scheme with delay of 2min

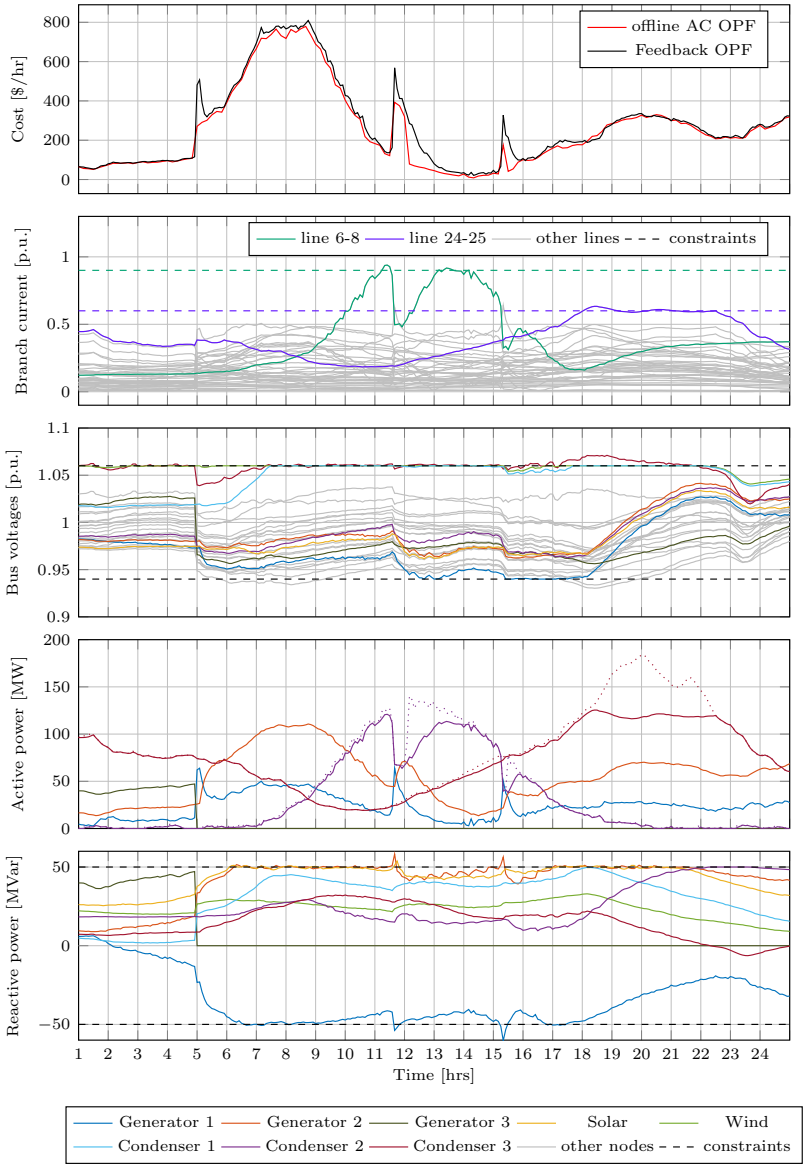

Figure B.7: Primal-dual scheme with delay of 5min

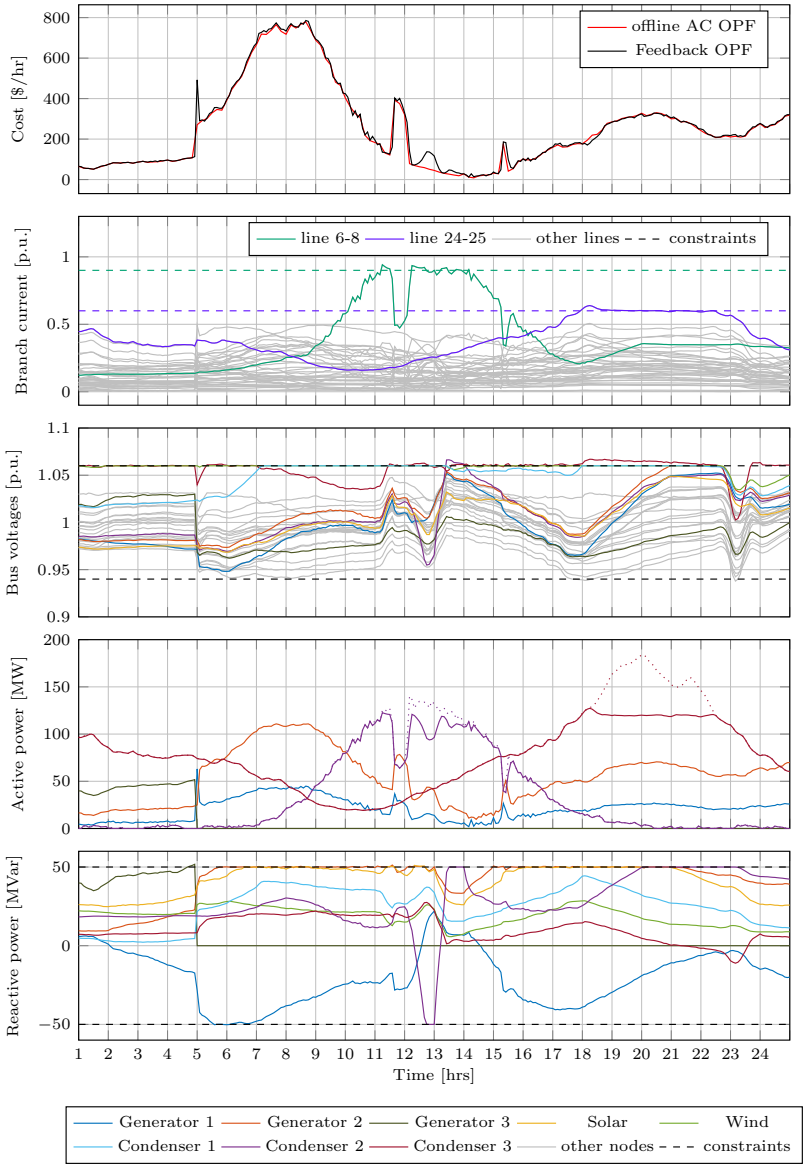

Figure B.8: Primal-dual scheme with linearization update every 60min

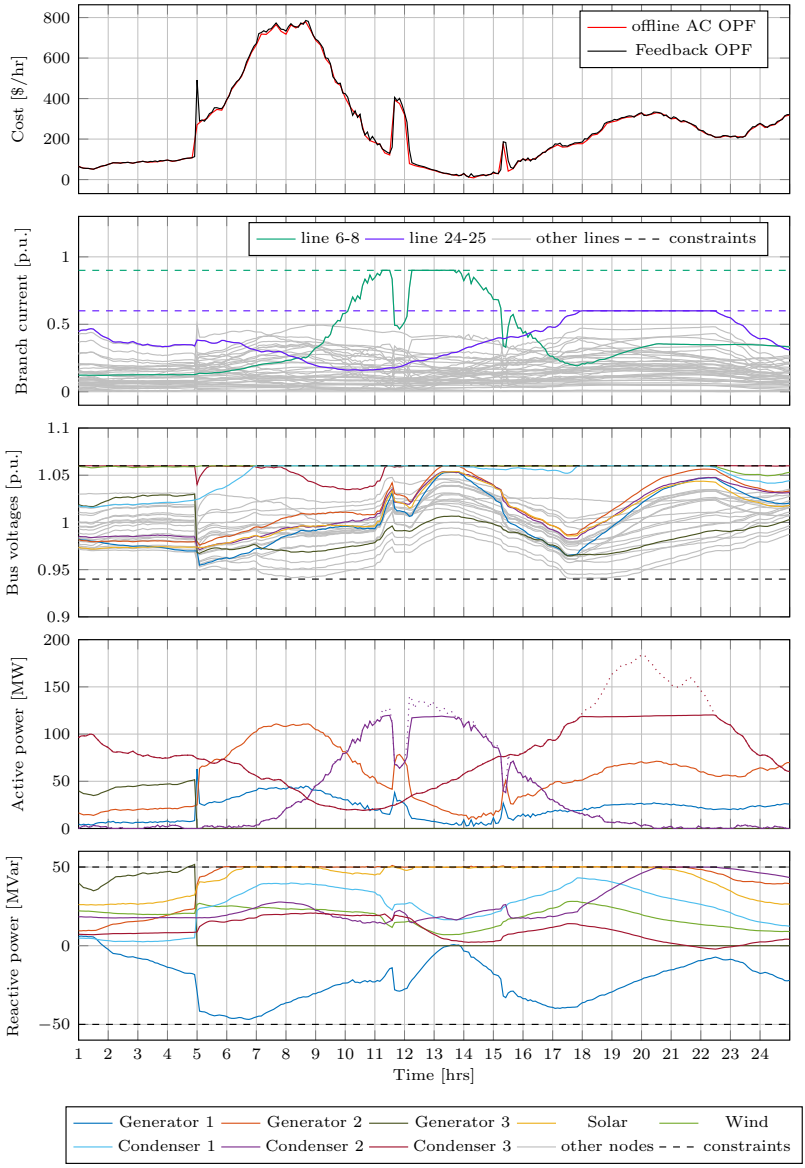

Figure B.9: LOP scheme under nominal conditions

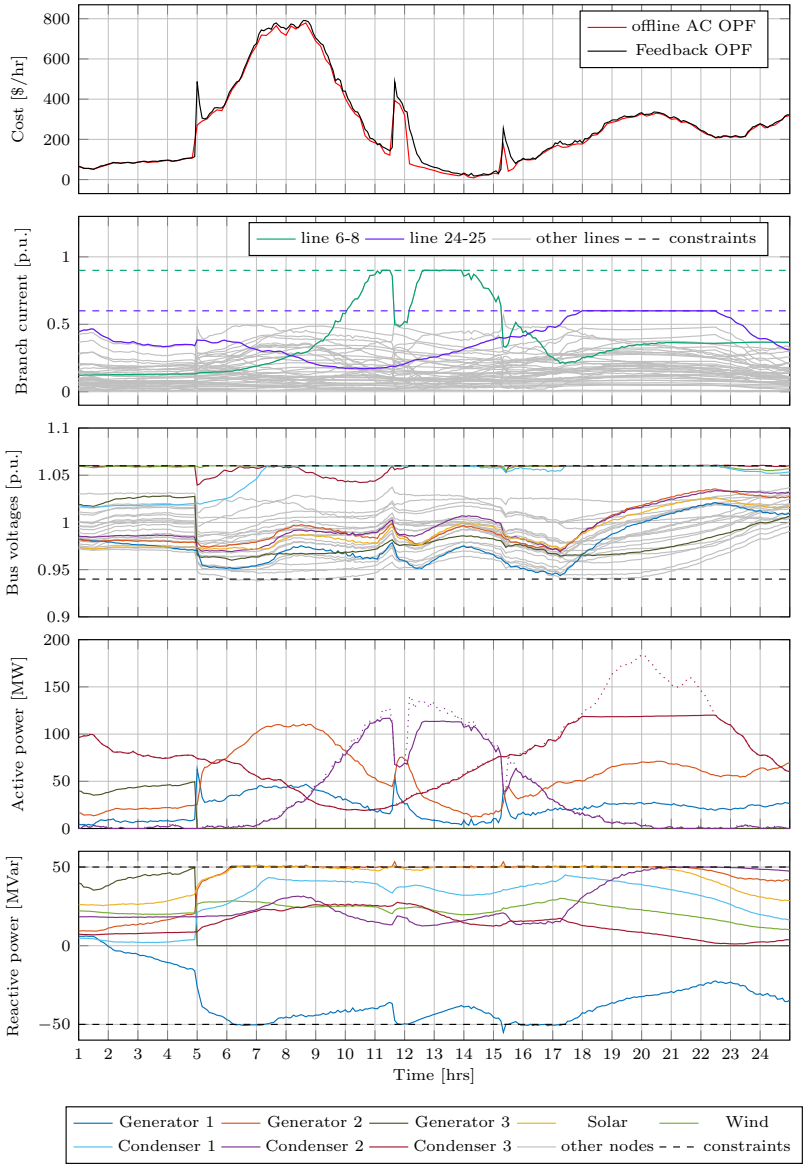

Figure B.10: LOP scheme with delay of 2min

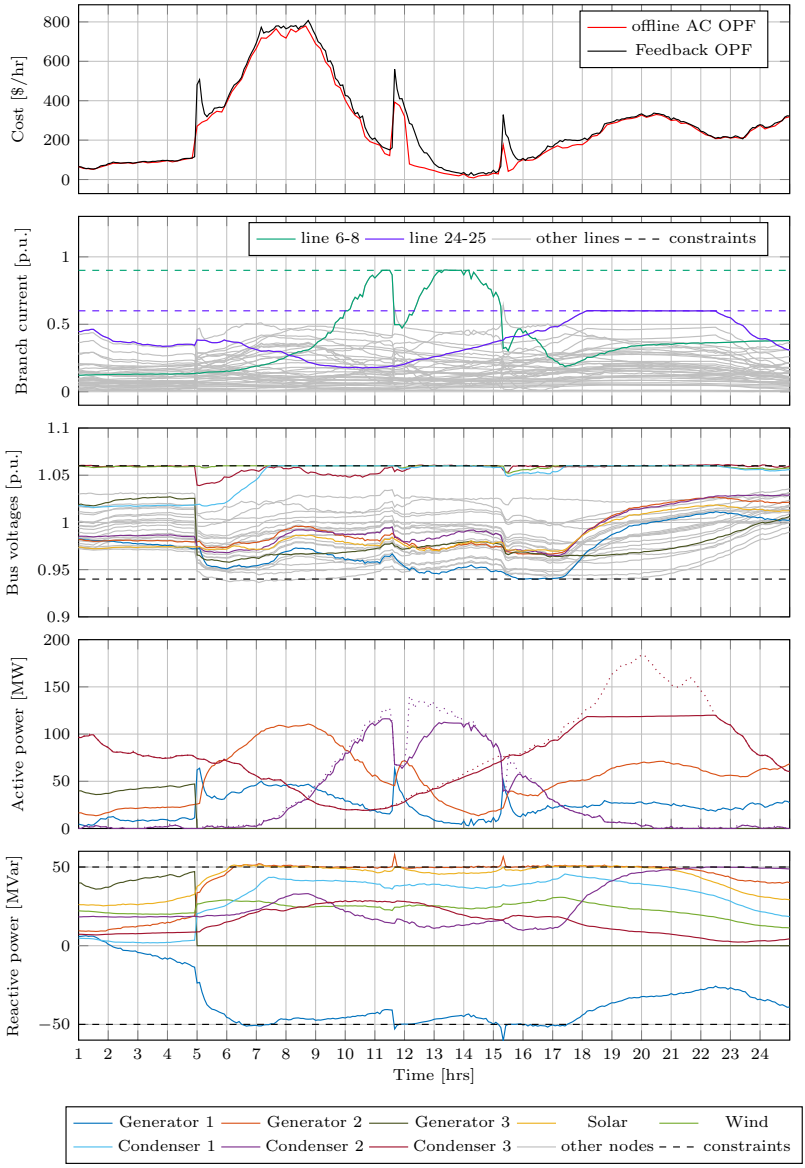

Figure B.11: LOP scheme with delay of 5min

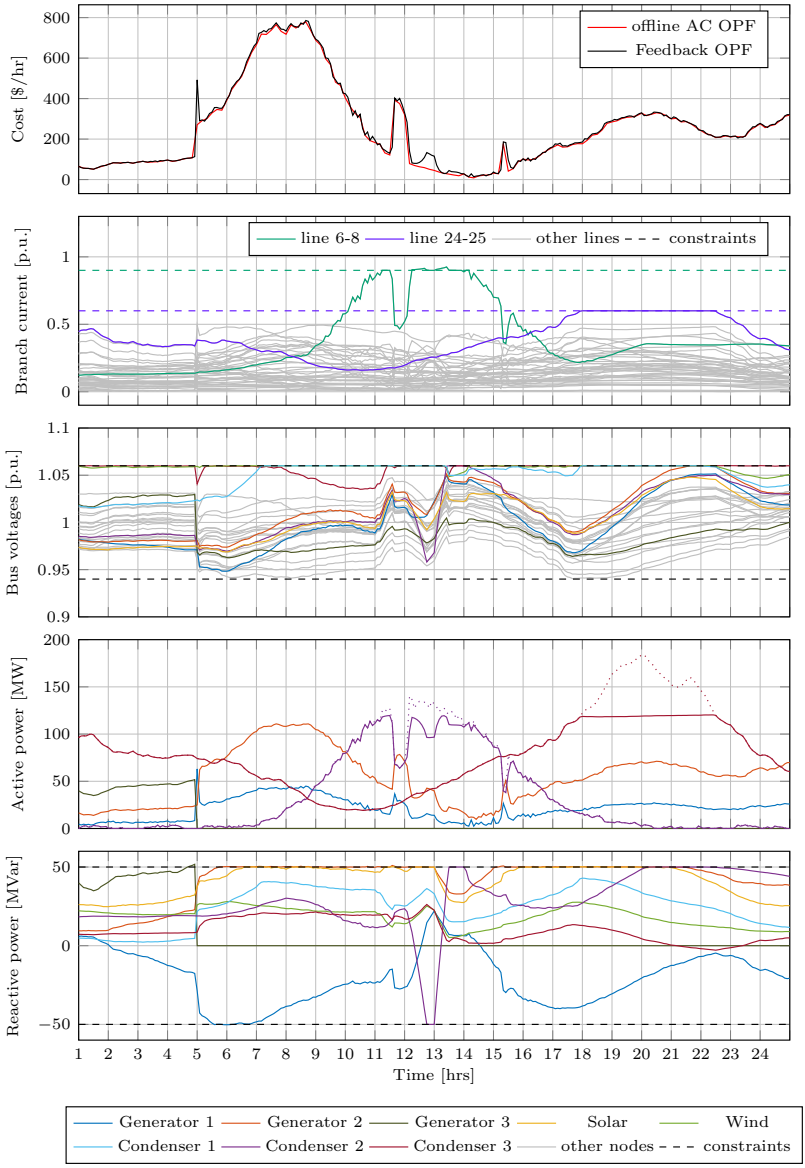

Figure B.12: LOP scheme with linearization update every 60min

Publications

The work presented in this thesis is a result of close collaboration with colleagues and is mainly based on the submitted or published articles below.

[Ha1] V. Häberle, **A. Hauswirth**, L. Ortmann, S. Bolognani, and F. Dörfler. "Non-Convex Feedback Optimization with Input and Output Constraints". In: *IEEE Control Systems Letters* 5.1 (2020), pp. 343–348.

[Ha2] **A. Hauswirth**, S. Bolognani, and F. Dörfler. "Projected Dynamical Systems on Irregular, Non-Euclidean Domains for Nonlinear Optimization". In: *SIAM Journal on Control and Optimization* (2020). in press [Online: arXiv 1905.06291].

[Ha3] **A. Hauswirth**, S. Bolognani, G. Hug, and F. Dörfler. "Timescale Separation in Autonomous Optimization". In: *IEEE Transactions on Automatic Control* (2020). in press [Online: IEEExplore Early Access].

[Ha4] **A. Hauswirth**, F. Dörfler, and A. R. Teel. "Anti-Windup Approximations of Oblique Projected Dynamical Systems for Feedback-Based Optimization". In: *SIAM Journal on Control and Optimization* (2020). under review [Online: arXiv 2003.00478].

[Ha5] **A. Hauswirth**, F. Dörfler, and A. R. Teel. "On the Differentiability of Projected Trajectories and the Robust Convergence of Non-Convex Anti-Windup Gradient Flows". In: *IEEE Control Systems Letters* 4.3 (2020), pp. 620–625.

