

A Lyapunov-like Function for Analysis of Model Predictive Control and Moving
Horizon Estimation

By

Douglas A. Allan

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

Doctor of Philosophy
(Chemical and Biological Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2020

Date of final oral examination: 8/24/2020

The dissertation is approved by the following members of the Final Oral
Committee:

- James B. Rawlings, Professor, Chemical Engineering (University of
California, Santa Barbara)
- Christos Maravelias, Professor, Chemical and Biological Engineering
(University of Wisconsin, Madison)
- Victor M. Zavala, Associate Professor, Chemical and Biological
Engineering (University of Wisconsin, Madison)
- Stephen J. Wright, Professor, Computer Science (University of Wisconsin,
Madison)
- Andrew R. Teel, Professor, Electrical Engineering (University of California,
Santa Barbara)

© Copyright by Douglas A. Allan 2020
All rights reserved.

Ad maiorem Dei gloriam.

CONTENTS

Contents	ii
Nomenclature	v
Abstract	x
Acknowledgements	xii
1 Introduction	1
1.1 Outline	5
2 Nonlinear detectability and incremental input/output to state stability	13
2.1 Literature review	14
2.2 i-IOSS and robustly stable state estimators	17
2.3 Exponential i-IOSS Converse Theorem	24
2.4 Quadratic exp i-IOSS Lyapunov function	32
2.5 Conclusions	33
3 MPC and Steady State Regulation	35
3.1 Infinite-horizon optimal control problem	38
3.2 Semiglobal practical stability with a finite horizon	53
3.3 Stability with a finite horizon	59
3.4 Quadratic stage costs and stabilizability of the linearized system	65
3.5 Practical stability cannot be avoided	73
3.6 Three methodologies, one winner	79
3.6.1 <i>Terminal equality constraint</i>	79
3.6.2 <i>Terminal cost and terminal region</i>	82

3.6.3	<i>Comparison to MPC without terminal conditions</i>	84
3.6.4	<i>Conclusions</i>	86
4	Output Tracking	88
4.1	Literature review	90
4.2	Infinite horizon controller	92
4.3	Stability with a finite horizon	97
4.4	Conclusions	104
5	State estimation	106
5.1	Literature review	108
5.2	Full information estimation	114
5.2.1	<i>Nominal Stability of FIE</i>	116
5.2.2	<i>From stability to robustness</i>	125
5.3	Robust global asymptotic stability of FIE	127
5.3.1	<i>Preliminary Results</i>	136
5.3.2	<i>Main result</i>	140
5.3.3	<i>Discussion</i>	144
5.4	Robust global exponential stability of FIE	150
5.4.1	<i>Assumptions</i>	150
5.4.2	<i>Robust exponential stability</i>	154
5.5	Moving horizon estimation	155
5.5.1	<i>Choice of prior weighting</i>	156
5.5.2	<i>Robust stability of MHE</i>	159
5.5.3	<i>Discussion</i>	160
5.6	Conclusions	162
6	Conclusions and Future Work	163
6.1	Conclusions	163
6.2	Future work	167
6.2.1	<i>Is stabilizability really necessary for robust stability of FIE and MHE?</i>	167
6.2.2	<i>Stability and robustness of MHE for asymptotically i-IOSS systems</i>	167
6.2.3	<i>Smoothed estimates with FIE and smoothing prior update for MHE</i>	169
6.2.4	<i>Return to probability and statistics</i>	170
6.2.5	<i>Probabilistic interpretation of i-IOSS</i>	173

Bibliography

NOMENCLATURE

Acronyms

AD	Algorithmic differentiation
CLF	Control Lyapunov function
EKF	Extended Kalman filter
i-IOSS	Incrementally input/output-to-state stable
ISS	Input-to-state stable
FIE	Full information estimation
KKT	Karush-Kuhn-Tucker
LICQ	Linear independence constraint qualification
LQR	Linear quadratic regulator
MHE	Moving horizon estimation
MPC	Model predictive control
NLP	Nonlinear program
OCP	Optimal control problem
ODE	Ordinary differential equation

PID	Proportional-integral-derivative
RGAS	Robustly globally asymptotically stable
RGES	Robustly globally exponentially stable
RHS	Right-hand side

Notation

$\mathbb{I}_{\geq 0}$	Nonnegative integers
$\mathbb{R}_{\geq 0}$	Nonnegative reals
\mathbb{B}	Unit ball of appropriate dimension
$ \cdot $	Euclidean norm of a vector or 2-norm of a matrix
$ x _{\mathcal{A}}$	Point to set distance $\inf_{y \in \mathcal{A}} x - y $
$\mathbf{0}$	Sequence of zeros of appropriate vector dimension
\mathbf{x}	Sequence of vector-valued variables
$\Delta \mathbf{x}$	Difference of two sequences \mathbf{x}_1 and \mathbf{x}_2
$\mathbf{x}(0 : k)$	Restriction of \mathbf{x} to $(x(j))$ for $j \in \mathbb{I}_{0:k}$
$\ \mathbf{x}\ $	Supremum norm of a sequence \mathbf{x} , $\sup_{k \in \mathbb{I}_{\geq 0}} x(k) $
$\ \mathbf{x}\ _{0:k-1}$	Maximum of first k terms of \mathbf{x} , $\max_{j \in \mathbb{I}_{0:k-1}} x(j) $
$x \frown \mathbf{x}$	The sequence with x inserted at the beginning of \mathbf{x}
\mathbb{X}^{∞}	Countably infinite product of sets \mathbb{X} , i.e., the space of sequences \mathbf{x} with $x(k) \in \mathbb{X}$ for all $k \in \mathbb{I}_{\geq 0}$
$a \oplus b$	Binary maximum of $a, b \in \mathbb{R}$, $\max(a, b)$
$\lceil b \rceil$	Ceiling function, i.e., least integer $k \geq b$.

$A > B$	For symmetric $A, B \in \mathbb{R}^{n \times n}$, the symmetric matrix $A - B$ is positive definite
\mathcal{K}	Class of continuous, strictly increasing functions $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\alpha(0) = 0$
\mathcal{K}_{∞}	Subset of \mathcal{K} functions such that $\lim_{s \rightarrow \infty} \alpha(s) = \infty$
\mathcal{L}	Class of strictly decreasing functions $\theta : \mathbb{I}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\lim_{k \rightarrow \infty} \theta(k) = 0$
\mathcal{KL}	Class of functions $\beta : \mathbb{R}_{\geq 0} \times \mathbb{I}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\beta(\cdot, k)$ is a \mathcal{K} function for fixed k and $\beta(s, \cdot)$ is nonincreasing and satisfies $\lim_{k \rightarrow \infty} \beta(s, k) = 0$ for all fixed s

Variables

$x \in \mathbb{X} \subseteq \mathbb{R}^n$	System state
$u \in \mathbb{U} \subseteq \mathbb{R}^m$	System input
$w \in \mathbb{W} \subseteq \mathbb{R}^g$	Process disturbance
$y \in \mathbb{Y} \subseteq \mathbb{R}^p$	System output
$v \in \mathbb{R}^p$	Measurement disturbance
χ	Decision variable corresponding to x
μ	Decision variable corresponding to u
ω	Decision variable corresponding to w
ν	Decision variable corresponding to v
$f(\cdot)$	System evolution function

$h(\cdot)$	Output function
$\phi(\cdot)$	Open-loop system trajectory
$\Psi_k(\cdot)$	Generic state estimator for $x(k)$
$\Lambda(\cdot)$	i-IOSS Lyapunov function
$\ell(\cdot)$	MPC (Chapter 3 and Chapter 4) or MHE (Chapter 5) stage cost
$\ell^*(x)$	$\min_{u \in \mathcal{U}} \ell(x, u)$
$V_f(\cdot)$	MPC terminal cost
N	MPC or MHE horizon length
$V_k(\cdot)$	MPC unoptimized cost function with horizon length k (Chapters 3 and 4) or FIE unoptimized cost function with k measurements (Chapter 5)
$\mathbb{P}_k(\cdot)$	MPC optimal control problem (Chapters 3 and 4) or FIE optimal estimation problem (Chapter 5)
$V_k^0(\cdot)$	MPC optimal value function (Chapters 3 and 4) or FIE optimal cost function (Chapter 5)
$x^0(j k; x)$	Optimal state forecast at time j for MPC with a horizon length k originating from state x
$u^0(j k; x)$	Optimal input forecast at time j for MPC with a horizon length k originating from state x
\mathcal{X}_N^ρ	Set of x such that $V_N^0(x) \leq \rho$
$\mathcal{K}_N(\cdot)$	MPC optimal control law (possibly set-valued)
$\alpha(\cdot, \cdot)$	Semimetric that is used in bounds (Chapter 4)
\mathbf{x}_r	Reference state trajectory

\mathbf{u}_r	Reference input trajectory
$Y_N(\cdot)$	Sum of optimal cost and i-IOSS Lyapunov function (Chapter 4)
$\ell_x(\cdot)$	Prior weighting for FIE
$\Gamma_k(\cdot)$	Prior weighting for MHE
$\hat{V}_k(\cdot)$	MHE unoptimized cost function at time k with horizon length N
$\hat{\mathbb{P}}_k$	MHE estimation problem
$L_k(\cdot)$	MHE arrival cost
$\hat{V}_k^0(\cdot)$	MHE optimal cost function
$E_{j:k}$	Disturbance “energy” injected into system between times j and k
$Q(\cdot)$	Q function
$Z(\cdot)$	Semidefinite precursor to Q function
\bar{x}	Prior estimate of x
$\hat{x}(j k)$	Estimate of $x(j)$ at time k
$\hat{w}(j k)$	Estimate of disturbance $w(j)$ at time k
$\hat{v}(j k)$	Estimate of disturbance $v(j)$ at time k
$\hat{\mathbf{x}}(k)$	Estimate of sequence $\mathbf{x}(0 : k)$ at time k
$\hat{\mathbf{w}}(k)$	Estimate of sequence $\mathbf{w}(0 : k)$ at time k
$\hat{\mathbf{v}}(k)$	Estimate of sequence $\mathbf{v}(0 : k)$ at time k
\bar{e}	Error in prior estimate $x(0) - \bar{x}$
$e(j k)$	Estimation error $x(j) - \hat{x}(j k)$

ABSTRACT

Model predictive control (MPC) is an established method for control of multiple-input/multiple-output systems. Because of the natural way constraints are handled in the problem formulation, it has found widespread adoption in the process industries. However, in order to accurately forecast system behavior, some method of estimating the system's state is needed. Moving horizon estimation (MHE) is a successful method of state estimation for nonlinear systems. Although MHE has been applied successfully to many types of systems in practice, at present the theory of MHE is still ill-understood. Although recently conditions under which MHE is a robustly stable estimator have been proposed, these conditions are hard to interpret and, for example, it is not clear from existing theory that a longer estimation horizon that includes more data produces a better estimator.

Here, this present gap in understanding of MHE is addressed. A new type of Lyapunov-like function, termed a Q function, is introduced for analysis of MHE, and its idealized counterpart in which the horizon length grows to include all available data, full information estimation (FIE). However, it turns out that this new tool is not just applicable to FIE and MHE. It is possible to analyze the stability of

MPC using *Q* functions as well. In particular, a class of results about the stability of *MPC* without use of stabilizing terminal constraints has grown in recent years. *Q* functions are able to provide a clean exposition of this sort of result that qualitatively produces the same sort of bounds that presently exist in the literature. Furthermore, a problem of tracking a time-varying reference trajectory using *MPC* is proposed to serve as a bridge problem between the setpoint regulation problem familiar to many researchers and *MHE*.

Another tool necessary to understanding the stability of *MHE* is a notion of detectability. For linear systems, detectability is a necessary condition for a stable estimator. For nonlinear systems, a popular detectability assumption is incremental input/output-to-state stability (*i*-IOSS). Here, it is shown that any system that admits a robustly stable state estimator is necessarily *i*-IOSS, and a characterization of *i*-IOSS using a storage function, termed an *i*-IOSS Lyapunov function is provided.

ACKNOWLEDGMENTS

In my academic career, I have been fortunate to have many excellent teachers and collaborators. Chief among these figures is Professor Rawlings. He has been an invaluable sounding board for new ideas during my time at Wisconsin (and Santa Barbara). Through our conversation, I have been exposed to a wide variety of concepts, both from academia and industry. The norm that all students in our group do at least one internship provided me a great opportunity for personal growth. Although my output was somewhat intermittent due to health issues, he remained supportive through my entire tenure in his group. His special attention that Travis Arnold and I were both taken care of after the move to Santa Barbara was critical to our success.

The other members of the Rawlings group have been a source of fellowship and support, whether that support is pedagogical, technical, or moral. I have been blessed to work with people such as Min Yao, Michael Risbeck, Nishith Patel, Travis Arnold, Pratyush Kumar, Cody McAllister, Chris Kuo-LeBlanc, and Steven Kuntz. In particular, I could never have gotten through my PhD without Michael's willingness to share his mastery of both the theoretical and practical aspects of MPC

with me. The visitors to our group: Florian Bayer, Marco Vaccari, and especially Joel Andersson also provided invaluable assistance during my time working with them. Joel was willing to spend time taking me step by step through things like adjoint sensitivity analysis, and without CasADi, we would, in the words of Julier and Uhlmann (2004), be left with “...many pages of dense algebra that must be converted to code...” whenever we required derivatives for optimal control problems.

The members of my committee deserve special thanks for taking time out of their busy schedules preparing for the uncertain prospects of a fall semester (or quarter) in the time of COVID-19. Professor Maravelias already did me a great service by managing my status as a student at the University of Wisconsin after Professor Rawlings moved to UCSB. Professor Wright’s optimization class was excellent, and he was always willing to discuss my research with me after class. Professor Zavala was a pleasure to TA undergraduate controls for. Professor Teel was an insightful collaborator, and I was thrilled when he agreed to be a reader for my thesis despite being from a different institution.

I could not have even got my foot in the door at UW were it not for the research opportunities I received at the University of Minnesota. Professor Daoutidis was a great mentor, and the list of potential academic advisors he provided to me was the reason I applied to UW in the first place. The research I conducted in his lab, under the supervision of Alex Marvin and Srinivas Rangarajan, gave me a first-author publication as an undergraduate, and I am thankful that they gave me that opportunity. Professor Siepmann also generously offered me the chance to conduct research as an undergraduate in his group. Angel Cortes-Morales mentored me

in using Unix and Peng Bai ensured that the data I collected found a home in a publication.

My parents also deserve special thanks. Without their help and encouragement throughout the entire length of my PhD I would never have made it through.

Douglas Allan

Omaha, NE

August 2020

CHAPTER **1**

INTRODUCTION

In simple control systems, like PID controllers, measured variables (“outputs”) are paired with manipulated variables (“inputs”) in such a way that deviation of an output from its desired setpoint results in a corresponding control action to correct it. However, this methodology runs into problems for systems of even modest complexity. When multiple inputs strongly affect multiple outputs simultaneously, controllers can fight one another if their actions are not coordinated, and phenomena like inverse responses can result in a controller taking the wrong initial response to a disturbance.

Model predictive control (MPC) presents a solution to these problems. It uses a process model to predict how a system will respond to an input policy over a certain horizon length, and aggregates the behavior of all outputs into an objective function.

Mathematical optimization then chooses the control policy that produces the best output behavior. The first element of the input sequence is then used to control the plant, until a period of time passes when the optimization problem is repeated and an entirely new input policy is computed. The fact that the control policy is revised online is both one of its greatest strengths and one of its greatest perils; the mismatch between the open-loop policies computed and closed-loop behavior of the controller can cause serious problems if the controller is not designed with care.

In order to predict the behavior of a system, MPC requires information about the system's past. In the academic literature, this information is called the system state. Kalman (1963, p 154), an early (in the English-speaking literature, at least) proponent of the state space approach to systems and control, defined the system's state thus:

Intuitively speaking, the state is the minimal amount of information about the past history of the system which suffices to predict the effect of the past upon the future.

Therefore state estimation is a vital part of any implementation of MPC. The most famous state estimator is the Kalman filter. Because it combined a state-space formalism that is significantly easier to understand than the frequency domain formalism used in earlier results like Wiener filtering with a simple recursive update formula, according to Kailath (1974, p 150):

Kalman developed [64], [68], [69], a somewhat more restricted algorithm

than Swerling's, but it was one that seemed particularly matched to the dynamical state estimation problems that were brought forward by the advent of the space age. Groups at the NASA Ames Laboratory [71], [72], and at the M.I.T. Draper Laboratories [70], [83], took up Kalman's ideas and developed them into programs that were successfully used in many space applications [139], [145], [138].

However, the Kalman filter applies only to linear systems. When confronted with a nonlinear system, the obvious solution is to linearize it, then apply the Kalman filtering equations. When linearization is done by taking partial derivatives, the extended Kalman filter (EKF) is the result. The EKF has the advantage of having a relatively easy, recursive implementation. However, according to Julier and Uhlmann (2004):

The extended Kalman filter (EKF) is probably the most widely used estimation algorithm for nonlinear systems. However, more than 35 years of experience in the estimation community has shown that is difficult to implement, difficult to tune, and only reliable for systems that are almost linear on the time scale of the updates.