[Ha6] **A. Hauswirth**, F. Dörfler, and A. R. Teel. "On the Implementation of Projected Dynamical Systems with Anti-Windup Controllers". In: *American Control Conference (ACC)*. Denver, CO, July 2020, pp. 1286–1291.

[Ha7] **A. Hauswirth**, L. Ortmann, S. Bolognani, and F. Dörfler. "Limit Behavior and the Role of Augmentation in Projected Saddle Flows for Convex Optimization". In: *21st IFAC World Congress*. Berlin, Germany, July 2020.

[Ha8] **A. Hauswirth**, S. Bolognani, G. Hug, and F. Dörfler. "Optimization Algorithms as Robust Feedback Controllers". In: *Annual Reviews in Control* (2020). in preparation.

[Ha9] L. Ortmann, **A. Hauswirth**, I. Caduff, F. Dörfler, and S. Bolognani. "Experimental Validation of Feedback Optimization in Power Distribution Grids". In: *Power Systems Computation Conference (PSCC)*. Porto, Portugal, June 2020.

[Ha10] I. Subotić, **A. Hauswirth**, and F. Dörfler. "Quantitative Sensitivity Bounds for Nonlinear Programming and Time-Varying Optimization". In: *IEEE Transactions on Automatic Control* (June 2020). under review [Online: arXiv 2006.10693].

[Ha11] **A. Hauswirth**, S. Bolognani, G. Hug, and F. Dörfler. "Generic Existence of Unique Lagrange Multipliers in AC Optimal Power Flow". In: *IEEE Control Systems Letters* 2.4 (Oct. 2018), pp. 791–796.

[Ha12] **A. Hauswirth**, I. Subotić, S. Bolognani, G. Hug, and F. Dörfler. "Time-Varying Projected Dynamical Systems with Applications to Feedback Optimization of Power Systems". In: *IEEE Conference on Decision and Control (CDC)*. Miami Beach, FL, Dec. 2018, pp. 3258–3263.

[Ha13] S. Menta, **A. Hauswirth**, S. Bolognani, G. Hug, and F. Dörfler. "Stability of Dynamic Feedback Optimization with Applications to Power Systems". In: *56th Annual Allerton Conference on Communication, Control, and Computing*. Monticello, IL, Oct. 2018, pp. 136–143.

[Ha14] **A. Hauswirth**, A. Zanardi, S. Bolognani, F. Dörfler, and G. Hug. "Online Optimization in Closed Loop on the Power Flow Manifold". In: *IEEE PowerTech Conference*. Manchester, UK, June 2017.

[Ha15] **A. Hauswirth**, S. Bolognani, G. Hug, and F. Dörfler. "Projected Gradient Descent on Riemannian Manifolds with Applications to Online Power System Optimization". In: *54th Annual Allerton Conference on Communication, Control, and Computing*. Monticello, IL, Sept. 2016, pp. 225–232.

References

[1] P. Absil, R. Mahony, and B. Andrews. "Convergence of the Iterates of Descent Methods for Analytic Cost Functions". In: *SIAM Journal on Optimization* 16.2 (Jan. 2005), pp. 531–547.

[2] P.-A. Absil and K. Kurdyka. "On the Stable Equilibrium Points of Gradient Systems". In: *Systems & Control Letters* 55.7 (2006), pp. 573–577.

[3] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.

[4] P.-A. Absil and J. Malick. "Projection-like Retractions on Matrix Manifolds". In: *SIAM Journal on Optimization* 22.1 (Jan. 2012), pp. 135–158.

[5] S. Adly, F. Nacry, and L. Thibault. "Preservation of Prox-Regularity of Sets with Applications to Constrained Optimization". In: *SIAM Journal on Optimization* 26.1 (Jan. 2016), pp. 448–473.

[6] S. Adly and B. K. Le. "Non-Convex Sweeping Processes Involving Maximal Monotone Operators". In: *Optimization* 66.9 (Sept. 2017), pp. 1465–1486.

[7] A. Alessio and A. Bemporad. "A Survey on Explicit Model Predictive Control". In: *Nonlinear Model Predictive Control*. Vol. 384. Lecture Notes in Control and Information Sciences. Berlin Heidelberg, Germany: Springer, 2009, pp. 345–369.

[8] F. Alvarez, J. Bolte, and O. Brahic. "Hessian Riemannian Gradient Flows in Convex Programming". In: *SIAM Journal on Control and Optimization* 43.2 (Jan. 2004), pp. 477–501.

[9] H. Amann. *Ordinary Differential Equations: An Introduction to Nonlinear Analysis*. De Gruyter Studies in Mathematics 13. Berlin, Germany: De Gruyter, 1990.

[10] M. Andreasson, D. V. Dimarogonas, H. Sandberg, and K. H. Johansson. "Distributed PI-Control with Applications to Power Systems Frequency Control". In: *American Control Conference (ACC)*. Portland, OR, June 2014, pp. 3183–3188.

[11] K. B. Ariyur and M. Krstic. *Real Time Optimization by Extremum Seeking Control*. First. Hoboken, NJ: Wiley Interscience, 2003.

[12] D. B. Arnold, M. Negrete-Pincetic, M. D. Sankur, D. M. Auslander, and D. S. Callaway. "Model-Free Optimal Control of VAR Resources in Distribution Systems: An Extremum Seeking Approach". In: *IEEE Transactions on Power Systems* 31.5 (Sept. 2016), pp. 3583–3593.

[13] K. J. Arrow, L. Hurwicz, and H. Uzawa. *Studies in Linear and Nonlinear Programming*. Stanford, CA: Stanford University Press, 1958.

[14] K. J. Åström and R. M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ: Princeton University Press, 2008.

[15] H. Attouch and Z. Chbani. "Fast Inertial Dynamics and FISTA Algorithms in Convex Optimization. Perturbation Aspects". In: *arXiv 1507.01367 [math]* (July 2015).

[16] H. Attouch, J. Peypouquet, and P. Redont. "Fast Convergence of an Inertial Gradient-like System with Vanishing Viscosity". In: *arXiv 1507.04782 [math]* (July 2015).

[17] J.-P. Aubin and A. Cellina. *Differential Inclusions: Set-Valued Maps and Viability Theory*. Grundlehren Der Mathematischen Wissenschaften 264. Berlin Heidelberg, Germany: Springer, 1984.

[18] J. P. Aubin. *Viability Theory*. Systems & Control: Foundations & Applications. New York, NY: Springer, 1991.

[19] M. E. Baran and F. F. Wu. "Optimal Capacitor Placement on Radial Distribution Systems". In: *IEEE Transactions on Power Delivery* 4.1 (1989), pp. 725–734.

[20] M. E. Baran and F. F. Wu. "Optimal Sizing of Capacitors Placed on a Radial Distribution System". In: *IEEE Transactions on Power Delivery* 4.1 (1989), pp. 735–743.

[21] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. New York, NY: Springer, 2011.

[22] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Hoboken, NJ: John Wiley & Sons, Inc., 2006.

[23] A. Beck. *First-Order Methods in Optimization*. Philadelphia, PA: SIAM, Oct. 2017.

[24] A. Beck and M. Teboulle. "Mirror Descent and Nonlinear Projected Subgradient Methods for Convex Optimization". In: *Operations Research Letters* 31.3 (May 2003), pp. 167–175.

[25] A. Beck and M. Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems". In: *SIAM Journal on Imaging Sciences* 2.1 (Jan. 2009), pp. 183–202.

[26] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos. "The Explicit Linear Quadratic Regulator for Constrained Systems". In: *Automatica* 38 (2002).

[27] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton, NJ: Princeton University Press, Oct. 2009.

[28] F. Bernicot and J. Venel. "Sweeping Process by Prox-Regular Sets in Riemannian Hilbert Manifolds". In: *Journal of Differential Equations* 259.8 (Oct. 2015), pp. 4086–4121.

[29] A. Bernstein, E. Dall'Anese, and A. Simonetto. "Online Primal-Dual Methods With Measurement Feedback for Time-Varying Convex Optimization". In: *IEEE Transactions on Signal Processing* 67.8 (Apr. 2019), pp. 1978–1991.

[30] D. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas*. Second. Princeton, NJ: Princeton University Press, 2009.

[31] D. P. Bertsekas. *Nonlinear Programming*. Second. Belmont, MA: Athena Scientific, Sept. 1999.

[32] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Fourth. Vol. 1. Belmont, MA: Athena Scientific, 2017.

[33] D. P. Bertsekas and S. K. Mitter. "A Descent Numerical Method for Optimizaiton Problems with Non-Differentiable Cost Functionals". In: *SIAM Journal on Control* 11.4 (Nov. 1973), pp. 637–652.

[34] D. Bertsimas and R. Freund. *Data, Models, and Decisions: The Fundamentals of Management Science*. Charlestown, MA: Dynamic Ideas, June 2004.

[35] O. Besbes, Y. Gur, and A. Zeevi. "Non-Stationary Stochastic Optimization". In: *Operations Research* 63.5 (Oct. 2015), pp. 1227–1244.

[36] P. Binetti, K. B. Ariyur, M. Krstic, and F. Bernelli. "Formation Flight Optimization Using Extremum Seeking Feedback". In: *Journal of Guidance, Control, and Dynamics* 26.1 (Jan. 2003), pp. 132–142.

[37] C. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. New York, NY: Springer, 2009.

[38] A. M. Bloch, R. W. Brockett, and T. S. Ratiu. "Completely Integrable Gradient Flows". In: *Communications in Mathematical Physics* 147.1 (June 1992), pp. 57–74.

[39] A. M. Bloch, R. W. Brockett, and T. S. Ratiu. "On the Geometry of Saddle Point Algorithms". In: *IEEE Conference on Decision and Control (CDC)*. Tucson, AZ, Dec. 1992, pp. 1482–1487.

[40] H. G. Bock, M. Diehl, D. Leineweber, and J. P. Schlöder. "A Direct Multiple Shooting Method for Real-Time Optimization of Nonlinear DAE Processes". In: *Nonlinear Model Predictive Control*. Vol. 26. Progress in Systems and Control Theory. Basel, Switzerland: Springer, 2000, pp. 245–268.