Since the publication of that paper, the advent of algorithmic differentiation (AD) has obviated the need for the "...many pages of dense algebra that must be converted to code..." in order to calculate Jacobian matrices. However, their point still stands that

Linearized transformations are only reliable if the error propagation can be well approximated by a linear function. If this condition does not hold, the linearized approximation can be extremely poor. At best, this undermines the performance of the filter. At worst, it causes its estimates to diverge altogether. However, determining the validity of this assumption is extremely difficult because it depends on the transformation, the current state estimate, and the magnitude of the covariance. This problem is well documented in many applications such as the estimation of ballistic parameters of missiles [1], [9][12] and computer vision [13].

The solution they propose is to use *statistical* linearization, rather than pointwise linearization using derivatives. The resulting filter is the unscented Kalman filter. This transformation has the benefit of not having to evaluate derivatives directly, but has to evaluate the state evolution and measurement equations at certain special points, termed “sigma points”. Because derivatives do not need to be evaluated, the resulting filter equations can be applied to nonsmooth or even discontinuous systems (though how much we should trust the results of those equations is another question).

Moving horizon estimation (MHE) offers a different approach to nonlinear state estimation. Like in MPC, a sum of stage costs is optimized over a finite horizon, and an estimate of the state trajectory over that horizon is returned. This form of objective can be justified by its relationship to maximum likelihood estimation; given a certain prior distribution of the system state and distribution for the disturbances

affecting the system, the maximum likelihood estimate of the system trajectory is given by solving an optimal control problem minimizing a sum of stage costs (Bryson and Frazier, 1963; Cox, 1964).

Although a great variety of stability and robustness results exist for MPC, the corresponding literature for (nonlinear) MHE has lagged behind. This state of affairs is problematic for any practical implementation of MPC. Although simulations of MHE show that it is a good estimator for many nonlinear systems, without an understanding of why and under what conditions it is robustly stabilizing, we have no idea of when it is likely to fall short and how (and whether) these shortcomings can be addressed. In analysis of MHE, it is useful to consider an idealized state estimation method that is not limited to measurements that fall within a certain horizon, but rather uses all measurements generated since it was brought online. Termed full information estimation (FIE), it represents a limiting case for the sort of estimation that can be accomplished by MHE—that in which no data is excluded.

1.1 Outline

The primary focus of this dissertation is the robust stability of FIE and MHE. However, in order to develop the tools necessary to analyze this problem, several other topics must be covered first. The biggest tool is a new type of Lyapunov-like function, termed a Q function, that serves as an analysis tool that permits us to conclude robust stability of FIE and, in the exponential case, MHE so long as a sufficiently long horizon is used. However, it turns out that Q functions lend

themselves readily to analysis of the stability of MPC given a sufficiently long control horizon. Because regulation to the origin using MPC is a problem that is much more familiar to many readers than estimation, we introduce Q functions in that context in Chapter 3 instead. Chapter 4 uses the problem of tracking a time varying trajectory of inputs and outputs, rather than to a time-invariant steady state, as a bridge between regulation and estimation. In order to analyze that problem, characterization of nonlinear detectability using a storage function is necessary. Such a storage function, termed an i-IOSS Lyapunov function, is introduced in Chapter 2. Finally, all these tools are combined in analysis of FIE and MHE in Chapter 5.

Chapter 2 Nonlinear detectability and i-IOSS

In order for FIE and MHE to be well-posed problems, we need some notion of nonlinear detectability. For linear systems, detectability can be characterized in terms of the Kalman decomposition or a matrix rank condition, and it is known that time-invariant linear systems admit a robustly stable estimator if and only if they are detectable. A detectability assumption popular in optimization-based state estimation literature is that of incremental input/output-to-state stability (i-IOSS). Introduced by Sontag and Wang (1997), they showed that a nonlinear system admits a robustly stable full-order state estimator, i.e., a dynamical system evolving in the same state-space as the system stabilized by output injection, only if it is i-IOSS. While considering such Luenberger observers is sufficient to characterize detectability for linear systems, methods like FIE and MHE do not fall in this

framework. In Chapter 2, a more general notion of state estimator in terms of maps from inputs and outputs to states is proposed. This definition can easily be applied to FIE, and MHE can be analyzed in similar terms with modest effort. It is shown in Proposition 2.4 that a system admits a robustly stable state estimator only if it is i-IOSS.

Properties similar to i-IOSS, such as input-to-state stability (ISS), have been characterized in the literature by storage functions. For example, Jiang and Wang (2001) demonstrate that, under mild regularity conditions, a discrete-time system is ISS if and only if it admits such a storage function, termed ISS Lyapunov function. However, no such result existed in the literature for i-IOSS before (Allan et al., 2021) was submitted. Here, that result is stated, and a full converse theorem, Theorem 2.10, is provided for the particular case of exponential i-IOSS. Furthermore, Theorem 2.14 demonstrates that, for a Lipschitz continuous system, the resulting i-IOSS Lyapunov function is Lipschitz. Finally, a useful result, Corollary 2.15 is given for changing the supply rate of exponential i-IOSS Lyapunov functions.

Chapter 3 MPC and setpoint regulation

The primary aim of this chapter is to introduce the major new analysis tool, the Q function, in a context that's likely to be familiar to many researchers in the field—regulatory control of a system to a steady-state. The dominant paradigm for guaranteeing the closed-loop stability of MPC is that formalized by Mayne et al. (2000), the use of stabilizing terminal conditions, whether that be a terminal equality constraint, such as that used in (Keerthi and Gilbert, 1988; Rawlings and

Muske, 1993), or a terminal region, such as that used in (De Nicolao et al., 1998; Chen and Allgöwer, 1998). A parallel literature about MPC without stabilizing terminal conditions has grown up, such as that featured in (Grimm et al., 2005; Tuna et al., 2006; Grune and Rantzer, 2008) and formalized in (Grüne and Pannek, 2016). The literature for MPC without stabilizing terminal conditions, however, is somewhat unapproachable to the uninitiated. Grüne and Pannek (2016) derive tight bounds for how long a horizon must be used in order to produce a stabilizing controller, but the formalization in terms of an abstract optimization problem fails to give much intuition about what is occurring with the optimal cost as the horizon is extended.

Here, we begin with a general stabilizability assumption and an infinite horizon optimal control problem. Through a variation on the arguments of Keerthi and Gilbert (1985), it can be shown that this infinite horizon problem has a solution. Furthermore, it can be shown that the optimal cost function is continuous, and therefore the system is robustly stabilizable. As stated in Theorem 3.16, the optimal cost function can serve as a control Lyapunov function, producing a converse theorem more general than that in (Kellett and Teel, 2004a) in that it does not have certain difficulties in application to systems with nonconvex sets of inputs (such as those which describe MPC with discrete actuators).

The form that the proof of continuity of the infinite horizon optimal cost function takes is rather important, however. Continuity is proven by showing that the finite-horizon optimal cost functions, which can be shown to be continuous through results in optimization, converge uniformly to the infinite horizon optimal cost function. In order to guarantee uniformity of convergence, the first Q

function in this work is constructed. In addition to allowing the infinite-horizon cost function to inherit continuity from the finite-horizon cost functions, it allows the finite-horizon cost functions to inherit approximate satisfaction of Bellman's equation from the infinite-horizon cost function.

Unfortunately, without further assumptions on the stage cost, the most that can be guaranteed is semiglobal practical stability (Theorem 3.21). This limitation is not a fault in analysis, but rather a feature of MPC schemes without stabilizing terminal conditions, as two examples show. If the finite horizon cost functions converge *exponentially* to the infinite horizon cost function, however, global stability of MPC with a finite horizon can be shown (Theorem 3.26). Although robustness is not explicitly analyzed here, it is a natural consequence of the continuity of the optimal cost functions.

This chapter ends with a comparison between MPC with a terminal equality constraint, MPC with a terminal region, and MPC without stabilizing terminal conditions. If MPC is used to regulate to a steady state without any active constraints, in my judgment MPC with a terminal region is superior to the other two methods. In order for a terminal equality constraint to satisfy constraint qualifications needed to guarantee good behavior of gradient-based solvers, the linearized system must be controllable. That is precisely the case in which there are good methods for constructing a terminal region. Similarly, Theorem 3.26 is satisfied with a quadratic cost function only if the linearized system is stabilizable. In attempting to avoid the local design problem of a terminal cost and terminal region, one ends up with the global design problem of designing a stage cost.

Chapter 4 Output tracking with MPC

In an attempt to bridge the gap between the problem of regulation to a setpoint familiar to all MPC researchers and using MHE for state estimation, we follow the prescient suggestion by Rawlings and Ji (2012):

Another issue that may hinder many researchers from making quick exploratory forays into state estimation is the inherent problem complexity. The simplest control problem is regulation to the origin, and the study of that simplest problem is still far from complete. There is no serious state estimation problem without an entire sequence of measurements. And the appearance of a sequence of measurements significantly complicates the notation and basic state estimation problem statements. It's more difficult for newcomers to get a foothold in this subject. It would be as if every control problem were stated as a tracking problem with a time-varying setpoint sequence. If that were how one had to get started in control problems, then state estimation would be about the same complexity and would probably compete better for new researchers' attention.

The situation is even worse than they describe, however, because there is the additional complication in that the stage cost used is no longer positive definite. In the case of linear problems, semidefinite stage costs were considered since the introduction of the linear quadratic regulator by Kalman (1960b), in which he required the assumption of uniform observability. In the case of nonlinear

problems, despite attempts like Property O introduced by Keerthi and Gilbert (1988), characterization of detectability through a storage function by Grimm et al. (2005), and the IOSS Lyapunov function used by Rawlings et al. (2017, Sec. 2.4.4), examples of MPC with semidefinite cost functions remain few and far between. Here, an i-IOSS Lyapunov function from Chapter 2 and a Q function from Chapter 3 are combined to first prove exponential convergence of finite horizon control problems to an infinite horizon control problem (Proposition 4.6) and then prove, as a result MPC can stabilize an arbitrary feasible output trajectory, i.e., an output trajectory generated by some real sequence of states and inputs, with a finite horizon (Theorem 4.7).

Chapter 5 Robust stability of FIE and MHE

Finally, we come to the central point of this work: analysis of the robust stability of FIE and MHE. As mentioned above, analysis of the robust stability of FIE and MHE has, for a long time, fallen short that of MPC. Because of the fundamental differences between MPC, where the closed-loop behavior of a system under a certain control law is analyzed, and FIE, where a sequence of open loop estimation problems is analyzed, even results for nominal stability of FIE have fallen short of the rigor that exists for MPC. Q functions were introduced in (Allan and Rawlings, 2019b) in order to remedy this lack of rigor. Two other novel features, the characterization of detectability using an i-IOSS Lyapunov function and the addition of a *stabilizability* assumption, distinguish that work from the previous literature. The result there is given in Section 5.2.1.

An example is then given in Section 5.2.2 in order to illustrate the sort of behav-

ior that we can and cannot expect FIE to have in the case of persistent disturbances. In particular, we cannot guarantee the good behavior of smoothed estimates of states under the assumptions common in the estimation literature. However, we can guarantee that the filtered estimates, i.e., estimates of the current state, remain bounded, and a proof of the robust stability of FIE is provided (Theorem 5.18). Unfortunately, the assumptions necessary for this result are somewhat unintuitive. A tighter characterization is possible for systems that are exponentially detectable and stabilizable. In particular, it is possible to use the robust stability of FIE (Theorem 5.28) to conclude the robust stability of MHE (Theorem 5.30). This work is submitted for publication in (Allan and Rawlings, 2021), and abbreviated proofs are given here. This chapter is concluded with some thoughts on why there is no obvious way to extend the robust stability of asymptotic FIE to asymptotic MHE.

CHAPTER **2**

NONLINEAR DETECTABILITY AND INCREMENTAL INPUT/OUTPUT TO STATE STABILITY

In order to treat the problem of nonlinear state estimation, we should first investigate the conditions under which it is well-posed. For linear systems, detectability is a necessary and sufficient condition for the existence of a robustly stable state estimator. For nonlinear systems, the most useful definition of detectability turns out to be incremental input/output-to-state stability (i-IOSS). Because this property is somewhat abstract, we motivate it by a comparison to extant results for linear systems, and demonstrate it is a necessary condition for the existence of a robustly stable state estimator. Then, it is useful to introduce a characterization of i-IOSS using a storage function, termed an i-IOSS Lyapunov function. A proof that a

system is i-IOSS if and only if it admits an i-IOSS Lyapunov function is contained in (Allan et al., 2021) and it is not repeated here, but a proof for the particular case of exponential i-IOSS is proved because it is significantly streamlined and provides an i-IOSS Lyapunov function with a particular form.

2.1 Literature review

The truly major accomplishment made by Kalman (1960a) in the introduction of the filter that bears his name was *not* the recursive form of the filtering equations. Swerling (1959) published a solution of a special case in which no process noise existed. A larger contribution was articulating the filtering problem using a state space formalism (Kailath, 1974, p 151) that allowed these recursive equations to be stated in a particularly simple way. The major accomplishment, though, was the introduction of the notions of controllability and observability to the filtering literature:

As Kalman has often stressed [68] the major contribution of his work is not perhaps the actual filter algorithm, elegant and useful as it no doubt is, but the proof that under certain technical conditions called “controllability” and “observability,” the optimum filter is “stable” or “robust” in the sense that the effects of initial errors and round-off and other computational errors will die out asymptotically. —Kailath (1974, p 152)

However, it quickly became clear that while observability is a sufficient condition

for estimator stability, it is not a necessary one. The notion of detectability was introduced by Wonham (1968a) and Anderson and Moore (1981) extended it to time-varying linear systems.

A literature about the definition of observability for nonlinear systems quickly grew. Brockett (1972) extended the notion of observability to a class of bilinear systems characterized by Lie algebras. Sussmann (1976) characterized observability for continuous-time analytic systems in the context of realization theory. Hermann and Krener (1977) characterized observability by whether or not states could be distinguished by local measurements and whether or not adjacent states could be distinguished. They called the property that adjacent states can be distinguished by local measurements “local weak observability”, showed that, if the matrix of Lie derivatives is full rank, then the system is locally weakly observable, and also showed that if a system is locally weakly observable, the rank condition is fulfilled almost everywhere. Aeyels (1981) showed that almost all smooth nonlinear sampled systems are observable so long as $2n + 1$ measurements are used, in which n is the dimension of the state space. Sontag (1984) gave a sufficient condition for local observability of analytic nonlinear sampled systems in terms of a rank condition for the matrix of Lie derivatives, and showed that all such systems that are locally observable satisfy the rank condition almost everywhere. Interestingly, in order to obtain certain results about local observability to hold he requires the system be locally controllable as well. One feature that these characterizations of nonlinear observability share is that, in general, a certain input sequence may be required to distinguish two states.

The introduction of input-to-state stability (ISS) by Sontag (1989) produced a paradigm shift in terms of how robust stability properties are considered. In an investigation of related properties, Sontag and Wang (1997) introduced i-IOSS as a notion of nonlinear detectability. They motivated it by demonstrating that any system that admits a full-order state observer, i.e., a parallel dynamical system evolving in the same state space that is stabilized by output injection, must necessarily be i-IOSS. Rawlings and Mayne (2009, Ch. 4) introduced it to the optimization-based state estimation literature, where it has been used as a standard detectability assumption ever since.

Characterization of ISS using a storage function, termed an ISS Lyapunov function, alongside a corresponding converse theorem, was provided by Sontag and Wang (1995) for continuous time systems and Jiang and Wang (2001) for discrete-time systems. ISS Lyapunov functions for discrete-time systems with discontinuous state evolution maps were characterized by Grüne and Kellett (2014). A Lyapunov characterization and corresponding converse theorem were provided for output-to-state stability (OSS) by Sontag and Wang (1997) and a continuous-time converse theorem for (non-incremental) input/output-to-state stability (IOSS) was provided by Krichman et al. (2001). A discrete-time counterpart was provided by Cai and Teel (2008), and Allan and Rawlings (2018b) extended the result from IOSS systems defined on all of \mathbb{R}^n to those defined on closed positive invariant subsets of \mathbb{R}^n . A Lyapunov characterization to incremental ISS was given by Angeli (2002), and a converse theorem was provided for compact input domains. Incremental ISS and a converse theorem for systems with compact input domains were provided for

discrete-time systems by Tran et al. (2016). The notions of incremental passivity (Pavlov and Marconi, 2008) and incremental dissipativity (Santoso et al., 2012) are also related to i-IOSS.

2.2 *i-IOSS and robustly stable state estimators*

By the definition introduced by Wonham (1968a), a linear time-invariant system

$$x^+ = Ax + Bu \quad y = Cx$$

is detectable if and only if there exists some matrix L such that $A - LC$ is Schur¹ stable, i.e., having all eigenvalues strictly inside the unit circle. Other equivalent definitions exist in terms of matrix rank conditions and observable subspaces. Suppose we are given an initial estimate of x , $\hat{x}(0)$, and are given a forecast u_f of the input sequence u such that $u = u_f + w$ and a measurement y_m of the output y such that $y_m = y + v$. Such an L can be used to construct a robustly stable observer

$$\hat{x}^+ = A\hat{x} + Bu_f + L(y_m - C\hat{x}).$$

Because the system is linear, we can combine this evolution equation with that of the true state to obtain the evolution equation of the estimation error $e(k) := x(k) - \hat{x}(k)$:

$$e^+ = (A - LC)e + Bw - Lv.$$

This equation can be applied recursively to obtain the estimate error in terms of the initial error and a convolution sum of disturbances w and v :

¹Wonham (1968a) worked in continuous-time, so he required $A - LC$ to be Hurwitz stable.

$$e(k) = (A - LC)^k e(0) + \sum_{j=0}^{k-1} (A - LC)^{k-j-1} B w(j) - \sum_{j=0}^{k-1} (A - LC)^{k-j-1} L v(j).$$

By using the well-known fact that there exists $c > 0$ and $\lambda \in (0, 1)$ such that $|(A - LC)^k| \leq c\lambda^k$ because $A - LC$ is Schur stable, we can obtain

$$|e(k)| \leq c\lambda^k |e(0)| + \sum_{j=0}^{k-1} c\lambda^{k-j-1} |B| |w(j)| + \sum_{j=0}^{k-1} c\lambda^{k-j-1} |L| |v(j)|. \quad (2.1)$$

which gives us a complete understanding of the error dynamics of the system. The effects of initial error in the state estimate exponentially decay, as do the effects of past forecasting errors $w(j)$ and output noise $v(j)$. Furthermore, for all bounded disturbance sequences the estimation error is bounded. We can go one step further with this error expression and derive the bound

$$|e(k)| \leq c\lambda^k |e(0)| + \frac{c|B|}{1-\lambda} \|\mathbf{w}\|_{0:k-1} + \frac{c|L|}{1-\lambda} \|\mathbf{v}\|_{0:k-1},$$

in which $\|\mathbf{d}\|_{0:k-1} := \max_{j \in \mathbb{I}_{0:k}} |d(j)|$, but this expression is less descriptive than the previous one.

Now, consider the nonlinear system

$$x^+ = f(x, u) \quad y = h(x) \quad (2.2)$$

in which $x \in \mathbb{X} \subseteq \mathbb{R}^n$ is the system state, x^+ is the successor state, $u \in \mathbb{U} \subset \mathbb{R}^m$ is the system input, and $y \in \mathbb{Y}$ is the system output. We denote the k^{th} element of a trajectory starting from x evolving with input sequence \mathbf{u} as $\phi(k; x, \mathbf{u})$. We denote the k^{th} output of this trajectory by $y(k; x, \mathbf{u})$ and denote the entire output sequence by $\mathbf{y}(x, \mathbf{u})$.

When Sontag and Wang (1997) introduced the concept of incremental input/output-to-state stability (i-IOSS) as the definition of nonlinear detectability, they were motivated by this sort of “full-order observer”, i.e., a parallel dynamical system that evolves in the same state space as the system state and is stabilized by output injection. While the design of this type of observer is an active area of research (see (Rajamani, 2017a,b) for some recent results), there are many other types of nonlinear state estimation. The extended Kalman filter (EKF) has a state that consists of a state estimate and an approximate covariance matrix. MHE has a complicated internal state that varies with implementation, while FIE admits no convenient (i.e., finite dimensional) state-space representation. We define a general state estimator in terms of input/output maps.

Definition 2.1 (State estimator). A state estimator is a sequence of functions $\Psi_k : \mathbb{X} \times \mathbb{U}^k \times \mathbb{Y}^k \rightarrow \mathbb{X}$ that takes an initial state estimate, \bar{x} , a sequence of forecast inputs $\mathbf{u}_f \in \mathbb{U}^\infty$, and a sequence of measured outputs $\mathbf{y}_m \in \mathbb{Y}^\infty$ to produce a state estimate

$$\hat{x}(k) := \Psi_k(\bar{x}, \mathbf{u}_f(0 : k - 1), \mathbf{y}_m(0 : k - 1))$$

for all $k \in \mathbb{I}_{\geq 0}$.

Note that we frequently abuse notation and write $\Psi_k(\bar{x}, \mathbf{u}_f, \mathbf{y}_m)$ for $\Psi_k(\bar{x}, \mathbf{u}_f(0 : k - 1), \mathbf{y}_m(0 : k - 1))$. While a function $\Psi_k : \mathbb{X} \times \mathbb{U}^\infty \times \mathbb{Y}^\infty \rightarrow \mathbb{X}$ can be distinguished from one $\Psi_k : \mathbb{X} \times \mathbb{U}^k \times \mathbb{Y}^k \rightarrow \mathbb{X}$, the additional verbosity adds little clarity.

This definition is natural for FIE, and other methods like full-order observers, MHE, and the EKF can be analyzed in this fashion. Sontag and Wang (1997, Remark

25) remarked that an input/output mapping would be a more general definition of state estimator, but declined to do so because full-order observers were more common in the literature at the time.

We want an expression like (2.1) to characterize robust estimator stability for nonlinear systems. However, many nonlinear systems have subexponential decay rates. Towards this end, we can substitute a general \mathcal{KL} function $\beta(s, t)$ for the particular \mathcal{KL} function $cs\lambda^t$. However, many \mathcal{KL} functions decay slowly enough such that

$$\lim_{k \rightarrow \infty} \sum_{j=0}^{k-1} \beta(|d(j)|, k-j-1) = \infty$$

or bounded \mathbf{d} , e.g., $\beta(s, t) = s/(t+1)$. A solution is to substitute maximization for summation. Let $a \oplus b := \max(a, b)$ for $a, b \in \mathbb{R}$. Note that \mathcal{K} functions distribute across this operation, i.e., for $\alpha \in \mathcal{K}$, we have that $\alpha(a \oplus b) = \alpha(a) \oplus \alpha(b)$.

Definition 2.2 (Robust global asymptotic stability (RGAS)). An estimator (Ψ_k) is RGAS if there exist $\beta_x, \beta_w, \beta_v \in \mathcal{KL}$ such that

$$\begin{aligned} |x(k) - \hat{x}(k)| \leq & \beta_x(|x(0) - \bar{x}|, k) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_w(|w(j)|, k-j-1) \\ & \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_v(|v(j)|, k-j-1), \end{aligned}$$

in which

$$\begin{aligned} x(j) &= \phi(j; x(0), \mathbf{u}) & \hat{x}(j) &= \Psi_j(\bar{x}, \mathbf{u}_f(0:j-1), \mathbf{y}_m(0:j-1)) \\ y(j) &= h(x(j)) & \mathbf{v} &= \mathbf{y}_m - \mathbf{y} \\ & & \mathbf{w} &= \mathbf{u} - \mathbf{u}_f, \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$, $x(0), \bar{x} \in \mathbb{X}$, $\mathbf{u}, \mathbf{u}_f \in \mathbb{U}^\infty$, and $\mathbf{y}, \mathbf{y}_m \in \mathbb{Y}^\infty$.

With this definition of RGAS in mind, i-IOSS is now defined. In order to streamline the presentation of incremental stability properties, let

$$\begin{aligned} x_1(j) &:= \phi(j; x_1(0), \mathbf{u}_1) & y_1(j) &:= h(x_1(j)) \\ x_2(j) &:= \phi(j; x_2(0), \mathbf{u}_2) & y_2(j) &:= h(x_2(j)) \\ \Delta x(j) &:= x_1(j) - x_2(j) & \Delta y(j) &:= y_1(j) - y_2(j) & \Delta u(j) &:= u_1(j) - u_2(j) \end{aligned}$$

for some $x_1(0), x_2(0) \in \mathbb{X}$, $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^\infty$, and $j \in \mathbb{I}_{\geq 0}$.

Definition 2.3 (Incremental input/output-to-state stability (i-IOSS)). A system (2.2) is i-IOSS if there exist $\beta_x, \beta_u, \beta_y \in \mathcal{KL}$ such that

$$\begin{aligned} |\Delta x(k)| &\leq \beta_x(|\Delta x(0)|, k) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_u(|\Delta u(j)|, k - j - 1) \\ &\quad \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_y(|\Delta y(j)|, k - j - 1), \end{aligned}$$

for all $x_1(0), x_2(0) \in \mathbb{X}$, $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^\infty$, and $k \in \mathbb{I}_{\geq 0}$.

This definition of i-IOSS is not the one common in the literature. Sontag and Wang (1997, Def. 22) defined i-IOSS in terms of a \mathcal{KL} function $\beta(\cdot)$ and asymptotic gain functions $\gamma_u, \gamma_y \in \mathcal{K}$ with the bound

$$|\Delta x(k)| \leq \beta(|\Delta x(0)|, k) \oplus \gamma_u(\|\Delta \mathbf{u}\|_{0:k-1}) \oplus \gamma_y(\|\Delta \mathbf{y}\|_{0:k-1}) \quad (2.3)$$

and this definition is predominant in the literature. Similarly, they define robust stability in terms of a \mathcal{KL} function $\beta(\cdot)$ and asymptotic gain functions $\gamma_w, \gamma_v \in \mathcal{K}$ such that

$$|x(k) - \hat{x}(k)| \leq \beta(|x(0) - \bar{x}|, k) \oplus \gamma_w(\|\mathbf{w}\|_{0:k-1}) \oplus \gamma_v(\|\mathbf{v}\|_{0:k-1}). \quad (2.4)$$

This definition of robust stability is adequate for the full-order state observers they were considering, although proving that an observer converges when disturbances

converge is not straightforward. However, for the broader class of estimators given by Definition 2.1, the implication that the estimator converges if disturbances converge is lost. For full-order state observers, (2.4) can be applied repeatedly to move the effect of past disturbances from the asymptotic gain terms, which persist, to the \mathcal{KL} function term, which decays. However, the initial state is given a special role in Definition 2.1, and frequently we have that

$$\Psi_2(\bar{x}, \mathbf{y}_m(0 : 1), \mathbf{u}_f(0 : 1)) \neq \Psi_1(\Psi_1(\bar{x}, y_m(0), u_f(0)), y_m(1), u_f(1))$$

in estimators such as FIE and MHE. Several recent works (Ji et al., 2016; Hu, 2017) are able to show that a variation of FIE satisfies (2.4) but unable to show convergence in the case of convergent disturbances. That is why Definition 2.2 is used instead.

It is shown in the appendix of (Allan et al., 2020) that i-IOSS as defined using (2.3) implies Definition 2.3, which explicitly reveals the sort of convergence properties enjoyed by i-IOSS systems. However, that proof is tedious, and requires pushing the techniques developed by Krichman et al. (2001), Cai and Teel (2008), Grüne and Kellett (2014), and Allan and Rawlings (2018b) to their limits. A much more direct argument shows that any system that admits an RGAS estimator as defined by Definition 2.2 must necessarily be i-IOSS as defined by Definition 2.3.

Proposition 2.4 ((Allan et al., 2021, Prop. 2.5)). *Any system (2.2) that admits an RGAS estimator (Ψ_k) must be i-IOSS.*

The key to this proof is noting that if a robustly stable state estimator is fed a feasible initial condition, sequence of inputs, and corresponding outputs, it returns the corresponding state. In other words,

$$\phi(k; x(0), \mathbf{u}) = \Psi_k(x(0), \mathbf{u}, \mathbf{y}(x(0), \mathbf{u}))$$

by application of Definition 2.2. We can then consider the input sequences in two ways. The first is as a true system trajectory, the second is as a disturbed system trajectory originating from some different initial state, input sequences, and disturbance sequences.

Proof. Let $x_1(0), x_2(0) \in \mathbb{X}$ and $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^\infty$ be arbitrary. By Definition 2.2 we have that

$$\begin{aligned} |x_1(k) - \Psi_k(x_1(0), \mathbf{u}_1, \mathbf{y}_1)| &\leq \beta_x(|x_1(0) - x_1(0)|, k) \oplus \beta_u(|u_1(j) - u_1(j)|, k - j - 1) \\ &\quad \oplus \beta_y(|y_1(j) - y_1(j)|, k - j - 1) \\ &= 0 \end{aligned}$$

and therefore $\Psi_k(x_1(0), \mathbf{u}_1, \mathbf{y}_1) = x_1(k)$. Similarly, we have that $\Psi_k(x_2(0), \mathbf{u}_2, \mathbf{y}_2) = x_2(k)$. Next, consider what would happen if we attempted to estimate $x_2(k)$ using $\bar{x} = x_1(0)$, $\mathbf{u}_f = \mathbf{u}_1$, and $\mathbf{y}_m = \mathbf{y}_1$. Those choices of forecast and measurements produce $\mathbf{w} = \mathbf{u}_2 - \mathbf{u}_1$ and $\mathbf{v} = \mathbf{y}_1 - \mathbf{y}_2$. Application of Definition 2.2 produces

$$\begin{aligned} |x_2(k) - \Psi_k(x_1(0), \mathbf{u}_1, \mathbf{y}_1)| &\leq \beta_x(|x_2(0) - \bar{x}|, k) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_w(|w(j)|, k - j - 1) \\ &\quad \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_v(|v(j)|, k - j - 1) \end{aligned}$$

and therefore

$$\begin{aligned} |x_2(k) - x_1(k)| &\leq \beta_x(|x_2(0) - x_1(0)|, k) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_w(|\Delta u(j)|, k - j - 1) \\ &\quad \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_v(|\Delta y(j)|, k - j - 1). \end{aligned}$$

Because $x_1(0), x_2(0) \in \mathbb{X}$ and $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^\infty$ were arbitrary, the system is i-IOSS by Definition 2.3. ■

Although numerous ISS-like properties have received characterization with storage functions, until recently there was no such characterization of i-IOSS in the literature. To my knowledge, an i-IOSS Lyapunov function was first defined in (Allan and Rawlings, 2019b).

Definition 2.5 (i-IOSS Lyapunov function). A function $\lambda : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0}$ is an i-IOSS Lyapunov function for a system (2.2) if there exist $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ and $\gamma_u, \gamma_y \in \mathcal{K}$ such that

$$\begin{aligned} \alpha_1(|\Delta x|) &\leq \lambda(x_1, x_2) \leq \alpha_2(|\Delta x|) \\ \lambda(f(x_1, u_1), f(x_2, u_2)) &\leq \lambda(x_1, x_2) - \alpha_3(|\Delta x|) + \gamma_u(|\Delta u|) + \gamma_y(|\Delta y|) \end{aligned}$$

for all $x_1, x_2 \in \mathbb{X}$ and $u_1, u_2 \in \mathbb{U}$.

Theorem 2.6 ((Allan et al., 2021, Thm. 1)). *A system (2.2) is i-IOSS if and only if it admits an i-IOSS Lyapunov function.*

The proof of this theorem is rather technical, and is not duplicated here. However, a streamlined proof is available for the special case of exponential i-IOSS. Exponential i-IOSS has recently been used by Knüfer and Müller (2018) to prove the robust stability of MHE with a certain cost function.

2.3 Exponential i-IOSS Converse Theorem

For asymptotic i-IOSS, it is difficult to produce a converse theorem from (2.3) directly (though that is accomplished in the appendix of (Allan et al., 2020)), and I know of no method to show continuity for the resulting i-IOSS Lyapunov function.

However, a much more direct proof from the asymptotic gain definition is possible for exponential i-IOSS, so we illustrate it here.

Definition 2.7 (Exponential i-IOSS). A system (2.2) is said to be exponentially i-IOSS if there exist $C_x, C_u, C_y > 0$ and $\lambda \in (0, 1)$ such that

$$|\Delta x(k)| \leq C_x |\Delta x(0)| \lambda^k + C_u \|\Delta \mathbf{u}\|_{0:k-1} + C_y \|\Delta \mathbf{y}\|_{0:k-1} \quad (2.5)$$

for all $x_1, x_2 \in \mathbb{X}$, all $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^\infty$, and all $k \in \mathbb{I}_{\geq 0}$.

Definition 2.8 (Exponential i-IOSS Lyapunov function). A function $V : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is called an exponential i-IOSS Lyapunov function if there exist $c_1, c_2, c_3, c_u, c_y, \sigma > 0$ such that

$$c_1 |\Delta x|^\sigma \leq \lambda(x_1, x_2) \leq c_2 |\Delta x|^\sigma \quad (2.6)$$

$$\lambda(f(x_1, u_1), f(x_2, u_2)) \leq \lambda(x_1, x_2) - c_3 |\Delta x|^\sigma + c_u |\Delta u|^\sigma + c_y |\Delta y|^\sigma, \quad (2.7)$$

for all $x_1, x_2 \in \mathbb{X}$ and $u_1, u_2 \in \mathbb{U}$.