[41] S. Bolognani and F. Dörfler. "Fast Power System Analysis via Implicit Linearization of the Power Flow Manifold". In: *53rd Annual Allerton Conference on Communication, Control, and Computing*. Monticello, IL, Sept. 2015, pp. 402–409.

[42] S. Bolognani, G. Cavraro, R. Carli, and S. Zampieri. "A Distributed Feedback Control Strategy for Optimal Reactive Power Flow with Voltage Constraints". In: *arXiv 1303.7173* (2013).

[43] S. Bolognani, G. Cavraro, R. Carli, and S. Zampieri. "Distributed Reactive Power Feedback Control for Voltage Regulation and Loss Minimization". In: *IEEE Transactions on Automatic Control* 60.4 (Apr. 2015), pp. 966–981.

[44] S. Bolognani and S. Zampieri. "On the Existence and Linear Approximation of the Power Flow Solution in Power Distribution Networks". In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016), pp. 163–172.

[45] J. Bolte, A. Daniilidis, and A. Lewis. “The Lojasiewicz Inequality for Nonsmooth Subanalytic Functions with Applications to Subgradient Dynamical Systems”. In: *SIAM Journal on Optimization* 17.4 (Jan. 2007), pp. 1205–1223.

[46] J. F. Bonnans. *Convex and Stochastic Optimization*. Universitext. Cham, Switzerland: Springer, 2019.

[47] J. F. Bonnans, R. Cominetti, and A. Shapiro. “Sensitivity Analysis of Optimization Problems Under Second Order Regular Constraints”. In: *Mathematics of Operations Research* 23.4 (Nov. 1998), pp. 806–831.

[48] W. M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Vol. 120. Gulf Professional Publishing, 2003.

[49] S. Bose, S. H. Low, T. Teeraratkul, and B. Hassibi. “Equivalent Relaxations of Optimal Power Flow”. In: *IEEE Transactions on Automatic Control* 60.3 (Mar. 2014), pp. 729–742.

[50] R. K. Brayton and J. K. Moser. “A Theory of Nonlinear Networks - I”. In: *Quarterly of Applied Mathematics* 22.1 (1964).

[51] R. W. Brockett. “Dynamical Systems That Sort Lists, Diagonalize Matrices and Solve Linear Programming Problems”. In: *IEEE Conference on Decision and Control (CDC)*. Vol. 1. Austin, TX, Dec. 1988, pp. 799–803.

[52] B. Brogliato, A. Daniilidis, C. Lemaréchal, and V. Acary. “On the Equivalence between Complementarity Systems, Projected Systems and Differential Inclusions”. In: *Systems & Control Letters* 55.1 (Jan. 2006), pp. 45–51.

[53] F. D. Brunner, H. Dürr, and C. Ebenbauer. “Feedback Design for Multi-Agent Systems: A Saddle Point Approach”. In: *IEEE Conference on Decision and Control (CDC)*. Maui, HI, Dec. 2012, pp. 3783–3789.

[54] A. Y. Burtscher. “Length Structures on Manifolds with Continuous Riemannian Metrics”. In: *New York J. Math* 21 (2015), pp. 273–296.

[55] C. Castaing and M. D. P. Monteiro Marques. “Evolution Problems Associated with Nonconvex Closed Moving Sets with Bounded Variation”. In: *Portugaliae Mathematica* 53.1 (1996), pp. 73–88.

[56] B. Chachuat, B. Srinivasan, and D. Bonvin. "Adaptation Strategies for Real-Time Optimization". In: *Computers & Chemical Engineering* 33.10 (Oct. 2009), pp. 1557–1567.

[57] C.-Y. Chang, M. Colombino, J. Cortes, and E. Dall'Anese. "Saddle-Flow Dynamics for Distributed Feedback-Based Optimization". In: *IEEE Control Systems Letters* 3.4 (Oct. 2019), pp. 948–953.

[58] A. Cherukuri, E. Mallada, S. Low, and J. Cortés. "The Role of Convexity on Saddle-Point Dynamics: Lyapunov Function and Robustness". In: *IEEE Transactions on Automatic Control* 63.8 (2017), pp. 2449–2464.

[59] A. Cherukuri, B. Gharesifard, and J. Cortés. "Saddle-Point Dynamics: Conditions for Asymptotic Stability of Saddle Points". In: *SIAM Journal on Control and Optimization* 55.1 (Jan. 2017), pp. 486–511.

[60] A. Cherukuri, E. Mallada, and J. Cortés. "Asymptotic Convergence of Constrained Primal–Dual Dynamics". In: *Systems & Control Letters* 87 (Jan. 2016), pp. 10–15.

[61] CIGRÉ. *Coordinated Voltage Control in Transmission Networks*. Tech. rep. Paris, France: Task Force C4.602, Feb. 2007.

[62] CIGRÉ. *Defense Plan against Extreme Contingencies*. Tech. rep. Paris, France: Task Force C2.02.24, Apr. 2007.

[63] CIGRÉ. *Ancillary Services: An Overview of International Practices*. Tech. rep. Paris, France: Working Group C5.06, Oct. 2010.

[64] CIGRÉ. *Lessons Learnt from Recent Emergencies and Blackout Incidents*. Tech. rep. Paris, France: Working Group C2.21, Jan. 2015.

[65] F. Coito, J. Lemos, and S. Alves. "Stochastic Extremum Seeking in the Presence of Constraints". In: *16th IFAC World Congress*. Prague, Czech Republic, Aug. 2005, pp. 276–281.

[66] M. Colombino, E. Dall'Anese, and A. Bernstein. "Online Optimization as a Feedback Controller: Stability and Tracking". In: *IEEE Transactions on Control of Network Systems* (2019).

[67] M. Colombino, J. W. Simpson-Porco, and A. Bernstein. "Towards Robustness Guarantees for Feedback-Based Optimization". In: *arXiv 1905.07363 [math]* (May 2019).

[68] B. Cornet. “Existence of Slow Solutions for a Class of Differential Inclusions”. In: *Journal of Mathematical Analysis and Applications* 96.1 (Oct. 1983), pp. 130–147.

[69] S. Costello, G. François, D. Bonvin, and A. Marchetti. “Modifier Adaptation for Constrained Closed-Loop Systems”. In: *19th IFAC World Congress*. Cape Town, South Africa, Aug. 2014, pp. 11080–11086.

[70] S. Cyrus, B. Hu, B. Van Scoy, and L. Lessard. “A Robust Accelerated Optimization Algorithm for Strongly Convex Functions”. In: *American Control Conference (ACC)*. Milwaukee, WI, June 2018, pp. 1376–1381.

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

[72] A. Daniilidis. “Gradient Dynamical Systems, Tame Optimization and Applications”. Lecture Notes "Spring School on Variational Analysis". Paseky nad Jizerou, Czech Republic, Apr. 2009.

[73] C. De Persis and S. Grammatico. “Distributed Averaging Integral Nash Equilibrium Seeking on Networks”. In: *Automatica* 110 (Dec. 2019), p. 108548.

[74] D. DeHaan and M. Guay. “Extremum-Seeking Control of State-Constrained Nonlinear Systems”. In: *Automatica* 41.9 (Sept. 2005), pp. 1567–1574.

[75] O. Devolder, F. Glineur, and Y. Nesterov. “First-Order Methods of Smooth Convex Optimization with Inexact Oracle”. In: *Mathematical Programming* 146.1-2 (Aug. 2014), pp. 37–75.

[76] N. K. Dhingra, S. Z. Khong, and M. R. Jovanovic. “The Proximal Augmented Lagrangian Method for Nonsmooth Composite Optimization”. In: *IEEE Transactions on Automatic Control* 64.7 (July 2019), pp. 2861–2868.

[77] M. Diehl, H. Bock, and J. Schlöder. “A Real-Time Iteration Scheme for Nonlinear Optimization in Optimal Feedback Control”. In: *SIAM Journal on Control and Optimization* 43.5 (Jan. 2005), pp. 1714–1736.

[78] M. Diehl, H. Bock, H. Diedam, and P.-B. Wieber. “Fast Direct Multiple Shooting Algorithms for Optimal Robot Control”. In: *Fast Motions in Biomechanics and Robotics*. Vol. 340. Berlin Heidelberg, Germany: Springer, 2006, pp. 65–93.

[79] M. Diehl, R. Findeisen, F. Allgöwer, H. Bock, and J. Schlöder. “Nominal Stability of Real-Time Iteration Scheme for Nonlinear Model Predictive Control”. In: *IEE Proceedings - Control Theory and Applications* 152.3 (May 2005), pp. 296–308.

[80] M. Diehl, H. G. Bock, J. P. Schlöder, R. Findeisen, Z. Nagy, and F. Allgöwer. “Real-Time Optimization and Nonlinear Model Predictive Control of Processes Governed by Differential-Algebraic Equations”. In: *Journal of Process Control* 12.4 (June 2002), pp. 577–585.

[81] M. Diehl, R. Findeisen, and F. Allgöwer. “A Stabilizing Real-Time Implementation of Nonlinear Model Predictive Control”. In: *Real-Time PDE-Constrained Optimization*. Philadelphia, PA: SIAM, Jan. 2007, pp. 25–52.

[82] M. Diehl, R. Findeisen, S. Schwarzkopf, I. Uslu, F. Allgöwer, H. G. Bock, E. D. Gilles, and J. P. Schlöder. “An Efficient Algorithm for Nonlinear Model Predictive Control of Large-Scale Systems Part II: Experimental Evaluation for a Distillation Column”. In: *at - Automatisierungstechnik* 51.1 (Jan. 2003), pp. 22–29.

[83] I. Dobson. “The Irrelevance of Electric Power System Dynamics for the Loading Margin to Voltage Collapse and Its Sensitivities”. In: *Nonlinear Theory and Its Applications, IEICE* 2.3 (2011), pp. 263–280.

[84] A. Dontchev and F. Lempio. “Difference Methods for Differential Inclusions: A Survey”. In: *SIAM Review* 34.2 (June 1992), pp. 263–294.