We define

$$\lambda(x_1, x_2) := \sup_{k, \mathbf{u}_1, \mathbf{u}_2} \lambda^{-k/2} [|\Delta x(k)| - C_u \|\Delta \mathbf{u}\|_{0:k-1} - C_y \|\Delta \mathbf{y}\|_{0:k-1}] \quad (2.8)$$

as our exponential i-IOSS Lyapunov function candidate. By substituting (2.5) into this equation, we obtain

$$\lambda(x_1, x_2) \leq \sup_{k, \mathbf{u}_1, \mathbf{u}_2} \lambda^{-k/2} [C_x \lambda^k |\Delta x|] = C_x |\Delta x|, \quad (2.9)$$

so $\lambda(x_1, x_2)$ admits an upper bound linear in $|x_1 - x_2|$ for every $x_1, x_2 \in \mathbb{X}$. Furthermore, by evaluating the value function for $k = 0$, we obtain

$$\Lambda(x_1, x_2) \geq [|\Delta x| - C_u \|\Delta \mathbf{u}\|_{0:-1} - C_y \|\Delta \mathbf{y}\|_{0:-1}] = |\Delta x|, \quad (2.10)$$

so $\Lambda(\cdot)$ satisfies (2.6). In order to derive the cost-decrease condition, we require an upper bound on the “optimal” k in (2.8).² In order to more easily discuss feasible but suboptimal sequences, for some $\mathbf{u}_1^*, \mathbf{u}_2^* \in \mathbb{U}^\infty$, let

$$\begin{aligned} x_1^*(j) &:= \phi(j; x_1(0), \mathbf{u}_1^*) & y_1^*(j) &:= h(x_1^*(j)) \\ x_2^*(j) &:= \phi(j; x_2(0), \mathbf{u}_2^*) & y_2^*(j) &:= h(x_2^*(j)) \\ \Delta x^*(j) &:= x_1^*(j) - x_2^*(j) & \Delta y^*(j) &:= y_1^*(j) - y_2^*(j) & \Delta \mathbf{u}^*(j) &:= \mathbf{u}_1^*(j) - \mathbf{u}_2^*(j). \end{aligned}$$

Proposition 2.9. *There exists some \bar{k} , such that for every $x_1, x_2 \in \mathbb{X}$,*

$\varepsilon \in (0, |x_1 - x_2|/2]$, $k^ \in \mathbb{N}_{\geq 0}$, and $\mathbf{u}_1^*, \mathbf{u}_2^* \in \mathbb{U}^\infty$ that satisfy*

$$\Lambda(x_1, x_2) \leq \varepsilon + \lambda^{-k^*/2} [|\Delta x^*(k^*)| - C_u \|\Delta \mathbf{u}^*\|_{0:k^*-1} - C_y \|\Delta \mathbf{y}^*\|_{0:k^*-1}] \quad (2.11)$$

we have that $k^ \leq \bar{k}$.*

Proof. Suppose we have $x_1, x_2 \in \mathbb{X}$. If $x_1 = x_2$, then $(0, |x_1 - x_2|/2]$ is empty, and the proposition holds trivially. Now suppose that $x_1 \neq x_2$. We can substitute (2.5) into (2.11) to obtain

$$\Lambda(x_1, x_2) \leq \varepsilon + C_x |x_1 - x_2| \lambda^{k^*/2}.$$

From (2.10) and the assumption on ε , we obtain

$$\begin{aligned} |x_1 - x_2| &\leq |x_1 - x_2|/2 + C_x |x_1 - x_2| \lambda^{k^*/2} \\ |x_1 - x_2|/2 &\leq C_x |x_1 - x_2| \lambda^{k^*/2} \\ 1/2 &\leq C_x \lambda^{k^*/2}, \end{aligned}$$

²There is no guarantee that the infimum is attained by any particular k , \mathbf{u}_1 , and \mathbf{u}_2 , but all sequences that are sufficiently close to the infimum must obey this bound.

in which the last step follows because $x_1 \neq x_2$. We can then obtain

$$k^* \leq \lceil -\log_\lambda(2C_x) \rceil := \bar{k},$$

in which $\lceil \cdot \rceil$ denotes the ceiling function. ■

With this result, we can demonstrate a cost-decrease condition for $\Lambda(\cdot)$.

Theorem 2.10. *A system (2.2) is exponentially i-IOSS if and only if it admits an exponential i-IOSS Lyapunov function.*

Proof. The sufficiency of such a Lyapunov function can be shown by a straightforward specialization of existing arguments for asymptotic i-IOSS such as (Allan and Rawlings, 2019b, Prop. 5). Showing the necessity of such a function requires considerably more work.

Because of (2.9) and (2.10), we have that (2.6) holds with $\sigma = 1$, $c_1 = 1$, and $c_2 = C_x$. Next, we derive a cost decrease condition.

Let $x_1, x_2 \in \mathbb{X}$ and $u_1, u_2 \in \mathbb{U}$. Let $x_1^+ := f(x_1, u_1)$, $x_2^+ := f(x_2, u_2)$, $y_1 := h(x_1)$, and $y_2 := h(x_2)$. Suppose that $x_1^+ \neq x_2^+$ and fix $\varepsilon \in (0, |x_1 - x_2|/2]$. There exist $\tilde{k} \leq \bar{k}$, $\tilde{\mathbf{u}}_1, \tilde{\mathbf{u}}_2 \in \mathbb{U}^\infty$ such that

$$\begin{aligned} \Lambda(x_1^+, x_2^+) \leq & \varepsilon + \lambda^{-\tilde{k}/2} \left[\left| \phi(\tilde{k}; x_1^+, \tilde{\mathbf{u}}_1) - \phi(\tilde{k}; x_2^+, \tilde{\mathbf{u}}_2) \right| - C_u \|\mathbf{u}_1 - \mathbf{u}_2\|_{0:\tilde{k}-1} \right. \\ & \left. - C_y \max_{j \in \mathbb{I}_{0:\tilde{k}-1}} \left| h \circ \phi(j; x_1^+, \tilde{\mathbf{u}}_1) - h \circ \phi(j; x_2^+, \tilde{\mathbf{u}}_2) \right| \right] \end{aligned} \quad (2.12)$$

and by Proposition 2.9 we have that $k^* \leq \tilde{k}$. By considering trajectories beginning from x_1 and x_2 , we have that

$$\begin{aligned} \phi(\tilde{k}; x_1^+, \tilde{\mathbf{u}}_1) &= \phi(\tilde{k} + 1; x_1, u_1 \widehat{\tilde{\mathbf{u}}}_1) \\ \phi(\tilde{k}; x_2^+, \tilde{\mathbf{u}}_2) &= \phi(\tilde{k} + 1; x_2, u_2 \widehat{\tilde{\mathbf{u}}}_1) \end{aligned}$$

in which the operation $s \frown \mathbf{s}$ inserts s at the beginning of the sequence \mathbf{s} . For brevity,

let $\mathbf{u}_1^* := u_1 \frown \tilde{\mathbf{u}}_1$ and $\mathbf{u}_2^* := u_2 \frown \tilde{\mathbf{u}}_2$. Then we can rewrite (2.12) as

$$\begin{aligned} \Lambda(x_1^+, x_2^+) &\leq \varepsilon + \lambda^{-\tilde{k}/2} \left[|\Delta x^*(\tilde{k} + 1)| - C_u \|\Delta \mathbf{u}^*\|_{1:\tilde{k}} - C_y \|\Delta \mathbf{y}^*\|_{1:\tilde{k}} \right] \\ &\leq \varepsilon + \sqrt{\lambda} \left(\lambda^{-(\tilde{k}+1)/2} \left[|\Delta x^*(\tilde{k} + 1)| - C_u \|\Delta \mathbf{u}^*\|_{0:\tilde{k}} - C_y \|\Delta \mathbf{y}^*\|_{0:\tilde{k}} \right] \right) \\ &\quad + \lambda^{-\tilde{k}/2} (C_u |\Delta u^*(0)| + C_y |\Delta y^*(0)|), \end{aligned}$$

in which the last step follows because

$$\|\mathbf{s}\|_{1:n} = \|\mathbf{s}\|_{1:k} + |s(0)| - |s(0)| \geq (\|\mathbf{s}\|_{1:k} \oplus |s(0)|) - |s(0)| = \|\mathbf{s}\|_{0:k} - |s(0)|$$

for any sequence \mathbf{s} . We note that the quantity

$$\lambda^{-(k^*+1)/2} \left[|\tilde{x}_1(k^* + 1) - \tilde{x}_2(k^* + 1)| - C_u \|\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2\|_{0:k^*} - C_y \|\tilde{\mathbf{y}}_1 - \tilde{\mathbf{y}}_2\|_{0:k^*} \right]$$

is the value function (2.8) evaluated at (x_1, x_2) with particular k , \mathbf{u} , and \mathbf{d} . The supremum (2.8) must be greater than this feasible quantity, hence

$$\Lambda(x_1^+, x_2^+) \leq \varepsilon + \lambda^{1/2} \Lambda(x_1, x_2) + \lambda^{-\tilde{k}/2} (C_u |\Delta u^*(0)| + C_y |\Delta y^*(0)|).$$

Furthermore, $\tilde{k} \leq \bar{k}$ by Proposition 2.9, $u_i^*(0) = u_i$, and $y_i^*(0) = y_i$ for $i = 1, 2$.

Therefore, we have that

$$\begin{aligned} \Lambda(x_1^+, x_2^+) &\leq \varepsilon + \lambda^{1/2} \Lambda(x_1, x_2) + \lambda^{-\tilde{k}/2} (C_u |u_1 - u_2| + C_y |y_1 - y_2|) \\ &= \varepsilon + \Lambda(x_1, x_2) - (1 - \lambda^{1/2}) \Lambda(x_1, x_2) + c_u |u_1 - u_2| + c_y |y_1 - y_2| \\ &\leq \varepsilon + \Lambda(x_1, x_2) - (1 - \lambda^{1/2}) |x_1 - x_2| + c_u |u_1 - u_2| + c_y |y_1 - y_2|, \end{aligned}$$

in which $c_u := \lambda^{-\tilde{k}/2} C_u$, $c_y := \lambda^{-\tilde{k}/2} C_y$, and the final inequality follows from (2.10).

In this final expression, we have removed any quantity that depends on ε . Thus, we can take the limit as ε goes to zero to obtain

$$\Lambda(x_1^+, x_2^+) \leq \Lambda(x_1, x_2) - (1 - \lambda^{1/2}) |x_1 - x_2| + c_u |u_1 - u_2| + c_y |y_1 - y_2|,$$

which gives us the dissipation inequality (2.7) with $\sigma = 1$ and $c_3 = 1 - \lambda^{1/2}$.

Now suppose that $x_1^+ = x_2^+$. Then $\Lambda(x_1^+, x_2^+) = 0$, and the cost-decrease condition

$$\Lambda(x_1^+, x_2^+) \leq \Lambda(x_1, x_2) - (1 - \lambda^{1/2}) |x_1 - x_2| + c_u |u_1 - u_2| + c_y |y_1 - y_2|$$

follows from (2.10) and the non-negativity of the supply rate. Thus $\Lambda(\cdot)$ is an exponential i-IOSS Lyapunov function. ■

Although the existence of an i-IOSS Lyapunov function can be enough for analysis purposes, like in (Allan and Rawlings, 2019b), it is often desirable to have a continuous i-IOSS Lyapunov function. In the special case of exponential i-IOSS, a (globally) Lipschitz continuous i-IOSS Lyapunov function can be produced so long as the system (2.2) is itself Lipschitz continuous.

Assumption 2.11 (Lipschitz continuity). Both $f(\cdot)$ and $h(\cdot)$ are globally Lipschitz continuous, i.e., there exist $L_f, L_h > 0$ such that

$$\begin{aligned} |f(x_1, u_1) - f(x_2, u_2)| &\leq L_f |(x_1, u_1) - (x_2, u_2)| \\ |h(x_1) - h(x_2)| &\leq L_h |x_1 - x_2| \end{aligned}$$

for all $x_1, x_2 \in \mathbb{X}$ and all $u_1, u_2 \in \mathbb{U}$.

Remark 2.12. By applying (2.5) with $k = 1$, we obtain a bound

$$|f(x_1, u_1) - f(x_2, u_2)| \leq C_x \lambda |x_1 - x_2| + C_u |u_1 - u_2| + C_y |h(x_1) - h(x_2)|$$

so if $h(\cdot)$ is Lipschitz continuous, as occurs when a subset of states is measured, $f(\cdot)$ is necessarily Lipschitz continuous as well.

We require a minor result suggested, but not named, by Angeli (2009).

Proposition 2.13 (Quadrille inequality (suggested in (Angeli, 2009))). *For vectors $u_1, u_2, v_1, v_2 \in \mathbb{R}^n$, we have that*

$$\left| |u_1 - u_2| - |v_1 - v_2| \right| \leq |u_1 - v_1| + |u_2 - v_2|.$$

Theorem 2.14. *The function $\lambda(\cdot)$, as defined in (2.8), is globally Lipschitz continuous if Assumption 2.11 holds.*

Proof. Fix $x_1, x_2, z_1, z_2 \in \mathbb{X}$ and suppose that both $x_1 \neq x_2$ and $z_1 \neq z_2$. Let $\varepsilon \in (0, |x_1 - x_2|/2]$. There exist $k^* \in \mathbb{I}_{0:\bar{k}}$, $\mathbf{u}_1^*, \mathbf{u}_2^* \in \mathbb{U}^\infty$, and $\mathbf{d} \in \mathbb{D}^\infty$ such that

$$\lambda(x_1, x_2) \leq \varepsilon + \lambda^{-k^*/2} \left[|\Delta x^*(k^*)| - C_u \|\Delta \mathbf{u}^*\|_{0:k^*-1} - C_y \|\Delta \mathbf{y}^*\|_{0:k^*-1} \right] \quad (2.13)$$

and by Proposition 2.9 $k^* \leq \bar{k}$. For brevity, let

$$\begin{aligned} z_1^*(j) &:= \phi(j; z_1, \mathbf{u}_1^*) & \eta_1^*(j) &:= h(z_1^*(j)) \\ z_2^*(j) &:= \phi(j; z_2, \mathbf{u}_2^*) & \eta_2^*(j) &:= h(z_2^*(j)) \\ \Delta z^*(j) &:= z_1^*(j) - z_2^*(j) & \Delta \eta^*(j) &:= \eta_1^*(j) - \eta_2^*(j). \end{aligned}$$

Because this set of variables is feasible in the optimization that produces $\lambda(z_1, z_2)$, we have that

$$\lambda(z_1, z_2) \geq \lambda^{-k^*/2} \left[|\Delta z^*(k^*)| - C_u \|\Delta \mathbf{u}^*\|_{0:k^*-1} - C_y \|\Delta \boldsymbol{\eta}\|_{0:k^*-1} \right]. \quad (2.14)$$

We can subtract (2.14) from (2.13) to obtain

$$\begin{aligned} \Lambda(x_1, x_2) - \Lambda(z_1, z_2) &\leq \varepsilon + \lambda^{-k^*/2} \left[|x_1^*(k^*) - x_2^*(k^*)| - |z_1^*(k^*) - z_2^*(k^*)| \right] \\ &\quad - C_y \left[\|\mathbf{y}_1^* - \mathbf{y}_2^*\|_{0:k^*-1} - \|\boldsymbol{\eta}_1^* - \boldsymbol{\eta}_2^*\|_{0:k^*-1} \right], \end{aligned}$$

in which the terms involving \mathbf{u} have been canceled. We can then apply Proposition 2.13 to obtain

$$\begin{aligned} \Lambda(x_1, x_2) - \Lambda(z_1, z_2) &\leq \varepsilon + \lambda^{-k^*/2} \left[|x_1^*(k^*) - z_1^*(k^*)| + |x_2^*(k^*) - z_2^*(k^*)| \right] \\ &\quad + C_y \left[\|\mathbf{y}_1^* - \boldsymbol{\eta}_1^*\|_{0:k^*-1} + \|\mathbf{y}_2^* - \boldsymbol{\eta}_2^*\|_{0:k^*-1} \right]. \quad (2.15) \end{aligned}$$

By Assumption 2.11, we have that

$$|x_i^*(j+1) - z_i^*(j+1)| \leq L_f |x_i^*(j) - z_i^*(j)|$$

for $i = 1, 2$ and $j \in \mathbb{0:k^*-1}$ because x_i^* and z_i^* both receive the same inputs. Without loss of generality, we can assume $L_f \geq 1$. We thus have that

$$|x_i^*(k^*) - z_i^*(k^*)| \leq L_f^{k^*} |x_i - z_i| \leq L_f^{\bar{k}} |x_i - z_i|.$$

Similarly, we have that

$$|y_i^*(j) - \eta_i^*(j)| \leq L_h |x_i(j) - z_i(j)|$$

and hence

$$\|\mathbf{y}_i^* - \boldsymbol{\eta}_i^*\|_{0:k^*-1} \leq L_h L_f^{k^*} |x_i - z_i| \leq L_h L_f^{\bar{k}} |x_i - z_i|.$$

Finally, because $\lambda^{-k^*/2} \leq \lambda^{-\bar{k}/2}$, we can substitute these bounds into (2.15) to obtain

$$\begin{aligned} \Lambda(x_1, x_2) - \Lambda(z_1, z_2) &\leq \varepsilon + \lambda^{-\bar{k}/2} \left[L_f^{\bar{k}} |x_1 - z_1| + L_f^{\bar{k}} |x_2 - z_2| \right] \\ &\quad + C_y \left[L_h L_f^{\bar{k}} |x_1 - z_1| + L_h L_f^{\bar{k}} |x_2 - z_2| \right]. \end{aligned}$$

Let $L_\lambda := \lambda^{-\bar{k}/2} (L_f^{\bar{k}} + C_y L_h L_f^{\bar{k}})$. We have that

$$\Lambda(x_1, x_2) - \Lambda(z_1, z_2) \leq \varepsilon + L_\Lambda |x_1 - z_1| + L_\Lambda |x_2 - z_2|.$$

Because no remaining term of this expression depends upon ε , we can take the limit as ε goes to zero to obtain

$$\Lambda(x_1, x_2) - \Lambda(z_1, z_2) \leq L_\Lambda |x_1 - z_1| + L_\Lambda |x_2 - z_2|.$$

Finally, this same argument can be made for $\Lambda(z_1, z_2) - \Lambda(x_1, x_2)$, and we therefore have that

$$|\Lambda(x_1, x_2) - \Lambda(z_1, z_2)| \leq L_\Lambda |x_1 - z_1| + L_\Lambda |x_2 - z_2|.$$

Now suppose that $x_1 = x_2$ and $z_1 = z_2$. Then this bound holds trivially. Now suppose without loss of generality that $z_1 = z_2$ but $x_1 \neq x_2$. We have that

$$\begin{aligned} |\Lambda(x_1, x_2) - \Lambda(z_1, z_2)| &= \Lambda(x_1, x_2) \\ &\leq C_x |x_1 - x_2| \\ &= C_x |x_1 - z_1 + z_2 - x_2| \\ &\leq C_x |x_1 - z_1| + C_x |z_2 - x_2| \end{aligned}$$

Let $\tilde{L}_\Lambda := \max(L_\Lambda, C_x)$. We thus have that

$$|\Lambda(x_1, x_2) - \Lambda(z_1, z_2)| \leq \tilde{L}_\Lambda |x_1 - z_1| + \tilde{L}_\Lambda |x_2 - z_2|.$$

irrespective of $x_1, x_2, z_1, z_2 \in \mathbb{X}$, which concludes the proof. ■

2.4 Quadratic exp i-IOSS Lyapunov function

We have derived a Lipschitz continuous exponential i-IOSS Lyapunov function with power-law bounds with $\sigma = 1$. However, it is useful to generate an exp i-IOSS Lyapunov functions with $\sigma = 2$ due to its relationship with the ever-popular

quadratic stage cost. The key result that enables the generation of exponential i-IOSS Lyapunov functions with powers $\sigma > 1$ is that when a convex \mathcal{K}_∞ function ρ and an i-IOSS Lyapunov function $\lambda(\cdot)$ are composed, the result is another i-IOSS Lyapunov function (Allan et al., 2021, Cor. 5.4).

Consider the composition of an exponential i-IOSS Lyapunov function $\lambda(\cdot)$ with the convex \mathcal{K}_∞ function $\rho(s) = s^\sigma$ for some $\sigma > 1$. Let $\rho \circ \lambda(\cdot) := \tilde{\lambda}(\cdot)$. We immediately have that

$$c_1^\sigma |x_1 - x_2|^\sigma \leq \tilde{\lambda}(x_1, x_2) \leq c_2^\sigma |x_1 - x_2|^\sigma.$$

The dissipation condition is next. By applying the change of supply rate equations from (Grimm et al., 2005, Lemma 4) as applied in (Allan et al., 2021, Prop. 5.2), we can obtain

$$\begin{aligned} \tilde{\lambda}(x_1^+, x_2^+) &\leq \tilde{\lambda}(x_1, x_2) - \sigma \left(\frac{c_3 |x_1 - x_2|}{4} \right)^\sigma \\ &\quad + 2^\sigma \sigma \left(\frac{2c_2}{c_3} + 1 \right)^{\sigma-1} \left(c_u^\sigma |u_1 - u_2|^\sigma + c_y^\sigma |y_1 - y_2|^\sigma \right) \end{aligned}$$

which is an exp i-IOSS Lyapunov function with power σ . To summarize

Corollary 2.15. *If a system (2.2) is exponentially i-IOSS, then it admits an exponential i-IOSS Lyapunov function with any power $\sigma \geq 1$.*

2.5 Conclusions

In this chapter, we have seen that i-IOSS is a reasonable definition of nonlinear detectability, and that it can be characterized with a storage function, termed an

i-IOSS Lyapunov function. One of the key features about i-IOSS as a notion of detectability is that it does not require particular certain input sequences to be fed into the system in order to asymptotically distinguish two state trajectories. This structure permits a weak separation principle, in which a robustly stable estimator and stabilizing regulator robust to estimation errors can be designed separately from each other, then connected in cascade to generate a stable system.³

In the next chapter, we set the estimation problem aside for the moment and design such a controller. MPC is a well-established area of research, with a large variety of stability and robustness results. Recently, however, a subfield has grown in the design of an MPC regulator without the stabilizing terminal conditions that had grown popular in the larger field. Many of the results in this area, however, are dense and inaccessible to new researchers. A streamlined analysis with a Lyapunov-like function termed a Q function is introduced in the next chapter. Estimator design is returned to in Chapter 5, where it turns out that analysis with Q functions is what permits the conclusion that optimization-based state estimators like FIE and MHE are robustly stable.

³This weak separation principle, in which regulator and estimator are designed separately in order to achieve a robustly stable system, should be distinguished from the strong separation principle that exists for linear systems, in which an optimal regulator and estimator can be designed separately in order to create a system that is optimal overall. This strong sense of “separation principle” (or at least “separation theorem”) seems to be the original, see, e.g., (Wonham, 1968b). The more modern usage of “separation principle” and “separation theorem”, see, e.g., (Hespanha, 2018, p 208) and (Khalil, 2017, p 30) refers to this weaker notion.

CHAPTER 3

MPC AND STEADY STATE REGULATION

There is a rich literature about the regulation of a system's setpoint using MPC. Books, such as Rawlings et al. (2017) and Grüne and Pannek (2016), provide comprehensive treatment of the problem, as well as variations such as Economic MPC. The central problem of MPC research for the last thirty years of MPC research has been: how does one formulate a finite-horizon optimal control problem such that the finite-horizon controller inherits the stability properties of the infinite-horizon controller? Three approaches are dominant in the literature. The first is to use a terminal equality constraint at the desired steady-state. The second is to design a local approximation for the infinite-horizon cost to use as a terminal penalty, and constrain the terminal state to reside in this region. The third is to use a horizon long enough that the finite-horizon optimal cost function is a good enough approx-

imation to the infinite-horizon cost. The first two are covered in Rawlings et al. (2017), while all three are covered in Grüne and Pannek (2016).

While I had no trouble learning about MPC with stabilizing terminal conditions¹, I found learning about MPC stabilized using sufficiently long horizons out of Chapter 6 of Grüne and Pannek (2016) to be difficult. The authors are concerned with deriving tight lower bounds on the horizon length necessary to ensure that MPC is stabilizing, but their setting of an abstract optimization problem makes it easy to lose sight of the underlying control problem. I found that the long horizon result in Grimm et al. (2005) was easier to grasp, despite the complications introduced by use of a semidefinite but detectable stage cost (as opposed to one that is positive definite). It was to my great surprise, then that I found that Q functions, a tool that was introduced to be a Lyapunov-like function for state estimation problems in (Allan and Rawlings, 2019b), could also be used to understand this MPC design methodology. Later, I discovered similar results existed in Primbs and Nevistić (2000) for constrained linear systems and Jadbabaie and Hauser (2005) for continuous-time nonlinear systems. To my knowledge, however, there is no corresponding work in the setting of nonlinear discrete-time systems, and there are additional insights from Grimm et al. (2005) and Grüne and Pannek (2016) that can be applied using this method. Furthermore, such an analysis permits the introduction of Q functions in a context familiar to many in the control community, while far fewer are familiar with full information estimation, the context in which they were originally developed.

¹Taking a semester-long course taught out of Rawlings et al. (2017) with an author as the teacher certainly helped.

In this chapter, therefore, we analyze long-horizon MPC using Q functions. First, we formulate an infinite horizon optimal control problem (OCP). If the problem is formulated correctly and the system is stabilizable, this problem has a solution that attains a finite value. In the absence of hard state constraints, the optimal cost function is continuous, and therefore the infinite horizon control law robustly stabilizes the origin.

We can rarely calculate the infinite horizon control law, however, and so must use a finite horizon approximation. If the finite horizon optimal cost functions converge exponentially to the infinite horizon optimal cost, then MPC globally stabilizes the origin with a finite horizon. However, with a slower convergence rate, asymptotic stability of the origin cannot be guaranteed, and we provide an example in which the controller does converge to a state away from the origin. However, we can prove a weaker property, “semiglobal practical stability”, in which the system converges to a neighborhood of the origin.

Finally, we provide a comparison of the three methods to guarantee the stability of MPC. Although the results in this section are largely about long-horizon MPC, I think that using a terminal region is much more practical for nonlinear MPC. Under certain regularity conditions, the systems for which we can expect MPC formulated with a naive quadratic cost to work are precisely those for which we have a good design methodology for a terminal region and terminal cost, namely, those for which the linearized system is stabilizable. The outlook for a terminal equality constraint is even worse, because in order for gradient-based NLP solvers to behave well, the linearized system must be controllable. However, both of those

design methodologies are still relevant. For tracking a reference trajectory, terminal region formulations are awkward, and for estimation problems any sort of terminal constraint is out of the question, and only long-horizon formulations remain.

3.1 *Infinite-horizon optimal control problem*

Consider a system

$$x^+ = f(x, u) \tag{3.1}$$

in which $x \in \mathbb{X} \subseteq \mathbb{R}^n$ is the system state, $u \in \mathbb{U} \subseteq \mathbb{R}^m$ is the control input, and $f : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{X}$ is the state transition map. Note that \mathbb{X} does not represent a set of state constraints, but rather the domain on which $f(\cdot)$ is defined. For example, in chemical systems, it usually makes sense to consider only nonnegative concentrations and (absolute) temperatures. Therefore note that the codomain of $f(\cdot)$ is limited to \mathbb{X} as well, and thus \mathbb{X} is positive invariant under admissible controls. We typically write $x(k)$ for $\phi(k; x, \mathbf{u})$ when the initial state $x(0)$ and input sequence are clear from context.

We consider the problem of controlling this system to a steady state. Without loss of generality, we can assume that this steady state is the origin, i.e., $f(0, 0) = 0$. In order for this problem to be well-defined, we require a stabilizability assumption (compare (Kellett and Teel, 2004c, Definition 7) and (Grüne and Pannek, 2016, Definition 4.2)).

Definition 3.1 (Stabilizability with small controls). The origin is said to be stabilizable with small controls if there exist $\beta_x, \beta_u \in \mathcal{KL}$ such that for every $x \in \mathbb{X}$ there exists $\mathbf{u} \in \mathbb{U}^\infty$ such that

$$\begin{aligned} |\phi(k; x, \mathbf{u})| &\leq \beta_x(|x|, k) \\ |\mathbf{u}(k)| &\leq \beta_u(|x|, k) \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$.

We desire to achieve setpoint regulation using optimal control. Let $\ell : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ be some stage cost. We define the infinite-horizon cost function

$$V_\infty(\chi, \mu) := \sum_{k=0}^{\infty} \ell(\chi(k), \mu(k)) \quad \chi^+ = f(\chi, \mu)$$

as a function from $\mathbb{X} \times \mathbb{U}^\infty \rightarrow [0, \infty]$. The infinite-horizon OCP is then

$$\mathbb{P}_\infty(x) : \quad \inf_{\mathbf{u} \in \mathbb{U}^\infty} V_\infty(x, \mathbf{u}) := V_\infty^0(x).$$

Because $\ell(\cdot)$ is nonnegative, this infinite-horizon problem is the monotone limit of a sequence of finite horizon problems (for the more general stage costs found in economic MPC, Carlson et al. (1991) and Dong and Angeli (2018) are good references).

We define

$$V_N(\chi, \mu) := \sum_{k=0}^{N-1} \ell(\chi(k), \mu(k)) \quad \chi^+ = f(\chi, \mu)$$

as a cost function for a horizon length N , and

$$\mathbb{P}_N(x) : \quad \min_{\mathbf{u} \in \mathbb{U}^N} V_N(x, \mathbf{u}) := V_N^0(x).$$

as the OCP for a horizon length N . We do not distinguish between $V_N(\cdot)$ as a function of an infinite control sequence in \mathbb{U}^∞ and that of a finite sequence in

\mathbb{U}^N . We denote any² optimal solution for $\mathbb{P}_N(x)$ as $\mathbf{u}^0(N; x)$, the j^{th} control in this sequence as $u^0(j|N; x)$, the corresponding state sequence as $\mathbf{x}^0(N; x)$, the j^{th} forecast state as $x^0(j|N; x)$, and the optimal control law, i.e., the set of all $u^0(0|N; x)$ as $\mathcal{K}_N^0(x)$.

Under suitable assumptions, the infimum of the infinite-horizon OCP is finite and attained by some $\mathbf{u} \in \mathbb{U}^\infty$.

Assumption 3.2 (Continuity). The functions $f : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{X}$ and $\ell : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ are continuous, \mathbb{X} and \mathbb{U} are closed, and $f(0, 0) = 0$.

Assumption 3.3 (Positive definite $\ell(\cdot)$). There exist $\alpha_x, \alpha_u, \alpha_2 \in \mathcal{K}_\infty$ such that

$$\ell(x, u) \geq \alpha_x(|x|) + \alpha_u(|u|) \quad (3.2)$$

for all $x \in \mathbb{X}$ and $u \in \mathbb{U}$.

Assumption 3.4 (Stabilizability with respect to $\ell(\cdot)$). For every $x \in \mathbb{X}$ there exists $\mathbf{u}^* \in \mathbb{U}^\infty$ such that

$$V_\infty(x, \mathbf{u}^*) \leq \alpha_2(|x|). \quad (3.3)$$

We demonstrate that every system that is stabilizable with small controls admits a stage cost that satisfies Assumptions 3.3 and 3.4. We first require a version of Sontag's \mathcal{KL} -Lemma (Sontag, 1998, Proposition 7).

Proposition 3.5 (Lemma 5.3 in Kellett and Teel (2004c)). *For every $\beta \in \mathcal{KL}$ and $\lambda \in (0, 1)$ there exist $\theta_1, \theta_2 \in \mathcal{K}_\infty$ such that $\theta_1(\cdot)$ is Lipschitz continuous and*

$$\theta_1(\beta(s, k)) \leq \theta_2(s)\lambda^k$$

²In general, the solution to $\mathbb{P}_N(x)$ is not unique. However, it typically does not matter which solution is chosen, and the bounds derived are valid for *any* optimal solution.

for all $s \in \mathbb{R}_{\geq 0}$ and $k \in \mathbb{I}_{\geq 0}$.

Proposition 3.6. *If the origin is stabilizable with small controls for the system $x^+ = f(x, u)$, then there exists some continuous function $\ell : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ that satisfies Assumptions 3.3 and 3.4.*

Proof. By Proposition 3.5, there exist $\theta_{1,x}, \theta_{1,u}, \theta_{2,x}, \theta_{2,u} \in \mathcal{K}_\infty$ such that

$$\begin{aligned}\theta_{1,x}(\beta_x(s, k)) &\leq \theta_{2,x}(s)\lambda^k \\ \theta_{1,u}(\beta_u(s, k)) &\leq \theta_{2,u}(s)\lambda^k\end{aligned}$$

in which $\lambda \in (0, 1)$, for all $s \in \mathbb{R}_{\geq 0}$ and $k \in \mathbb{I}_{\geq 0}$. Let

$$\ell(x, u) := \theta_{1,x}(|x|) + \theta_{1,u}(|u|)$$

This choice of $\ell(\cdot)$ clearly satisfies (3.2), and, because of the definition of \mathcal{K} functions, is continuous. Finally, we note that because the origin is stabilizable with small controls, for each $x \in \mathbb{X}$ there exists $\mathbf{u} \in \mathbb{U}^\infty$ such that

$$\begin{aligned}|\phi(k; x, \mathbf{u})| &\leq \beta_x(|x|, k) \\ |\mathbf{u}(k)| &\leq \beta_u(|x|, k)\end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$. We can then evaluate $V_\infty(\cdot)$ with that \mathbf{u} to obtain

$$\begin{aligned}
V_\infty(x, \mathbf{u}) &= \sum_{k=0}^{\infty} \theta_{1,x}(|x(k)|) + \theta_{1,u}(|u(k)|) \\
&\leq \sum_{k=0}^{\infty} \theta_{1,x}(\beta_x(|x|, k)) + \theta_{1,u}(\beta_u(|x|, k)) \\
&\leq \sum_{k=0}^{\infty} \lambda^k (\theta_{2,x}(|x|) + \theta_{2,u}(|x|)) \\
&\leq \frac{1}{1-\lambda} (\theta_{2,x}(|x|) + \theta_{2,u}(|x|))
\end{aligned}$$

and, as a result, satisfies (3.3). ■

Remark 3.7. This technique to obtain a stage cost for an optimal control problem previously appeared in (Grüne and Nešić, 2003) and (Grüne and Pannek, 2016, Theorem 4.3).