[85] F. Dörfler and S. Grammatico. “Gather-and-Broadcast Frequency Control in Power Systems”. In: *Automatica* 79 (May 2017), pp. 296–305.

[86] H.-B. Dürr, M. Krstić, A. Scheinker, and C. Ebenbauer. “Extremum Seeking for Dynamic Maps Using Lie Brackets and Singular Perturbations”. In: *Automatica* 83 (Sept. 2017), pp. 91–99.

[87] H.-B. Dürr, M. S. Stanković, C. Ebenbauer, and K. H. Johansson. "Lie Bracket Approximation of Extremum Seeking Systems". In: *Automatica* 49.6 (June 2013), pp. 1538–1552.

[88] H.-B. Dürr, M. S. Stanković, K. H. Johansson, and C. Ebenbauer. "Extremum Seeking on Submanifolds in the Euclidian Space". In: *Automatica* 50.10 (Oct. 2014), pp. 2591–2596.

[89] H.-B. Dürr, C. Zeng, and C. Ebenbauer. "Saddle Point Seeking for Convex Optimization Problems". In: *9th IFAC Symposium on Nonlinear Control Systems*. Tolouse, France, Sept. 2013, pp. 540–545.

[90] J. Edmond and L. Thibault. "Relaxation of an Optimal Control Problem Involving a Perturbed Sweeping Process". In: *Mathematical Programming* 104.2-3 (Nov. 2005), pp. 347–373.

[91] M. Ellis, H. Durand, and P. D. Christofides. "A Tutorial Review of Economic Model Predictive Control Methods". In: *Journal of Process Control* 24.8 (Aug. 2014), pp. 1156–1178.

[92] M. Farivar and S. H. Low. "Branch Flow Model: Relaxations and Convexification." In: *IEEE Conference on Decision and Control (CDC)*. Maui, HI, 2012, pp. 3672–3679.

[93] T. Faulwasser, L. Grüne, and M. A. Müller. "Economic Nonlinear Model Predictive Control". In: *Foundations and Trends in Systems and Control* 5.1 (2018).

[94] L. Faybusovich. "Dynamical Systems Which Solve Optimization Problems with Linear Constraints". In: *IMA Journal of Mathematical Control and Information* 8.2 (1991), pp. 135–149.

[95] M. Fazlyab, S. Paternain, V. M. Preciado, and A. Ribeiro. "Interior Point Method for Dynamic Constrained Optimization in Continuous Time". In: *American Control Conference (ACC)*. Boston, MA, July 2016, pp. 5612–5618.

[96] M. Fazlyab, S. Paternain, V. M. Preciado, and A. Ribeiro. "Prediction-Correction Interior-Point Method for Time-Varying Convex Optimization". In: *IEEE Transactions on Automatic Control* 63.7 (July 2018), pp. 1973–1986.

[97] J. Feiling, A. Zeller, and C. Ebenbauer. "Derivative-Free Optimization Algorithms Based on Non-Commutative Maps". In: *IEEE Control Systems Letters* 2.4 (Oct. 2018), pp. 743–748.

[98] A. V. Fiacco. “Sensitivity Analysis for Nonlinear Programming Using Penalty Methods”. In: *Mathematical Programming* 10.1 (Dec. 1976), pp. 287–311.

[99] A. V. Fiacco and Y. Ishizuka. “Sensitivity and Stability Analysis for Nonlinear Programming”. In: *Annals of Operations Research* 27.1 (Dec. 1990), pp. 215–235.

[100] A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides: Control Systems*. Mathematics and Its Applications (Soviet Series). Dordrecht, The Netherlands: Springer, 1988.

[101] N. Fischer, R. Kamalapurkar, and W. E. Dixon. “LaSalle-Yoshizawa Corollaries for Nonsmooth Systems”. In: *IEEE Transactions on Automatic Control* 58.9 (Sept. 2013), pp. 2333–2338.

[102] S. Fitzpatrick and R. R. Phelps. “Differentiability of the Metric Projection in Hilbert Space”. In: *Transactions of the American Mathematical Society* 270.2 (Apr. 1982), pp. 483–501.

[103] G. Francois and D. Bonvin. “Measurement-Based Real-Time Optimization of Chemical Processes”. In: *Advances in Chemical Engineering* 43 (2013).

[104] G. François, B. Srinivasan, and D. Bonvin. “Use of Measurements for Enforcing the Necessary Conditions of Optimality in the Presence of Constraints and Uncertainty”. In: *Journal of Process Control* 15.6 (Sept. 2005), pp. 701–712.

[105] S. Frank and S. Rebennack. “An Introduction to Optimal Power Flow: Theory, Formulation, and Examples”. In: *IIE Transactions* 48.12 (Dec. 2016), pp. 1172–1197.

[106] S. Frank, I. Steponavice, and S. Rebennack. “Optimal Power Flow: A Bibliographic Survey I: Formulations and Deterministic Methods”. In: *Energy Systems* 3.3 (Sept. 2012), pp. 221–258.

[107] S. Frank, I. Steponavice, and S. Rebennack. “Optimal Power Flow: A Bibliographic Survey II: Non-Deterministic and Hybrid Methods”. In: *Energy Systems* 3.3 (Sept. 2012), pp. 259–289.

[108] G. F. Franklin, J. D. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Sixth. Upper Saddle River, NJ: Pearson, 2010.

[109] P. Frihauf, M. Krstic, and T. Basar. “Nash Equilibrium Seeking in Noncooperative Games”. In: *IEEE Transactions on Automatic Control* 57.5 (May 2012), pp. 1192–1207.

[110] P. Frihauf, M. Krstic, and T. Başar. "Finite-Horizon LQ Control for Unknown Discrete-Time Linear Systems via Extremum Seeking". In: *European Journal of Control* 19.5 (Sept. 2013), pp. 399–407.

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

[112] W. Gao and S. Engell. "Iterative Set-Point Optimization of Batch Chromatography". In: *Computers & Chemical Engineering* 29.6 (May 2005), pp. 1401–1409.

[113] G. Garrigos. "Descent Dynamical Systems and Algorithms for Tame Optimization, and Multi-Objective Problems". PhD thesis. Montpellier, France: Université de Montpellier, Nov. 2015.

[114] E. Ghadimi, H. R. Feyzmahdavian, and M. Johansson. "Global Convergence of the Heavy-Ball Method for Convex Optimization". In: *European Control Conference (ECC)*. Linz, Austria, July 2015, pp. 310–315.

[115] A. Ghaffari, M. Krstic, and S. Seshagiri. "Power Optimization and Control in Wind Energy Conversion Systems Using Extremum Seeking". In: *IEEE Transactions on Control Systems Technology* 22.5 (Sept. 2014), pp. 1684–1695.

[116] A. Ghaffari, S. Seshagiri, and M. Krstić. "Multivariable Maximum Power Point Tracking for Photovoltaic Micro-Converters Using Extremum Seeking". In: *Control Engineering Practice* 35 (Feb. 2015), pp. 83–91.

[117] R. Goebel. "Stability and Robustness for Saddle-Point Dynamics through Monotone Mappings". In: *Systems & Control Letters* 108 (Oct. 2017), pp. 16–22.

[118] R. Goebel, R. G. Sanfelice, and A. R. Teel. *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton, NJ: Princeton University Press, 2012.

[119] G. Grégory, F. Ois, and D. Bonvin. "Use of Transient Measurements for the Optimization of Steady-State Performance via Modifier Adaptation". In: *Industrial Engineering Chemistry Research* 53 (2014), pp. 5148–5159.

[120] V. Grushkovskaya, A. Zuyev, and C. Ebenbauer. "On a Class of Generating Vector Fields for the Extremum Seeking Problem: Lie Bracket Approximation and Stability Properties". In: *Automatica* 94 (Aug. 2018), pp. 151–160.

[121] M. Guay, D. Dochain, and M. Perrier. "Adaptive Extremum Seeking Control of Continuous Stirred Tank Bioreactors with Unknown Growth Kinetics". In: *Automatica* 40.5 (May 2004), pp. 881–888.

[122] M. Guay and T. Zhang. "Adaptive Extremum Seeking Control of Nonlinear Dynamic Systems with Parametric Uncertainties". In: *Automatica* 39.7 (July 2003), pp. 1283–1293.

[123] E. C. Hall and R. M. Willett. "Online Convex Optimization in Dynamic Environments". In: *IEEE Journal of Selected Topics in Signal Processing* 9.4 (June 2015), pp. 647–662.

[124] P. Hartman. "On the Local Uniqueness of Geodesics". In: *American Journal of Mathematics* 72.4 (1950), pp. 723–730.

[125] W. P. M. H. Heemels, M. K. Camlibel, and M. F. Heertjes. "Oblique Projected Dynamical Systems and Incremental Stability Under State Constraints". In: *IEEE Control Systems Letters* 4.4 (Oct. 2020), pp. 1060–1065.

[126] U. Helmke and J. B. Moore. *Optimization and Dynamical Systems*. Second. Communications and Control Engineering. London, UK: Springer, 1996.

[127] C. Henry. "An Existence Theorem for a Class of Differential Equations with Multivalued Right-Hand Side". In: *Journal of Mathematical Analysis and Applications* 41.1 (Jan. 1973), pp. 179–186.

[128] D. J. Hill and I. M. Y. Mareels. "Stability Theory for Differential / Algebraic Systems with Application to Power Systems". In: *IEEE Transactions on Circuits and Systems* 37.11 (Nov. 1990), pp. 1416–1423.

[129] F. S. Hillier and G. J. Lieberman. *Introduction to Operations Research*. Seventh. New York, NY: McGraw-Hill, 2001.

[130] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms I. Fundamentals*. Second. Grundlehren Der Mathematischen Wissenschaften 305. Berlin Heidelberg, Germany: Springer, 1996.

[131] J.-B. Hiriart-Urruty and C. Lemaréchal. *Fundamentals of Convex Analysis*. Grundlehren Text Editions. Berlin Heidelberg, Germany: Springer, 2012.