Now that we have a suitable stage cost, we can demonstrate that a solution exists to $\mathbb{P}_\infty(x)$ for all $x \in \mathbb{X}$. Demonstrating the existence of a solution requires some basic topology however. For a reader who is unfamiliar with topology, as I was when I first investigated this matter, suffice it to say that properties such as function continuity and the closedness or compactness of sets is relative to a particular topology. For the set $\mathbb{X} \times \mathbb{U}^\infty$, over which $V_\infty(\cdot)$ is defined, there is no natural topology in the way that there is for subsets of \mathbb{R}^n . For our purposes here, the most important topology turns out to be the *product topology*, which is the topology that has the least restrictive criteria for convergence, i.e., the coarsest topology, for which all projection maps are continuous.

One problem is that $V_\infty(\cdot)$ is not a continuous function in the product topology. To illustrate this unfortunate property of $V_\infty(\cdot)$, consider a system that is not open-loop stable, such as $x^+ = 2x + u$. From the starting point $x(0) = 1$, the control sequence $(-2, 0, 0, \dots)$ attains a finite infinite-horizon cost. The sequence $(-2 + \delta, 0, 0, \dots)$, however, attains an infinite cost for any $\delta \neq 0$. This problem cannot be fixed by restricting ourselves to sets for which $V_\infty(\cdot)$ attains a finite cost. A similar perturbation to a (necessarily nonlinear) system that has an attractive but unstable equilibrium would result in a finite perturbation of cost with an infinitesimal change in state or input. However, we *can* prove that $V_\infty(\cdot)$ is lower semicontinuous.

Proposition 3.8. *If Assumptions 3.2 and 3.3 hold, then the sublevel sets $\text{lev}_\rho V_\infty := \{(x, \mathbf{u}) \in \mathbb{X} \times \mathbb{U}^\infty \mid V_\infty(x, \mathbf{u}) \leq \rho\}$ are compact in the product topology. As a result, $V_\infty(\cdot)$ is lower semicontinuous.*

Proof. For $\rho < 0$, $\text{lev}_\rho V_\infty = \emptyset$ and is compact trivially. For all $\rho \geq 0$, $(0, \mathbf{0}) \in \text{lev}_\rho V_\infty$. Because the product topology of $\mathbb{X} \times \mathbb{U}^\infty$ is metrizable, it suffices to show that every sequence has a convergent subsequence (Morris, 2020, Ex. 7.2.22).

Let $(x_i, \mathbf{u}_i) \in \text{lev}_\rho V_\infty$ for all $i \in \mathbb{N}_{\geq 0}$. We have that $V_\infty(x_i, \mathbf{u}_i) \geq \ell(x_i(k), u_i(k))$ for all $i, k \in \mathbb{N}_{\geq 0}$. By Assumption 3.3, we have that

$$\begin{aligned} |x_i| &\leq \alpha_x^{-1}(\rho) \\ |u_i(k)| &\leq \alpha_u^{-1}(\rho) \end{aligned}$$

for all $i, k \in \mathbb{N}_{\geq 0}$. Because \mathbb{X} and \mathbb{U} are closed, we have that $\mathbb{X} \cap \alpha_{1,x}^{-1}(\rho)\mathbb{B} := \tilde{\mathbb{X}}$ and $\mathbb{U} \cap \alpha_{1,u}^{-1}(\rho)\mathbb{B} := \tilde{\mathbb{U}}$ are compact, and therefore the sequence (x_i, \mathbf{u}_i) is evolving in

$\tilde{\mathbb{X}} \times \tilde{\mathbb{U}}^\infty$, which, as a product of compact sets, is compact in the product topology by Tychonoff's theorem. Therefore, there exists a convergent subsequence, which we do not relabel, that has a limit $(\bar{x}, \bar{\mathbf{u}}) \in \mathbb{X} \times \mathbb{U}^\infty$. All that remains is to show that $V_\infty(\bar{x}, \bar{\mathbf{u}}) \leq \rho$.

By Assumption 3.2, we have that $f(\cdot)$, $l(\cdot)$, and therefore $V_N(\cdot)$ are continuous for all $N \in \mathbb{N}_{\geq 0}$. Therefore, $V_N(x_i, \mathbf{u}_i)$ converges to $V_N(\bar{x}, \bar{\mathbf{u}})$ for all $N \in \mathbb{N}_{\geq 0}$. We have that $V_N(x_i, \mathbf{u}_i) \leq V_\infty(x_i, \mathbf{u}_i) \leq \rho$ for all $i, N \in \mathbb{N}_{\geq 0}$. Therefore $V_N(\bar{x}, \bar{\mathbf{u}}) \leq \rho$ for all $N \in \mathbb{N}_{\geq 0}$. The sequence $(V_N(\bar{x}, \bar{\mathbf{u}}))$ indexed by N is nondecreasing and bounded above by ρ , and therefore it converges to $V_\infty(\bar{x}, \bar{\mathbf{u}}) \leq \rho$. Thus $(\bar{x}, \bar{\mathbf{u}}) \in \text{lev}_\rho V_\infty$ and $\text{lev}_\rho V_\infty$ is compact. Finally, because $\mathbb{X} \times \mathbb{U}^\infty$ is metrizable, compact sets must be closed (Morris, 2020, Cor. 7.2.6). As a result, all sublevel sets of $V_\infty(\cdot)$ are closed, and thus it is lower semicontinuous (Springer Verlag GmbH, European Mathematical Society, Theorem 4). ■

With this result, demonstrating that the infinite-horizon OCP has a solution is straightforward.

Proposition 3.9. *If Assumptions 3.2 to 3.4 are satisfied, then the infimum in $\mathbb{P}_\infty(x)$ is both finite and attained by some $\mathbf{u} \in \mathbb{U}^\infty$ for all $x \in \mathbb{X}$. Furthermore, by construction, $V_\infty^0(x) \leq \alpha_2(|x|)$.*

Proof. Let $x \in \mathbb{X}$. By Assumption 3.4, there exists a sequence $\mathbf{u}^* \in \mathbb{U}^\infty$ such that $V_\infty(x, \mathbf{u}^*) \leq \alpha_2(|x|)$. Therefore, for some ρ , we have that $(x, \mathbf{u}^*) \in \text{lev}_\rho V_\infty$. We have that $\mathcal{U}_\rho^\infty(x) := \{\mathbf{u} \mid (x, \mathbf{u}) \in \text{lev}_\rho V_\infty\}$ is compact-valued for all x . Because $V_\infty(x, \mathbf{u})$ is a lower semicontinuous function being minimized over a compact set,

the minimum is attained (Springer Verlag GmbH, European Mathematical Society, Theorem 6). Finally, by optimality, we have that $V_\infty^0(x) \leq V_\infty(x, \mathbf{u}^*) \leq \alpha_2(|x|)$. ■

Remark 3.10. Although this proof appears very different than that in Keerthi and Gilbert (1985), they are fundamentally rather similar. They generate a candidate optimal input sequence by a diagonalization argument, similar to some proofs of Tychonoff's theorem for countable products. They then show that, because the finite horizon control laws are increasing in horizon length and bounded above, their value converges to that of the optimal infinite-horizon control law. They consider a more general setting with (time-varying) joint state and input constraints, a time-varying system and time-varying lower semicontinuous stage cost.

One of the most attractive features of infinite-horizon optimal control is that, with a suitable stage cost, the infinite-horizon cost is a Lyapunov function.

Definition 3.11 (Lyapunov function). A function $V : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is a Lyapunov function for the autonomous system $x^+ = f(x)$ in the positive invariant set \mathbb{X} if there exist $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ such that

$$\begin{aligned} \alpha_1(|x|) &\leq V(x) \leq \alpha_2(|x|) \\ V(f(x)) &\leq V(x) - \alpha_3(|x|) \end{aligned}$$

for all $x \in \mathbb{X}$.

The existence of a Lyapunov function is important, because it implies that the system is asymptotically stable (Rawlings et al., 2017, Theorem B.13).

Definition 3.12 (Asymptotic stability). A system $x^+ = f(x)$ is asymptotically stable if there exists $\beta \in \mathcal{KL}$ such that

$$|x(k)| \leq \beta(|x(0)|, k)$$

for all $x \in \mathbb{X}$.

By Assumption 3.3 and Proposition 3.9, we immediately have that

$$\alpha_x(|x|) \leq V(x) \leq \alpha_2(|x|)$$

We define

$$\mathcal{U}_\infty^0(x) := \arg \min_{\mathbf{u} \in \mathbb{U}^\infty} V_\infty(x, \mathbf{u}),$$

the set of solutions to $\mathbb{P}(x)$. Clearly, V_∞^0 satisfies the Bellman equation

$$V_\infty^0(f(x, u)) + \ell(x, u) = V_\infty^0(x)$$

for all $u \in \mathcal{K}_\infty^0(x)$, and, as a result, we have that

$$V_\infty^0(f(x, u)) \leq V_\infty^0(x) - \alpha_x(|x|)$$

for all $u \in \mathcal{K}_\infty^0(x)$. Thus, when u is chosen to be optimal, $V_\infty^0(\cdot)$ is a Lyapunov function.

Now that we have established that $\mathcal{K}_\infty^0(\cdot)$ stabilizes the origin. We would like if it robustly stabilized the origin, however. It is known that a discrete-time system is robustly stable if and only if it admits a continuous Lyapunov function (Kellett and Teel, 2004b, Theorem 10).³ Therefore, we next show that $V_\infty^0(\cdot)$ is a continuous function, and, in so doing, require our first Q function.

³Note, however, that this Lyapunov function need not necessarily be the optimal value function from MPC. See Allan et al. (2017) and Allan and Rawlings (2018a) for examples of systems for which MPC is robustly stabilizing despite having a discontinuous optimal value function.

Proposition 3.13. *Under Assumptions 3.2 to 3.4, the function $V_\infty^0 : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is continuous.*

Proof. Consider the finite-horizon optimal control problem with horizon length k , $\mathbb{P}_k(x)$. Because there are no state constraints present, the set of feasible \mathbf{u} does not depend on x . As a result, by the application of Polak (1997, Theorem 5.4.3) to optimal control problems in Rawlings et al. (2017, Theorem C.28), each $V_k^0(\cdot)$ is a continuous function on \mathbb{X} . Let

$$Q(j|k; x) := V_\infty^0(x) - \sum_{i=0}^{j-1} \ell(x^0(i|k; x), u^0(i|k; x)). \quad (3.4)$$

Because a truncation of $\mathbf{u}_\infty^0(x)$ is feasible for $\mathbb{P}_k(x)$, we have that

$$V_\infty^0(x) \geq V_k^0(x) \geq \sum_{i=0}^{j-1} \ell(x^0(i|k; x), u^0(i|k; x))$$

and thus

$$Q(j|k; x) \geq 0$$

for all $j \in \mathbb{I}_{0:k}$. Note that

$$Q(k|k; x) = V_\infty^0(x) - V_k^0(x) \geq 0.$$

Thus, to prove that $\left(V_k^0(x)\right)_{k=1}^\infty$ converges uniformly to $V_\infty^0(x)$ we need only prove that $Q(k|k; x)$ converges uniformly to zero on compact subsets of \mathbb{X} . Because we have that $\mathbf{u}^0(k; x)$ is feasible for the infinite-horizon problem, we have that

$$\begin{aligned}
V_{\infty}^0(x) &\leq \sum_{i=0}^{j-1} \ell(x^0(i|k; x), u^0(i|k; x)) + \min_{\mu \in \mathbb{U}^{\infty}} \sum_{i=j}^{\infty} \ell(\chi(i), \mu(i)) \\
&\quad \text{subject to } \chi^+ = f(\chi, \mu) \\
&\quad \chi(j) = x^0(j|k; x) \\
&= \sum_{i=0}^{j-1} \ell(x^0(i|k; x), u^0(i|k; x)) + V_{\infty}^0(x^0(j|k; x))
\end{aligned}$$

and thus

$$Q(j|k; x) \leq V_{\infty}^0(x^0(j|k; x)) \leq \alpha_2 (|x^0(j|k; x)|) \quad (3.5)$$

for all $j \in \mathbb{I}_{0:k}$ and $k \in \mathbb{I}_{\geq 0}$. Finally, we have that

$$\begin{aligned}
Q(j+1|k; x) - Q(j|k; x) &= \sum_{i=0}^{j-1} \ell(x^0(i|k; x), u^0(i|k; x)) \\
&\quad - \sum_{i=0}^j \ell(x^0(i|k; x), u^0(i|k; x)) \\
&= -\ell(x^0(j|k; x), u^0(j|k; x)) \\
&\leq -\alpha_3 (|x^0(j|k; x)|). \quad (3.6)
\end{aligned}$$

for all $j \in \mathbb{I}_{0:k-1}$. We can substitute (3.5) into (3.6) to obtain

$$Q(j+1|k; x) \leq Q(j|k; x) - \alpha_3(\alpha_2^{-1}(Q(j|k; x)))$$

for all $j \in \mathbb{I}_{0:k-1}$. By a standard construction, provided in Rawlings et al. (2017, Theorem B.15), there exists $\sigma \in \mathcal{K}_{\infty}$ such that $s - \alpha_3(\alpha_2^{-1}(s)) \leq \sigma(s)$ and $\sigma(s) < s$ for $s > 0$. Thus we have that

$$Q(j+1|k; x) \leq \sigma(Q(j|k; x))$$

for $j \in \mathbb{I}_{0:k-1}$. We can then apply this equation recursively for $j \in \mathbb{I}_{0:k-1}$ in order to obtain

$$Q(k|k; x) \leq \sigma^k(Q(0|k; x)).$$

We can apply (3.4) for $j = 0$ to note that

$$Q(0|k; x) := V_\infty^0(x) \leq \alpha_2(|x|)$$

so we have that

$$Q(k|k; x) \leq \sigma^k(\alpha_2(|x|)) := \beta_Q(|x|, k)$$

and note that $\beta_Q \in \mathcal{KL}$. Let C be some compact subset of \mathbb{X} , and let $M := \max_{x \in C} |x|$. Fix $\epsilon > 0$. For all $x \in C$, we have that

$$Q(k|k; x) \leq \beta_Q(M, k).$$

Because $\beta_Q \in \mathcal{KL}$, there exists some T such that $\beta_Q(M, k) \leq \epsilon$ for all $k \geq T$. Thus $Q(k|k; x)$ converges to zero uniformly on compact subsets of \mathbb{X} , and thus, by Rudin (1976, Theorem 7.12), we have that $V_\infty^0(\cdot)$ is continuous on \mathbb{X} . ■

Remark 3.14. Note that the value of $Q(j|k; x)$ may depend on which $\mathbf{u} \in \mathcal{K}_k^0(x)$ is chosen. The bounds for $Q(j|k; x)$ hold along *all* possible choices for \mathbf{u} , and $Q(k|k; x)$ does *not* depend on the choice of $\mathbf{u} \in \mathcal{K}_k^0(x)$.

Since we have just derived our first Q function, it is useful to state its properties succinctly. We note that $Q(j|k; x) \geq 0$ for all $j \leq k \in \mathbb{N}_{\geq 0}$ and $x \in \mathbb{X}$, but the descent condition (3.6) allows us to strengthen this nonnegativity condition to $Q(j|k; x) \geq \alpha_x(|x^0(j|k)|)$ for $j \in \mathbb{N}_{0:k-1}$.⁴ We thus have that

⁴In retrospect, it is probably more natural to include a terminal penalty of $\ell^*(\chi(N)) := \min_{\mu \in \mathbb{U}} \ell(\chi, \mu)$ in $\mathbb{P}_N(x)$, but much of the literature, including Grüne and Pannek (2016, Ch. 6), does not include such a penalty. We note that $\mathbb{P}_N(x)$ without such a penalty is equivalent to $\mathbb{P}_{N-1}(x)$ with such a penalty.

$$\begin{aligned}\alpha_x(|x^0(j|k; x)|) &\leq Q(j|k; x) \leq \alpha_2(|x^0(j|k; x)|) \\ Q(j+1|k; x) &\leq Q(j|k; x) - \alpha_x(|x^0(j|k; x)|)\end{aligned}$$

for all $j \in \mathbb{0}_{:k-1}$, $k \in \mathbb{0}_{\geq 0}$, and $x \in \mathbb{X}$. Additionally we have that $Q(k|k; x) \leq \alpha_2(x^0(k|k; x))$. This function looks a lot like a Lyapunov function, but it tells us about *open-loop* state trajectories, rather than closed-loop ones. We could find $\beta \in \mathcal{KL}$ such that $|x(j|k; x)| \leq \beta(|x|, j)$ for all $j \in \mathbb{0}_{:k-1}$, $k \in \mathbb{0}_{\geq 0}$, and $x \in \mathbb{X}$, giving us the (rather obvious) conclusion that the forecast state in the open-loop trajectory converges to zero. It also gives us an explicit *convergence rate* for these forecast states, something, to my knowledge, unique in the literature. Convergence of these forecast states becomes much less obvious when merely a semidefinite stage cost is used in Chapter 4, and finding such a bound for *estimated* states is the goal of analysis in Chapter 5.

One useful consequence of the continuity of $V_\infty^0(\cdot)$ is that it constitutes a global control Lyapunov function (CLF). CLFs were first defined by Sontag (1983) to characterize stabilizability using a storage function, and became important in the study of MPC when local control Lyapunov functions became important for the design of terminal costs and terminal regions.