[132] T. Holding and I. Lestas. "On the Convergence to Saddle Points of Concave-Convex Functions, the Gradient Method and Emergence of Oscillations". In: *IEEE Conference on Decision and Control (CDC)*. Los Angeles, CA, Dec. 2014, pp. 1143–1148.

[133] S. Hosseini and M. Pouryayevali. "On the Metric Projection onto Prox-Regular Subsets of Riemannian Manifolds". In: *Proceedings of the American Mathematical Society* 141.1 (2013), pp. 233–244.

[134] M. Huneault and F. D. Galiana. "A Survey of the Optimal Power Flow Literature". In: *IEEE Transactions on Power Systems* 6.2 (May 1991), pp. 762–770.

[135] D. M. Hustig-Schultz and R. G. Sanfelice. "A Robust Hybrid Heavy Ball Algorithm for Optimization with High Performance". In: *American Control Conference (ACC)*. Philadelphia, PA, July 2019, pp. 151–156.

[136] A. Ioffe. "An Invitation to Tame Optimization". In: *SIAM Journal on Optimization* 19.4 (Jan. 2009), pp. 1894–1917.

[137] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan. "Online Optimization : Competing with Dynamic Comparators". In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. San Diego, CA, 2015, pp. 398–406.

[138] K. Jittorntrum. "Sequential Algorithms in Nonlinear Programming". PhD thesis. Canberra, Australia: Australian National University, July 1978.

[139] K. Jittorntrum. "Solution Point Differentiability without Strict Complementarity in Nonlinear Programming". In: *Sensitivity, Stability and Parametric Analysis*. Vol. 21. Mathematical Programming Studies. Berlin Heidelberg, Germany: Springer, 1984, pp. 127–138.

[140] A. Jokic, M. Lazar, and van den P.P.J. Bosch. "On Constrained Steady-State Regulation : Dynamic KKT Controllers". In: *IEEE Transactions on Automatic Control* 54.9 (2009), pp. 2250–2254.

[141] A. Jokić, M. Lazar, and P. van den Bosch. "Real-Time Control of Power Systems Using Nodal Prices". In: *International Journal of Electrical Power & Energy Systems* 31.9 (Oct. 2009), pp. 522–530.

[142] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability". In: *Journal of the Operational Research Society* 49.3 (Mar. 1998), pp. 237–252.

[143] H. K. Khalil. *Nonlinear Systems*. Third. Upper Saddle River, NJ: Prentice Hall, 2002.

[144] O. O. Khamisov, T. S. Chernova, J. W. Bialek, and S. H. Low. "Corrective Control: Stability Analysis of Unified Controller Combining Frequency Control and Congestion Management". In: *NEIS Conference* (2018), pp. 15–20.

[145] N. Killingsworth, S. Aceves, D. Flowers, F. Espinosa-Loza, and M. Krstic. "HCCI Engine Combustion-Timing Control: Optimizing Gains and Fuel Consumption Via Extremum Seeking". In: *IEEE Transactions on Control Systems Technology* 17.6 (Nov. 2009), pp. 1350–1361.

[146] P. Kokotovic, H. Khalil, and J. O'Reilly. *Singular Perturbation Methods in Control: Analysis and Design*. Classics in Applied Mathematics 25. Philadelphia, PA: SIAM, 1999.

[147] A. S. Kolarijani, P. M. Esfahani, and T. Keciczky. "Fast Gradient-Based Methods with Exponential Rate: A Hybrid Control Framework". In: *International Conference on Machine Learning (ICML)*. Stockholm, Sweden, 2018, pp. 2728–2736.

[148] A. S. Kolarijani, P. M. Esfahani, and T. Keviczky. "Continuous-Time Accelerated Methods via a Hybrid Control Lens". In: *arXiv 1807.07805 [cs, math]* (Sept. 2019).

[149] T. Kose. "Solutions of Saddle Value Problems by Differential Equations". In: *Econometrica* 24.1 (Jan. 1956), pp. 59–70.

[150] W. Krichene, A. Bayen, and P. L. Bartlett. "Accelerated Mirror Descent in Continuous and Discrete Time". In: *Advances in Neural Information Processing Systems (NIPS)*. Montreal, Canada, 2015, pp. 2845–2853.

[151] M. Krstic, A. Ghaffari, and S. Seshagiri. "Extremum Seeking for Wind and Solar Energy Applications". In: *World Congress on Intelligent Control and Automation*. Shenyang, China, June 2014, pp. 6184–6193.

[152] M. Krstic and H.-H. Wang. "Stability of Extremum Seeking Feedback for General Nonlinear Dynamic Systems". In: *Automatica* 36 (2000), pp. 595–601.

[153] J. B. Kruskal. "Two Convex Counterexamples: A Discontinuous Function and a Non-Differentiable Nearest-Point Mapping". In: *Proceedings of the American Mathematical Society* 23 (1969), pp. 697–703.

[154] M. Kunze and M. D. P. M. Marques. "An Introduction to Moreau's Sweeping Process". In: *Impacts in Mechanical Systems*. Lecture Notes in Physics. Berlin Heidelberg: Springer, 2000, pp. 1–60.

[155] H. Kwatny, A. Pasrija, and L. Bahar. "Static Bifurcations in Electric Power Networks: Loss of Steady-State Stability and Voltage Collapse". In: *IEEE Transactions on Circuits and Systems* 33.10 (Oct. 1986), pp. 981–991.

[156] J. LaSalle and Z. Artstein. *The Stability of Dynamical Systems*. Philadelphia, PA: SIAM, 1976.

[157] J. Lavaei and S. H. Low. "Zero Duality Gap in Optimal Power Flow Problem". In: *IEEE Transactions on Power Systems* 27.1 (Feb. 2012), pp. 92–107.

[158] L. S. P. Lawrence, Z. E. Nelson, E. Mallada, and J. W. Simpson-Porco. "Optimal Steady-State Control for Linear Time-Invariant Systems". In: *IEEE Conference on Decision and Control (CDC)*. Miami Beach, FL, Dec. 2018, pp. 3251–3257.

[159] L. S. P. Lawrence, J. W. Simpson-Porco, and E. Mallada. "Linear-Convex Optimal Steady-State Control". In: *arXiv 1810.12892 [math]* (May 2020).

[160] J. D. Lee, Y. Sun, and M. A. Saunders. "Proximal Newton-Type Methods for Minimizing Composite Functions". In: *SIAM Journal on Optimization* 24.3 (Jan. 2014), pp. 1420–1443.

[161] J. M. Lee. *Riemannian Manifolds - An Introduction to Curvature*. First. Graduate Texts in Mathematics 176. New York, NY: Springer, 1997.

[162] J. M. Lee. *Introduction to Smooth Manifolds*. Vol. 218. Graduate Texts in Mathematics. New York, NY: Springer, 2012.

[163] A. Lesage-Landry, I. Shames, and J. A. Taylor. "Predictive Online Convex Optimization". In: *Automatica* 113 (Mar. 2020), p. 108771.

[164] L. Lessard, B. Recht, and A. Packard. “Analysis and Design of Optimization Algorithms Via Integral Quadratic Constraints”. In: *SIAM Journal on Optimization* 26.1 (Jan. 2016), pp. 57–95.

[165] N. Li, C. Zhao, and L. Chen. “Connecting Automatic Generation Control and Economic Dispatch from an Optimization View”. In: *IEEE Transactions on Control of Network Systems* 3.3 (Sept. 2016), pp. 254–264.

[166] S. H. Low and D. E. Lapsley. “Optimization Flow Control. I. Basic Algorithm and Convergence”. In: *IEEE/ACM Transactions on Networking* 7.6 (Dec. 1999), pp. 861–874.

[167] S. H. Low, F. Paganini, and J. C. Doyle. “Internet Congestion Control”. In: *IEEE Control Systems Magazine* 22.1 (Feb. 2002), pp. 28–43.

[168] S. H. Low. *Analytical Methods for Network Congestion Control*. Synthesis Lectures on Communication Networks 18. Williston, VT: Morgan & Claypool, 2017.

[169] S. H. Low. “Convex Relaxation of Optimal Power Flow; Part I: Formulations and Equivalence”. In: *IEEE Transactions on Control of Network Systems* 1.1 (Mar. 2014), pp. 15–27.

[170] D. G. Luenberger, Y. Ye, et al. *Linear and Nonlinear Programming*. Fourth. Cham, Switzerland: Springer, 1984.

[171] M. Mansour and J. Ellis. “Comparison of Methods for Estimating Real Process Derivatives in On-Line Optimization”. In: *Applied Mathematical Modelling* 27.4 (Apr. 2003), pp. 275–291.

[172] A. Marchetti, B. Chachuat, and D. Bonvin. “Modifier-Adaptation Methodology for Real-Time Optimization”. In: *Industrial & Engineering Chemistry Research* 48.13 (July 2009), pp. 6022–6033.

[173] A. Marchetti, B. Chachuat, and D. Bonvin. “Real-Time Optimization with Estimation of Experimental Gradients”. In: *7th IFAC Symposium on Advanced Control of Chemical Processes*. Tolouse, France, Jan. 2009, pp. 524–529.

[174] N. Mazzi, B. Zhang, and D. S. Kirschen. “An Online Optimization Algorithm for Alleviating Contingencies in Transmission Networks”. In: *arXiv 1709.03965 [cs, math]* (Feb. 2018).

[175] C. D. Meyer. *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA: SIAM, 2000.

[176] H. Mohammadi, M. Razaviyayn, and M. R. Jovanovic. "Variance Amplification of Accelerated First-Order Algorithms for Strongly Convex Quadratic Optimization Problems". In: *IEEE Conference on Decision and Control (CDC)*. Miami Beach, FL, Dec. 2018, pp. 5753–5758.

[177] H. Mohammadi, M. Razaviyayn, and M. R. Jovanovic. "Performance of Noisy Nesterov's Accelerated Method for Strongly Convex Optimization Problems". In: *American Control Conference (ACC)*. Philadelphia, PA, July 2019, pp. 3426–3431.