Definition 3.15 (Control Lyapunov function (CLF)). A function $V : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is a control Lyapunov function with respect to some stage cost $\ell : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ if there exist $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ such that

$$\begin{aligned}
\alpha_1(|x|) &\leq V(x) \leq \alpha_2(|x|) \\
\forall u \in \mathbb{U} \quad \ell(x, u) &\geq \alpha_3(|x|) \\
\inf_{u \in \mathbb{U}} [V(f(x, u)) + \ell(x, u)] &\leq V(x)
\end{aligned}$$

for all $x \in \mathbb{X}$.

By construction of $V_\infty^0(\cdot)$, it is obvious that it's a continuous CLF.

Theorem 3.16. *A continuous system (3.1) is stabilizable with small controls if and only if it admits a continuous CLF $V : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$.*

As observed by Kellett and Teel (2004a), the infinite-horizon existence result by Keerthi and Gilbert (1985) results in a (possibly discontinuous) CLF. Furthermore, Grüne and Nešić (2003) produce a continuous infinite-horizon cost function for sampled data systems. However, their result is tangled up with additional assumptions required in order that the sampled controller stabilize the underlying continuous-time system.

To my knowledge, the most general converse theorem for discrete-time CLFs in the literature is given by Kellett and Teel (2004a). This theorem is problematic to apply when \mathbb{U} is not (star) convex. It requires that the difference inclusion $F(x) := f(x, \sigma(|x|)\mathbb{B} \cap \mathbb{U})$, with $\sigma(\cdot)$ being continuous and nondecreasing, be continuous. While outer semicontinuity of a difference inclusion is relatively easy to fulfill, full continuity is much harder. To guarantee $F(\cdot)$ is continuous without assumptions beyond continuity of $f(\cdot)$, $\sigma(|x|)\mathbb{B} \cap \mathbb{U}$ must be continuous. If $\sigma \in \mathcal{K}$, the system is called stabilizable with small controls, and in order to guarantee

continuity of $\sigma(|x|)\mathbb{B} \cap \mathbb{U}$, (star) convexity of \mathbb{U} is required. One might get lucky and obtain continuity for nonconvex \mathbb{U} , but they cannot be satisfied for disconnected \mathbb{U} with a nontrivial $f(\cdot)$. Disconnected input domains are important for describing MPC with discrete actuators (see (Rawlings and Risbeck, 2017)), and here we have made no assumption, explicit or implicit, about the connectedness of \mathbb{U} . If $\sigma(\cdot)$ is chosen such that $\sigma(0)\mathbb{B}$ contains all nonconvexities of \mathbb{U} , then the converse theorem from Kellett and Teel (2004a) can be applied. However, then the concept of controllability with small controls is lost. Because discrete decisions are inherently large control actions, we might prefer that continuous decisions are sufficient to stabilize the system in a neighborhood of the origin. In this case the theorem from Kellett and Teel (2004a) is not applicable.

There are three respects in which the result here is less general than the one in (Kellett and Teel, 2004a). The first is that they produce a smooth CLF, whereas here only a continuous CLF is produced. It seems that obtaining a continuous CLF is the difficult step and afterwards a result like (Kellett and Teel, 2004a, Lemma 22) can be used to produce a smooth function. However, it also seems like having \mathbb{X} open, or at least the closure of an open set, is required for that additional step. Being unfamiliar with smoothing results, I cannot comment further on how easy or difficult smoothing $V_\infty^0(\cdot)$ would be. The second respect is that they produce an exponential-decrease CLF, whereas here a CLF with a \mathcal{K}_∞ descent condition is produced. The third respect is that Kellett and Teel (2004a) characterize stability of a closed set \mathcal{A} that is positive invariant for the unforced system $x^+ = f(x, 0)$. The result presented here can be extended easily to a more general version of

stabilizability to the compact set $\mathcal{A} \subseteq \mathbb{X}$ with controls that are small relative to the compact set $\mathcal{B} \subseteq \mathbb{U}$.

Corollary 3.17. *Let $\mathcal{A} \subseteq \mathbb{X}$ and $\mathcal{B} \subseteq \mathbb{U}$ be compact. Suppose that the system (3.1) is stabilizable to \mathcal{A} with \mathcal{B} small controls, i.e., there exist $\beta_x, \beta_u \in \mathcal{KL}$ such that for every $x \in \mathbb{X}$ there exists $\mathbf{u} \in \mathbb{U}^\infty$ such that*

$$\begin{aligned} |x(k)|_{\mathcal{A}} &\leq \beta_x(|x(0)|_{\mathcal{A}}, k) \\ |u(k)|_{\mathcal{B}} &\leq \beta_u(|x(0)|_{\mathcal{A}}, k) \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$. Then there exists $\ell : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ and $\alpha_2, \alpha_x, \alpha_u \in \mathcal{K}_\infty$ such that

$$\ell(x, u) \geq \alpha_x(|x|_{\mathcal{A}}) + \alpha_u(|u|_{\mathcal{B}})$$

for all $(x, u) \in \mathbb{X} \times \mathbb{U}$ and $V_\infty^0 : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is well-defined and continuous, such that

$$\begin{aligned} \alpha_x(|x|_{\mathcal{A}}) &\leq V_\infty^0(x) \leq \alpha_2(|x|_{\mathcal{A}}) \\ \min_{u \in \mathbb{U}} [V_\infty^0(f(x, u)) + \ell(x, u)] &= V_\infty^0(x) \end{aligned}$$

for all $x \in \mathbb{X}$.

It is possible that this result can be extended to closed but unbounded \mathcal{A} , but Proposition 3.8 would have to be modified to deal with sublevel sets of $V_\infty(x, \cdot)$ for fixed $x \in \mathbb{X}$, rather than sublevel sets of $V_\infty(\cdot, \cdot)$.

3.2 Semiglobal practical stability with a finite horizon

While the stability and robustness of the infinite-horizon controller are important theoretically, in very few cases can it be implemented in practice. We wish to

understand the circumstances under which a finite horizon is sufficient to stabilize the origin. Fortunately, the proof of Proposition 3.13 has already furnished for us a method to proceed. There, we used the uniform convergence of $V_N^0(\cdot)$ to $V_\infty^0(\cdot)$ in order to show that the infinite-horizon cost inherits continuity from the finite-horizon cost. Now, we show that the finite-horizon cost inherits approximate satisfaction of the Bellman equation from the infinite-horizon cost.

Proposition 3.18. *We have both*

$$\begin{aligned} V_k^0(f(x, u)) &\leq V_k^0(x) - \alpha_x(|x|) + \beta_Q(|f(x, u)|, k - 1) \quad \text{and} \\ V_\infty^0(f(x, u)) &\leq V_\infty^0(x) - \alpha_x(|x|) + \beta_Q(|f(x, u)|, k - 1) \end{aligned}$$

for all $u \in \mathcal{K}_N^0(x)$ and $x \in \mathbb{X}$.

Proof. We have that

$$V_\infty^0(x) - V_k^0(x) := Q(k|k; x) \leq \beta_Q(|x|, k)$$

from the proof of Proposition 3.13. We then have that

$$V_{k-1}^0(f(x, u)) = V_k^0(x) - \ell(x, u)$$

for all $u \in \mathcal{K}_N^0(x)$ by the principle of optimality. We thus have that

$$\begin{aligned} V_k^0(f(x, u)) &\leq V_\infty^0(f(x, u)) \leq V_{k-1}^0(f(x, u)) + \beta_Q(|f(x, u)|, k - 1) \\ &= V_k^0(x) - \ell(x, u) + \beta_Q(|f(x, u)|, k - 1) \\ &\leq V_k^0(x) - \alpha_x(|x|) + \beta_Q(|f(x, u)|, k - 1) \\ &\leq V_\infty^0(x) - \alpha_x(|x|) + \beta_Q(|f(x, u)|, k - 1) \end{aligned}$$

for all $u \in \mathcal{K}_N^0(x)$ by Assumption 3.3 and because $V_k^0(x) \leq V_\infty^0(x)$ for all $x \in \mathbb{X}$.

Thus we have both sets of inequalities. ■

Unfortunately, this approximate satisfaction of the Bellman equation implies only approximate stability. Without further assumptions, the best result that is possible is semiglobal practical stability. Instead of being attracted to the origin, attraction to a neighborhood of the origin is all that is possible, and this neighborhood attracts only a compact subset of \mathbb{X} . However, the neighborhood of the origin can be made arbitrarily small and the region of attraction arbitrarily large by use of a sufficiently long control horizon.

In order to demonstrate this result, we first need a few preliminary results. Let $\mathcal{X}_N^\rho := \{x \mid V_N^0(x) \leq \rho\}$.

Proposition 3.19. *For $M \leq N \in \mathbb{N}_{\geq 0}$ and $\rho \geq 0$, we have that*

$$\alpha_2^{-1}(\rho)\mathbb{B} \subseteq \mathcal{X}_\infty^\rho \subseteq \mathcal{X}_N^\rho \subseteq \mathcal{X}_M^\rho \subseteq \alpha_x^{-1}(\rho)\mathbb{B}$$

Proof. Note that we have that

$$\alpha_x(|x|) \leq V_M^0(x) \leq V_N^0(x) \leq V_\infty^0(x) \leq \alpha_2(|x|)$$

for all $x \in \mathbb{X}$. This order is reversed for sublevel set inclusion. ■

Next, we require a proposition that allows us to guarantee that Proposition 3.18 implies a decrease in either $V_N^0(\cdot)$ or $V_\infty^0(\cdot)$. The basic idea is to choose N large enough to “bend” the curve of $\beta(s, N)$ beneath that of $(1/2)\alpha_x(s)$, i.e., choose N sufficiently large that $\beta(s, N) \leq (1/2)\alpha_x(s)$. The factor of $1/2$ allows us to absorb the term of $\beta(s, N)$ into $\alpha_x(s)$ while still leaving enough slack for a strict cost decrease. It turns out that, without further assumptions about $\beta(\cdot)$ and $\alpha(\cdot)$, we cannot guarantee this for all s , but only for s in a finite interval bounded away from zero.

Proposition 3.20. *Let $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$. Then, for any $0 < s_1 \leq s_2$, there exists some t^* such that for all $t \geq t^*$ we have that*

$$\beta(s, t) \leq \gamma(s)$$

for all $s \in [s_1, s_2]$.

Proof. This proof is inspired by part of the proof of (Hu, 2017, Lemma 6). By (Sontag, 1998, Lemma 8), there exist $\sigma \in \mathcal{K}$ and $\psi \in \mathcal{L}$ such that $\beta(s, t) \leq \sigma(s)\psi(t)$ for all $s \in \mathbb{R}_{\geq 0}$ and $t \in \mathbb{I}_{\geq 0}$. It is then sufficient to prove that $\sigma(s)\psi(t) \leq \gamma(s)$. Let

$$\mu := \max_{s \in [s_1, s_2]} \sigma(s)/\gamma(s)$$

and because the ratio $\gamma(s)/\sigma(s)$ is well-defined and continuous on $(0, \infty)$, this maximum exists and is finite. Then there exists some t^* such that $\psi(t) \leq 1/\mu$ for all $t^* \geq t$. Thus, for $s \in [s_1, s_2]$ and $t \geq t^*$, we have that

$$\sigma(s)\psi(t) \leq \mu\gamma(s)\psi(t) \leq (1/\mu)\mu\gamma(s) = \gamma(s)$$

and thus the statement is established. ■

With these preliminary results, we can demonstrate semiglobal practical stability.

Theorem 3.21. *There exists $\beta_V \in \mathcal{KL}$ such that for every $\Delta > \delta > 0$ there exists a horizon length $N^*(\Delta, \delta)$ such that for all $N \geq N^*$ we have that*

$$|x(k)| \leq \beta_V(|x(0)|, k) + \delta$$

under the feedback $u \in \mathcal{K}_N^0(x)$ for all $x(0) \in \Delta\mathbb{B}$ and $k \in \mathbb{I}_{\geq 0}$.

Proof. For a proof strategy, we use Proposition 3.20 in order to guarantee that Proposition 3.18 gives a strict descent condition in some annular region of \mathbb{X} . We

cannot guarantee that the descent conditions extends to the origin or to infinity without stronger assumptions. In the region inside the annulus, our (practical) Lyapunov function can increase, but not enough to bypass the annular region.

Let $x^+ := f(x, u)$ for some $u \in \mathcal{K}_N^0(x)$.

Claim 3.22. Let $\rho > 0$. Then, if we have N sufficiently large that $\beta_Q(s, N - 1) \leq (1/2)\alpha_3(s)$ for $s \in [\alpha_2^{-1}(\rho), \alpha_x^{-1}(\rho)]$, we have that the set \mathcal{X}_∞^ρ is positive invariant under the control law $\mathcal{K}_N^0(\cdot)$.

Proof. We want to use Proposition 3.18 in order to achieve a cost decrease. In order to do so, however, we need an upper bound for x^+ . The best upper bound available in the literature, to our knowledge, is the one used in (Grüne and Pannek, 2016, Theorem 6.35).⁵ We have that

$$\begin{aligned} \alpha_x(|x^0(j|k; x)|) &\leq \ell(x^0(j|k; x), u^0(j|k; x)) \\ &\leq \sum_{j=0}^{k-1} \ell(x^0(j|k; x), u^0(j|k; x)) := V_k^0(x) \leq V_\infty^0(x) \end{aligned}$$

Thus, if $x \in \mathcal{X}_\infty^\rho$, we have that $|x^+| \leq \alpha_x^{-1}(\rho)$. Suppose $\alpha_2^{-1}(\rho) < |x^+|$. Then $\alpha_2^{-1}(\rho) \leq |x^+| \leq \alpha_x^{-1}(\rho)$. Therefore, by assumption we have that

$$V_\infty^0(x^+) \leq V_\infty^0(x) - \alpha_x(|x|) + \beta_Q(|x^+|, N - 1) \leq V_\infty^0(x) - (1/2)\alpha_x(|x|) \leq \rho$$

so $x^+ \in \mathcal{X}_\infty^\rho$. Now suppose that $|x^+| \leq \alpha_2^{-1}(\rho)$. Then we have that

$$V_\infty^0(x^+) \leq \alpha_2(|x^+|) \leq \rho$$

and therefore $x^+ \in \mathcal{X}_\infty^\rho$. Thus \mathcal{X}_∞^ρ is positive invariant under $\mathcal{K}_N^0(\cdot)$. ■

⁵(Grimm et al., 2005, Theorem 2) uses a similar concept, but with additional complications because they use a stage cost that is merely detectable, but not positive definite.

This claim gives us a way of proceeding. If we choose $\rho_2 = \alpha_2(\Delta)$, we can ensure all states in $\Delta\mathbb{B}$ remain $\mathcal{X}_\infty^{\rho_2}$. Furthermore, if we choose $\rho_1 = \alpha_x(\delta)$, we can ensure that all states that arrive in $\mathcal{X}_\infty^{\rho_1}$ remain in $\delta\mathbb{B}$. All that remains is to ensure that $\mathcal{X}_\infty^{\rho_1}$ is attractive in $\mathcal{X}_\infty^{\rho_2}$.

By Proposition 3.20, there exists some N^* such that for all $N \geq N^*$, we have that $\beta(s, N-1) \leq (1/2)\alpha_3(s)$ for all $s \in [\alpha_2^{-1}(\rho_1), \alpha_x^{-1}(\rho_2)] = [\alpha_x \circ \alpha_2^{-1}(\delta), \alpha_2 \circ \alpha_x^{-1}(\Delta)]$. Suppose that $x \in \mathcal{X}_\infty^{\rho_1}$. Then, by Claim 3.22, we have that $x^+ \in \mathcal{X}_\infty^{\rho_1}$. Now suppose that $x \in \mathcal{X}_\infty^{\rho_2} \setminus \mathcal{X}_\infty^{\rho_1}$. We have that $x^+ \in \mathcal{X}_\infty^{\rho_2}$ by Claim 3.22. If $x^+ \in \mathcal{X}_\infty^{\rho_1}$, we have, by definition that $V_\infty^0(x^+) \leq \rho_1$. Suppose, then, that $x^+ \in \mathcal{X}_\infty^{\rho_2} \setminus \mathcal{X}_\infty^{\rho_1}$. By Proposition 3.19, we have that $\alpha_2^{-1}(\rho_1) < |x^+| \leq \alpha_1^{-1}(\rho_2)$. Therefore, by Proposition 3.18, we have that

$$V_\infty^0(x^+) \leq V_\infty^0(x) - \alpha_x(|x|) + \beta_Q(|x^+|, N-1) \leq V_\infty^0(x) - (1/2)\alpha_x(|x|)$$

By noting that $\alpha_2^{-1}(V_\infty^0(x)) \leq |x|$, we have that

$$V_\infty^0(x^+) \leq V_\infty^0(x) - \alpha_x \circ \alpha_2^{-1}(V_\infty^0(x))$$

for $x^+ \in \mathcal{X}_\infty^{\rho_2} \setminus \mathcal{X}_\infty^{\rho_1}$. By a standard construction, there exists $\sigma_V \in \mathcal{K}_\infty$ such that $\sigma_V(s) < s$ for $s > 0$ and $\sigma_V(s) \geq s - \alpha_x \circ \alpha_2^{-1}(s)$ for all $s \geq 0$ (see the proof of Theorem B.15 in Rawlings et al. (2017) for details). We thus have that either $V_\infty^0(x^+) \leq \rho_1$ or $V_\infty^0(x^+) \leq \sigma_V(V_\infty^0(x))$. Therefore, we have that

$$V_\infty^0(x^+) \leq \sigma_V(V_\infty^0(x)) \oplus \rho_1$$

This equation can be applied recursively to obtain

$$V_\infty^0(x(k)) \leq \sigma_V^k(V_\infty^0(x(0))) \oplus \rho_1$$

Finally, because $V_\infty^0(x(k)) \geq \alpha_x(|x(k)|)$ and $V_\infty^0(x(0)) \leq \alpha_2(|x(0)|)$, we have that

$$|x(k)| \leq \beta_V(|x(0)|, k) \oplus \delta$$

in which $\beta_V(s, k) := \alpha_x^{-1} \circ \sigma_V^k \circ \alpha_2(s)$ is a \mathcal{KL} function, and we recall $\rho_1 := \alpha_x(\delta)$.