[178] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro. "Online Optimization in Dynamic Environments: Improved Regret Rates for Strongly Convex Problems". In: *IEEE Conference on Decision and Control (CDC)*. Las Vegas, NV, Dec. 2016, pp. 7195–7201.

[179] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei. "A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems". In: *IEEE Transactions on Smart Grid* 8.6 (Nov. 2017), pp. 2941–2962.

[180] D. K. Molzahn, I. Hiskens, et al. "Moment-Based Relaxation of the Optimal Power Flow Problem". In: *Power Systems Computation Conference (PSCC)*. Wroclaw, Poland, Aug. 2014.

[181] D. K. Molzahn and I. A. Hiskens. "A Survey of Relaxations and Approximations of the Power Flow Equations". In: *Foundations and Trends® in Electric Energy Systems* 4.1-2 (2019), pp. 1–221.

[182] D. K. Molzahn, J. T. Holzer, B. C. Lesieutre, and C. L. DeMarco. "Implementation of a Large-Scale Optimal Power Flow Solver Based on Semidefinite Programming". In: *Power Systems, IEEE Transactions on* 28.4 (2013), pp. 3987–3998.

[183] J. M. Montenbruck, H.-B. Dürr, C. Ebenbauer, and F. Allgöwer. "Extremum Seeking and Obstacle Avoidance on the Special Orthogonal Group". In: *19th IFAC World Congress*. Cape Town, South Africa, 2014, pp. 8229–8234.

[184] J. J. Moreau. "Evolution Problem Associated with a Moving Convex Set in a Hilbert Space". In: *Journal of Differential Equations* 26.3 (Dec. 1977), pp. 347–374.

[185] M. Muehlebach and M. I. Jordan. "A Dynamic System's Perspective on Nesterov's Accelerated Gradient Method". In: *International Conference on Machine Learning (ICML)*. Los Angeles, CA, 2019.

[186] A. Nagurney and D. Zhang. *Projected Dynamical Systems and Variational Inequalities with Applications*. First. New York, NY: Springer, 1996.

[187] Z. E. Nelson and E. Mallada. "An Integral Quadratic Constraint Framework for Real-Time Steady-State Optimization of Linear Time-Invariant Systems". In: *American Control Conference (ACC)*. Milwaukee, WI, June 2018, pp. 597–603.

[188] A. S. Nemirovsky and D. B. Yudin. "Problem Complexity and Method Efficiency in Optimization". In: (1983).

[189] Y. Nesterov. "A Method for Solving the Convex Programming Problem with Convergence Rate O(1/K$\hat{2}$)". In: *Soviet Math. Dokl.* 27.2 (1983), pp. 372–376.

[190] Y. Nesterov. *Introductory Lectures on Convex Optimization*. Vol. 87. Applied Optimization. Boston, MA: Springer, 2004.

[191] J. Nocedal and S. J. Wright. *Numerical Optimization*. Second. Series in Operations Research. New York, NY: Springer, 2006.

[192] B. O'Donoghue and E. Candès. "Adaptive Restart for Accelerated Gradient Schemes". In: *Foundations of Computational Mathematics* 15.3 (June 2015), pp. 715–732.

[193] F. Paganini, Zhikui Wang, J. Doyle, and S. Low. "Congestion Control for High Performance, Stability, and Fairness in General Networks". In: *IEEE/ACM Transactions on Networking* 13.1 (Feb. 2005), pp. 43–56.

[194] J. J. Palis and W. De Melo. *Geometric Theory of Dynamical Systems: An Introduction*. New York, NY: Springer, 1982.

[195] P. Panciatici, F. Bena, P. Pruvot, N. Janssens, J. Deuse, and M. Stubbe. "Centralized Voltage Control: A Key Point for Optimal Operation of Power Systems". In: *CIGRÉ Session 1998 39-116*. Paris, 1998, p. 8.

[196] N. Parikh and S. Boyd. "Proximal Algorithms". In: *Foundations and Trends in Optimization* 1.3 (Jan. 2014), pp. 127–239.

[197] J. P. Paul, J. Y. Leost, and J. M. Tesseron. "Survey of the Secondary Voltage Control in France : Present Realization and Investigations". In: *IEEE Transactions on Power Systems* 2.2 (1987), pp. 505–511.

[198] J. Peypouquet. *Convex Optimization in Normed Spaces: Theory, Methods and Examples*. Briefs in Optimization. Cham, Switzerland: Springer, 2015.

[199] R. Poliquin and R. Rockafellar. "Prox-Regular Functions in Variational Analysis". In: *Transactions of the American Mathematical Society* 348.5 (1996), pp. 1805–1838.

[200] B. Polyak. "Some Methods of Speeding up the Convergence of Iteration Methods". In: *USSR Computational Mathematics and Mathematical Physics* 4.5 (Jan. 1964), pp. 1–17.

[201] A. Y. Popkov. "Gradient Methods for Nonstationary Unconstrained Optimization Problems". In: *Automation and Remote Control* 66.6 (June 2005), pp. 883–891.

[202] J. Poveda, M. Benosman, and A. Teel. "Distributed Extremum Seeking in Multi-Agent Systems with Arbitrary Switching Graphs". In: *20th IFAC World Congress*. Toulouse, France, July 2017, pp. 735–740.

[203] J. I. Poveda and N. Li. "Inducing Uniform Asymptotic Stability in Non-Autonomous Accelerated Optimization Dynamics via Hybrid Regularization". In: *IEEE Conference on Decision and Control (CDC)*. Nice, France, Dec. 2019, pp. 3000–3005.

[204] J. I. Poveda and A. R. Teel. "A Framework for a Class of Hybrid Extremum Seeking Controllers with Dynamic Inclusions". In: *Automatica* 76 (Feb. 2017), pp. 113–126.

[205] *Power Systems Test Case Archive, University of Washington, Electrical Engineering*. University of Washington, Electrical Engineering Dept. Online: http://www.ee.washington.edu/research/pstca. 2007.

[206] R. Quirynen, S. Gros, and M. Diehl. "Inexact Newton-Type Optimization with Iterated Sensitivities". In: *SIAM Journal on Optimization* 28.1 (Jan. 2018), pp. 74–95.

[207] S. Rahili and W. Ren. "Distributed Continuous-Time Convex Optimization With Time-Varying Cost Functions". In: *IEEE Transactions on Automatic Control* 62.4 (Apr. 2017), pp. 1590–1605.

[208] S. Rahili, W. Ren, and S. Ghapani. “Distributed Convex Optimization of Time-Varying Cost Functions with Swarm Tracking Behavior for Continuous-Time Dynamics”. In: *arXiv 1507.04773 [math]* (July 2015).

[209] J. B. Rawlings, D. Angeli, and C. N. Bates. “Fundamentals of Economic Model Predictive Control”. In: *IEEE Conference on Decision and Control (CDC)*. Maui, HI, Dec. 2012, pp. 3851–3861.

[210] S. M. Robinson. “Perturbed Kuhn-Tucker Points and Rates of Convergence for a Class of Nonlinear-Programming Algorithms”. In: *Mathematical Programming* 7.1 (Dec. 1974), pp. 1–16.

[211] R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Third. Grundlehren Der Mathematischen Wissenschaften 317. Berlin Heidelberg, Germany: Springer, 2009.

[212] H. Royden and P. Fitzpatrick. *Real Analysis*. Fourth. Upper Saddle River, NJ: Pearson, 1988.

[213] W. Rudin. *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics. New York, NY: McGraw-Hill, 1976.

[214] E. P. Ryan. “An Integral Invariance Principle for Differential Inclusions with Applications in Adaptive Control”. In: *SIAM Journal on Control and Optimization* 36.3 (May 1998), pp. 960–980.

[215] A. Saberi and H. Khalil. “Quadratic-Type Lyapunov Functions for Singularly Perturbed Systems”. In: *IEEE Transactions on Automatic Control* 29.6 (June 1984), pp. 542–550.

[216] R. G. Sanfelice and A. R. Teel. “Dynamical Properties of Hybrid Systems Simulators”. In: *Automatica* 46.2 (Feb. 2010), pp. 239–248.

[217] S. Sastry. *Nonlinear Systems*. Ed. by J. E. Marsden, L. Sirovich, and S. Wiggins. Vol. 10. Interdisciplinary Applied Mathematics. New York, NY: Springer, 1999.

[218] S. Shakkottai and R. Srikant. “Network Optimization and Control”. In: *Foundations and Trends® in Networking* 2.3 (2007), pp. 271–379.

[219] A. Shapiro. “Sensitivity Analysis of Nonlinear Programs and Differentiability Properties of Metric Projections”. In: *SIAM Journal on Control and Optimization* 26.3 (May 1988), pp. 628–645.

[220] A. Shapiro. "Directionally Nondifferentiable Metric Projection". In: *Journal of Optimization Theory and Applications* 81.1 (Apr. 1994), pp. 203–204.

[221] A. Shapiro. "Differentiability Properties of Metric Projections onto Convex Sets". In: *Journal of Optimization Theory and Applications* 169.3 (June 2016), pp. 953–964.

[222] A. Simonetto. "Time-Varying Convex Optimization via Time-Varying Averaged Operators". In: *arXiv 1704.07338 [math]* (Nov. 2017).

[223] A. Simonetto. "Time-Varying Optimization: Algorithms and Engineering Applications". In: *arXiv 1807.07032 [math]* (Aug. 2018).

[224] A. Simonetto and E. Dall'Anese. "Prediction-Correction Algorithms for Time-Varying Constrained Optimization". In: *IEEE Transactions on Signal Processing* 65.20 (Oct. 2017), pp. 5481–5494.

[225] A. Simonetto, E. Dall'Anese, S. Paternain, G. Leus, and G. B. Giannakis. "Time-Varying Convex Optimization: Time-Structured Algorithms and Applications". In: *arXiv 2006.08500 [cs, eess, math]* (June 2020).

[226] A. Simonetto, A. Koppel, A. Mokhtari, G. Leus, and A. Ribeiro. "Decentralized Prediction-Correction Methods for Networked Time-Varying Convex Optimization". In: *arXiv 1602.01716 [cs, math]* (Nov. 2016).

[227] J. W. Simpson-Porco. "Analysis and Synthesis of Low-Gain Integral Controllers for Nonlinear Systems with Application to Feedback-Based Optimization". In: *arXiv 2003.01348 [math]* (Mar. 2020).

[228] S. Smale. "On the Mathematical Foundations of Electrical Circuit Theory". In: *Journal of Differential Geometry* 7.1-2 (1972), pp. 193–210.

[229] E. D. Sontag. "Comments on Integral Variants of ISS". In: *Systems & Control Letters* 34.1-2 (May 1998), pp. 93–100.

[230] J. E. Spingarn and R. T. Rockafellar. "The Generic Nature of Optimality Conditions in Nonlinear Programming". In: *Mathematics of Operations Research* 4.4 (1979), pp. 425–430.

[231] M. S. Stankovic, K. H. Johansson, and D. M. Stipanovic. "Distributed Seeking of Nash Equilibria With Applications to Mobile Sensor Networks". In: *IEEE Transactions on Automatic Control* 57.4 (Apr. 2012), pp. 904–919.

[232] M. S. Stankovic and D. M. Stipanovic. “Discrete Time Extremum Seeking by Autonomous Vehicles in a Stochastic Environment”. In: *IEEE Conference on Decision and Control (CDC) and Chinese Control Conference*. Shanghai, China, Dec. 2009, pp. 4541–4546.

[233] M. S. Stanković and D. M. Stipanović. “Extremum Seeking under Stochastic Noise and Applications to Mobile Sensors”. In: *Automatica* 46.8 (Aug. 2010), pp. 1243–1251.

[234] T. W. Stegink, T. V. Damme, and C. D. Persis. “Convergence of Projected Primal-Dual Dynamics with Applications in Data Centers”. In: *7th IFAC Workshop on Distributed Estimation and Control in Networked Systems (NECSYS)*. Groningen, The Netherlands, Aug. 2018, pp. 88–93.

[235] B. Stott, J. Jardim, and O. Alsac. “DC Power Flow Revisited”. In: *IEEE Transactions on Power Systems* 24.3 (Aug. 2009), pp. 1290–1300.

[236] A. Stuart and A. Humphries. *Dynamical Systems and Numerical Analysis*. Cambridge Monographs on Applied and Computational Mathematics Bd. 8. Cambridge, UK: Cambridge University Press, 1998.

[237] W. Su, S. Boyd, and E. Candes. “A Differential Equation for Modeling Nesterov’s Accelerated Gradient Method: Theory and Insights”. In: *Advances in Neural Information Processing Systems (NIPS)*. Montreal, Canada, 2014, pp. 2510–2518.

[238] Y. Tan, W. H. Moase, C. Manzie, D. Nesic, and I. M. Y. Mareels. “Extremum Seeking from 1922 to 2010”. In: *Proceedings of the 29th Chinese Control Conference*. Beijing, China, July 2010, pp. 14–26.

[239] Y. Tan, D. Nešić, and I. Mareels. “On Non-Local Stability Properties of Extremum Seeking Control”. In: *Automatica* 42.6 (June 2006), pp. 889–903.

[240] A. Tang, J. Wang, S. H. Low, and M. Chiang. “Equilibrium of Heterogeneous Congestion Control: Existence and Uniqueness”. In: *IEEE/ACM Transactions on Networking* 15.4 (Aug. 2007), pp. 824–837.

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

[242] Y. Tang, E. Dall'Anese, A. Bernstein, and S. Low. "Running Primal-Dual Gradient Method for Time-Varying Nonconvex Problems". In: *arXiv 1812.00613 [math]* (Dec. 2018).

[243] Y. Tang, E. Dall'Anese, A. Bernstein, and S. H. Low. "A Feedback-Based Regularized Primal-Dual Gradient Method for Time-Varying Nonconvex Optimization". In: *IEEE Conference on Decision and Control (CDC)*. Miami Beach, FL, Dec. 2018, pp. 3244–3250.

[244] Y. Tang and S. Low. "Distributed Algorithm for Time-Varying Optimal Power Flow". In: *IEEE Conference on Decision and Control (CDC)*. Melbourne, Australia, Dec. 2017, pp. 3264–3270.

[245] S. Tarbouriech and M. Turner. "Anti-Windup Design: An Overview of Some Recent Advances and Open Problems". In: *IET Control Theory & Applications* 3.1 (Jan. 2009), pp. 1–19.

[246] A. Teel and D. Popovic. "Solving Smooth and Nonsmooth Multivariable Extremum Seeking Problems by the Methods of Nonlinear Programming". In: *American Control Conference (ACC)*. Arlington, VA, 2001, pp. 2394–2399.

[247] A. R. Teel, J. I. Poveda, and J. Le. "First-Order Optimization Algorithms with Resets and Hamiltonian Flows". In: *IEEE Conference on Decision and Control (CDC)*. Nice, France, Dec. 2019, pp. 5838–5843.

[248] L. Thibault. "Sweeping Process with Regular and Nonregular Sets". In: *Journal of Differential Equations* 193.1 (Sept. 2003), pp. 1–26.

[249] P. Tondel, T. A. Johansen, and A. Bemporad. "An Algorithm for Multi-Parametric Quadratic Programming and Explicit MPC Solutions". In: *Automatica* 39.3 (2003), pp. 489–497.

[250] G. Torrisi, S. Grammatico, R. S. Smith, and M. Morari. "A Projected Gradient and Constraint Linearization Method for Nonlinear Model Predictive Control". In: *SIAM Journal on Control and Optimization* 56.3 (Jan. 2018), pp. 1968–1999.

[251] T. Van Cutsem. "Voltage Instability: Phenomena, Countermeasures, and Analysis Methods". In: *Proceedings of the IEEE* 88.2 (2000), pp. 208–227.

[252] T. Van Cutsem and C. Vournas. *Voltage Stability of Electric Power Systems*. Dordrecht, The Netherlands: Springer, 1998.

[253] B. Van Scoy, R. A. Freeman, and K. M. Lynch. "The Fastest Known Globally Convergent First-Order Method for Minimizing Strongly Convex Functions". In: *IEEE Control Systems Letters* 2.1 (Jan. 2018), pp. 49–54.

[254] J. J. P. Veerman. "Navigating Around Convex Sets". In: *The American Mathematical Monthly* 127.6 (May 2020), pp. 504–517.

[255] V. Venets. "Continuous Algorithms for Solution of Convex Optimization Problems and Finding Saddle Points of Contex-Coneave Functions with the Use of Projection Operations". In: *Optimization* 16.4 (Jan. 1985), pp. 519–533.

[256] G. Vinnicombe. "On the Stability of Networks Operating TCP-like Congestion Control". In: *15th IFAC World Congress*. Barcelona, Spain, 2002, pp. 217–222.

[257] G. Wachsmuth. "On LICQ and the Uniqueness of Lagrange Multipliers". In: *Operations Research Letters* 41.1 (Jan. 2013), pp. 78–80.

[258] J. Wang and N. Elia. "A Control Perspective for Centralized and Distributed Convex Optimization". In: *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*. Orlando, FL, Dec. 2011, pp. 3800–3805.

[259] D. X. Wei, C. Jin, S. H. Low, and S. Hegde. "FAST TCP: Motivation, Architecture, Algorithms, Performance". In: *IEEE/ACM Transactions on Networking* 14.6 (2006), pp. 1246–1259.

[260] E. Wei, A. Ozdaglar, and A. Jadbabaie. "A Distributed Newton Method for Network Utility Maximization—Part Ii: Convergence". In: *IEEE Transactions on Automatic Control* 58.9 (Sept. 2013), pp. 2176–2188.

[261] J. T. Wen and M. Arcak. "A Unifying Passivity Framework for Network Flow Control". In: *IEEE Transactions on Automatic Control* 49.2 (Feb. 2004), pp. 162–174.

[262] A. Wibisono, A. C. Wilson, and M. I. Jordan. "A Variational Perspective on Accelerated Methods in Optimization". In: *Proceedings of the National Academy of Sciences* 113.47 (Nov. 2016), E7351–E7358.

[263] A. C. Wilson, B. Recht, and M. I. Jordan. "A Lyapunov Analysis of Momentum Methods in Optimization". In: *arXiv 1611.02635 [cs, math]* (Nov. 2016).

[264] L. Zaccarian and A. R. Teel. *Modern Anti-Windup Synthesis: Control Augmentation for Actuator Saturation*. Princeton, NJ: Princeton University Press, 2011.

[265] M. Zargham, A. Ribeiroy, and A. Jadbabaie. "Accelerated Dual Descent for Constrained Convex Network Flow Optimization". In: *IEEE Conference on Decision and Control (CDC)*. Florence, Italy, Dec. 2013, pp. 1037–1042.

[266] V. M. Zavala and L. T. Biegler. "Nonlinear Programming Strategies for State Estimation and Model Predictive Control". In: *Nonlinear Model Predictive Control*. Vol. 384. Berlin Heidelberg, Germany: Springer, 2009, pp. 419–432.

[267] C. Zhao, U. Topcu, N. Li, and S. Low. "Design and Stability of Load-Side Primary Frequency Control in Power Systems". In: *IEEE Transactions on Automatic Control* 59.5 (May 2014), pp. 1177–1189.

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

[269] M. Zinkevich. "Online Convex Programming and Generalized Infinitesimal Gradient Ascent". In: *International Conference on Machine Learning (ICML)*. Washington DC, 2003.

Curriculum Vitae

Adrian Robert Hauswirth
born on October 12, 1988 in Altdorf, UR, Switzerland

May'15–Aug'20	**Ph.D.** Information Technology and Electrical Engineering *Swiss Federal Institute of Technology (ETH) Zürich* Switzerland
Sep'12–Feb'15	**M.Sc.** Robotics, Systems and Control *Swiss Federal Institute of Technology (ETH) Zürich* Switzerland
Sep'09–Jun'12	**B.Sc.** Electrical and Electronics Engineering *École Polytechnique Fédérale de Lausanne (EPFL)* Switzerland