Therefore the theorem is proven. \blacksquare

Remark 3.23. Theorem 3.21 is phrased in a somewhat awkward way so it is clear that $\beta_V(\cdot)$ does not depend on Δ , δ , x , or N^* . We could go further and explicitly construct δ as an \mathcal{L} function of N that does not depend on Δ so long as we chose $\Delta \geq \Delta^*$, where Δ^* is some fixed constant, but such a construction adds a considerable amount of complexity while not really adding much clarity.

We chose $V_\infty^0(\cdot)$ as a practical Lyapunov function and sublevel sets of $V_\infty^0(\cdot)$ as invariant sets because it is conceptually easier to choose a horizon length N such that a fixed set \mathcal{X}_∞^ρ is positive invariant then to adjust N such that the (shrinking) set \mathcal{X}_N^ρ is positive invariant. Since the upper and lower bounds we use for $V_N^0(\cdot)$, namely $\alpha_2(|x|)$ and $\alpha_x(|x|)$, are the same as those for $V_\infty^0(\cdot)$, the choices of ρ_1 and ρ_2 would remain the same even if we used \mathcal{X}_N^ρ as invariant sets. Grüne and Pannek (2016) use a sequence of upper bounds for each $V_N^0(\cdot)$ rather than a single one for $V_\infty^0(\cdot)$, and, while their estimate of N^* may be somewhat tighter as a result, it adds a considerable amount of complexity to the resulting proofs.

3.3 Stability with a finite horizon

Although the result of semiglobal practical stability under very general conditions is nice, it is still somewhat unsatisfactory. The semiglobal part is tolerable—we are

used to requiring a larger horizon length to for a larger basin of attraction in MPC with stabilizing terminal conditions—the practical part is rather unsatisfactory. That the most we can guarantee is attraction to a neighborhood of the origin is not just a limitation of present theory; in a later section we provide an example of a system that satisfies Assumptions 3.2 to 3.4, but for which no horizon length is sufficient to attract all states in a neighborhood of the origin to the origin.

Examining the argument used in Theorem 3.21, our result is semiglobal because $\beta_Q(s, N)$ cannot be bent below $(1/2)\alpha_x(s)$ for large s , and it provides only practical stability because it cannot be bent below $(1/2)\alpha_x(s)$ uniformly in a neighborhood of zero (in general, any fraction of $\alpha_x(s)$ strictly less than one would work). By making assumptions about the particular forms of $\alpha_x(\cdot)$ and $\beta_Q(\cdot)$, we can guarantee that these properties hold.

Suppose that we use a standard quadratic stage cost

$$\ell(x, u) = |x|_Q^2 + |u|_R^2$$

for $Q, R > 0$. We would then have $\alpha_x(s) = c_1 s^2$, in which $c_1 > 0$ is the smallest eigenvalue of Q . If the system is globally exponentially stabilizable to the origin (with exponentially small controls), Assumption 3.4 holds with $\alpha_2(s) = c_2 s^2$. The construction in Proposition 3.13 then produces $\beta_Q(s, t) = C s^2 \lambda^k$, in which $C > 0$ and $\lambda \in (0, 1)$. We can then bend $\beta_Q(\cdot)$ beneath $(1/2)\alpha_x(\cdot)$ on $\mathbb{R}_{\geq 0}$ by choosing N^* large enough such that $c_1/2 \leq C\lambda^{N^*}$. MPC could then globally exponentially stabilize the origin with a finite horizon length. Unfortunately, this argument holds only for exponentially stabilizable systems.

As it turns out, it is sufficient for $\ell(x^0(j|k; x), u^0(j|k; x))$ to converge exponentially to zero. Let

$$\ell^*(x) := \min_{u \in \mathbb{U}} \ell(x, u)$$

and note that, because of Assumption 3.3, this minimum is always attained. We then replace Assumption 3.4 with a stronger assumption.

Assumption 3.24. There exists $C > 1$ such that for all $x \in \mathbb{X}$ there exists $\mathbf{u} \in \mathbb{U}^\infty$ such that

$$V_\infty(x, \mathbf{u}) \leq C \ell^*(x).$$

Furthermore, it can be shown that the sequence $(V_N^0(x))$ converges to $V_\infty^0(x)$ in an exponential fashion. We define the residual ratio

$$r(x, N) := 1 - \frac{V_N^0(x)}{V_\infty^0(x)}$$

for all $x \in \mathbb{X}$ such that $x \neq 0$. For $x = 0$, we define $r(0, N) = 0$ for all $N \in \mathbb{N}_{\geq 0}$.

Proposition 3.25. *Under Assumptions 3.2, 3.3 and 3.24, for every $\epsilon > 0$ there exists $N \in \mathbb{N}_{\geq 0}$ such that*

$$r(x, k) \leq \epsilon$$

for all $x \in \mathbb{X}$ and $k \geq N$.

Proof. We demonstrate this proposition by showing that this system admits an exponential Q function. We define $Q(j|k; x)$ as in Proposition 3.13.

We can apply Assumption 3.24 to obtain the new upper bound

$$Q(j|k; x) \leq V_\infty^0(x^0(j|k; x)) \leq C \ell^*(x^0(j|k; x)) \quad (3.7)$$

Furthermore, because we have

$$\begin{aligned} V(j+1|k; x) - V(j|k; x) &= \sum_{i=0}^j \ell(x(i|k; x), u(i|k; x)) - \sum_{i=0}^{j-1} \ell(x^0(j|k; x), u^0(j|k; x)) \\ &= \ell(x^0(j|k; x), u^0(j|k; x)) \end{aligned}$$

we can derive a descent condition in terms of $\ell^*(x)$

$$Q(j+1|k; x) \leq Q(j|k; x) - \ell(x^0(j|k; x), u^0(j|k; x)) \leq Q(j|k; x) - \ell^*(x^0(j|k; x))$$

for all $x \in \mathbb{X}$, $k \in \mathbb{L}_{\geq 0}$, and $j \in \mathbb{L}_{0:k-1}$. We can combine this equation with (3.7) to obtain

$$Q(j+1|k; x) \leq Q(j|k; x) - \frac{1}{C} Q(j|k; x) = \left(1 - \frac{1}{C}\right) Q(j|k; x).$$

We can apply this inequality repeatedly to obtain

$$Q(k|k; x) \leq \left(1 - \frac{1}{C}\right)^k Q(0|k; x)$$

in which $Q(0|k; x) := V_{\infty}^0(x)$ for all $k \in \mathbb{L}_{\geq 0}$. Also note that

$$Q(k|k; x) := V_{\infty}^0(x) - V_k^0(x)$$

for all $k \in \mathbb{L}_{\geq 0}$. We thus have that

$$\frac{Q(k|k; x)}{V_{\infty}^0(x)} = 1 - \frac{V_k^0(x)}{V_{\infty}^0(x)} := r(x, k)$$

and therefore

$$r(x, k) \leq \left(1 - \frac{1}{C}\right)^k.$$

Because $C > 1$, we have that $(1 - 1/C) \in (0, 1)$. Because this bound on $r(\cdot)$ is monotone in k , does not depend on x , and $\lim_{k \rightarrow \infty} (1 - 1/C)^k = 0$, for any ϵ we can choose an N such that the statement of this proposition holds. \blacksquare

With this result demonstrated, we can show that the origin can be globally stabilized by MPC with a finite horizon.

Theorem 3.26. *If Assumptions 3.2, 3.3 and 3.24 hold, then there exists some N^* such that the system*

$$x^+ = f(x, u),$$

in which $u \in \mathcal{K}_N(x)$, is globally asymptotically stable for all $N \geq N^$.*

Proof. By definition of $r(x, N)$, we have that

$$V_\infty^0(x) = \frac{1}{1 - r(x, N)} V_N^0(x)$$

if $r(x, N) \neq 1$. Let $x^+ := f(x, u)$. By the principle of optimality, we have that

$$V_{N-1}^0(x^+) = V_N^0(x) - \ell(x, u) \leq V_N^0(x) - \ell^*(x).$$

Suppose we have $r(x, N - 1) \leq \bar{r}(N - 1)$ for all $x \in \mathbb{X}$. Then we have that

$$V_N^0(x^+) \leq V_\infty^0(x^+) \leq \frac{1}{1 - \bar{r}(N - 1)} V_{N-1}^0(x^+) \leq \frac{1}{1 - \bar{r}(N - 1)} (V_N^0(x) - \ell^*(x)).$$

Furthermore, we have that

$$V_N^0(x) \leq V_\infty^0(x) \leq C \ell^*(x)$$

so we can combine these equations to obtain

$$V_N^0(x^+) \leq \frac{1}{1 - \bar{r}(N - 1)} \left(V_N^0(x) - \frac{1}{C} V_N^0(x) \right) = \left(\frac{1 - \frac{1}{C}}{1 - \bar{r}(N - 1)} \right) V_N^0(x).$$

In order to guarantee a cost decrease, we need $\bar{r}(N - 1) < 1/C$. However, Proposition 3.25 says that we can choose N^* such that for all $N \geq N^*$ that inequality holds.

By Assumption 3.3, we have that

$$\alpha_x(|x|) \leq \ell^*(x) \leq V_N^0(x)$$

for all $x \in \mathbb{X}$. Furthermore, because $\ell(\cdot)$ is continuous, there exists $\bar{\alpha}_x \in \mathcal{K}_\infty$ such that

$$\ell^*(x) \leq \ell(x, 0) \leq \bar{\alpha}_x(|x|)$$

for all $x \in \mathbb{X}$. As a result, we have that

$$\begin{aligned} \alpha_x(|x|) &\leq V_N^0(x) \leq C\bar{\alpha}_x(|x|) \\ V_N^0(x) &\leq \lambda V_N^0(x) \end{aligned}$$

in which $\lambda := (1 - 1/C)/(1 - \bar{r}(N - 1)) \in (0, 1)$ because $N \geq N^*$. Thus we have that $V_N^0(x)$ is an exponential-decrease Lyapunov function. We can easily derive the bound

$$|x(k)| \leq \alpha_x^{-1}(\lambda^k C\bar{\alpha}_x(|x|))$$

and thus the origin is asymptotically stable under the control law $\mathcal{K}_N(\cdot)$. ■

So if the system can be globally exponentially stabilizable with respect to the stage cost $\ell(\cdot)$, then we know that, with a sufficiently large horizon, MPC can stabilize the origin without terminal constraints. Intermediate results between Theorem 3.21 and Theorem 3.26 are possible. In particular, we can get semiglobal stability if the system is locally exponentially stabilizable with respect to $\ell(\cdot)$, i.e., Assumption 3.24 is satisfied not on all \mathbb{X} , but on compact sets containing the origin in their interior (Grüne and Pannek, 2016, Theorem 6.35). It is also possible to show global practical stability if a bound like Assumption 3.24 is possible on sets excluding compact regions containing the origin in their interior.

An obvious question to address, then, is when does a system admit a stage cost such that Assumption 3.24 holds either globally or locally. Grune and Rantzer (2008, Remark 4.8) suggest that a stage cost satisfying Assumption 3.24 can always be constructed from a global CLF. It is not obvious to me, however, how to perform such a construction from a CLF with a \mathcal{K}_∞ descent condition, as is constructed in Theorem 3.16. For an exponential-decrease CLF, such as the one constructed in (Kellett and Teel, 2004a), constructing such a stage cost is straightforward.

That observation at least settles the mathematical question of existence of such a stage cost. Of course, if we had a global CLF, then we could use it as a stabilizing terminal cost for MPC without any terminal constraint, or just use the control law it induces directly. Practical considerations in choosing a stage cost are covered in (Grüne and Pannek, 2016, Section 6.6). We next turn our attention to the conditions for which Assumption 3.24 holds locally.

3.4 Quadratic stage costs and stabilizability of the linearized system

One important class of stage costs are those for which $\ell(\cdot)$ is a quadratic function. Least squares objectives are often engineers' first choice when confronted with a tracking problem, and the linear quadratic regulator (LQR) has been long a subject of study in control theory. For nonlinear systems, a local controller can often be synthesized by linearization. When stabilizing terminal conditions are used in MPC, a terminal cost derived from the LQR cost of the linearized system is often

used, when possible (De Nicolao et al., 1996; Chen and Allgöwer, 1998). Finally, a quadratic stage cost satisfies Assumption 3.24 if and only if the origin is exponentially stabilizable (in $|x|$, not $\ell^*(x)$).

For an initial result, we derive a sufficient condition for Assumption 3.24 to be fulfilled semiglobally with a quadratic stage cost.

Proposition 3.27. *Suppose that Assumptions 3.2, 3.3 and 3.24 are satisfied, and, furthermore, we have that $\ell(x, u) = |x|_Q^2 + |x|_R^2$ in which $Q, R > 0$. Also suppose that $f(\cdot)$ is twice continuously differentiable, the linearization of $f(\cdot)$ at the origin (A, B) is stabilizable, and $\mathbb{X} \times \mathbb{U}$ contains the origin in its interior. Then, for every $\rho \geq 0$ there exists $C > 0$ such that*

$$V_\infty^0(x) \leq C |x|^2$$

for all $x \in \mathcal{X}_\infty^\rho$.

Proof. It is well-known (see, e.g., (Rawlings et al., 2017, Sec. 2.5.5)) that under these conditions a controller $u = Kx$ synthesized for the linear system (A, B) locally stabilizes the nonlinear system. Furthermore, there exists $P_f > 0$ such that if we let $V_f(x) := |x|_{P_f}^2$, we have that

$$V_f(f(x, Kx)) \leq V_f(x) - \ell(x, Kx)$$

for all $x \in \{x \mid V_f(x) \leq \alpha\} := \mathcal{X}_f$ for some $\alpha > 0$. Such a terminal cost and terminal region would constitute stabilizing terminal conditions for MPC. By (Grüne and Pannek, 2016, Eq. 5.20, Thm. 5.13), we have that $V_\infty^0(x) \leq V_f(x)$ for $x \in \mathcal{X}_f$. As a result, we have that $\mathcal{X}_f \subseteq \mathcal{X}_\infty^\alpha$, and because $V_f(x) > \alpha$ for $x \notin \mathcal{X}_f$ we have that $V_\infty^0(x) \leq V_f(x)$ for $x \in \mathcal{X}_\infty^\alpha$. Let

$$\tilde{C} := \max_{x \in \mathbb{X}} V_\infty^0(x)/|x|^2 \quad \text{such that } \alpha \leq V_\infty^0(x) \leq \rho$$

and note that because $V_\infty^0(x)$ and $|x|^2$ are continuous and strictly positive in this set, the maximization is well-defined. Let $C := \max(\tilde{C}, \bar{\lambda}(P_f))$, in which $\bar{\lambda}(P_f)$ is the largest eigenvalue of P_f . We have that $V_\infty^0(x) \leq \bar{\lambda}(P_f) |x|^2$ for $x \in \mathcal{X}_\infty^\alpha$ and that $V_\infty^0(x) \leq \tilde{C} |x|^2$ for $x \in \mathcal{X}_\infty^\rho \setminus \mathcal{X}_\infty^\alpha$. Thus $V_\infty^0(x) \leq C |x|^2$ for all $x \in \mathcal{X}_\infty^\rho$. ■

As it turns out, however, stabilizability of the linearized system is not only sufficient for Assumption 3.24 to hold, but also necessary, under appropriate conditions. To facilitate this proof, as well as a later example, we require a preliminary result about smoothness of (finite-horizon) optimal cost functions. Here we include a terminal cost in the formulation of $V_N^0(\cdot)$ because it is useful for further results.

Proposition 3.28. *Let $f : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{X}$ and $V_f : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ be smooth (i.e., infinitely differentiable), let $\ell(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$ with $Q, R > 0$, let there not exist any terminal constraint, and let the origin lie in the interior of $\mathbb{X} \times \mathbb{U}$. Then for every $N \geq 0$ there exists a neighborhood of the origin in which both the MPC optimal value function*

$$V_N^0(x) := \min_{\mu \in \mathbb{U}^N} \sum_{k=0}^{N-1} \ell(\chi(k), \mu(k)) + V_f(\chi(N)) \quad \chi^+ = f(\chi, \mu) \quad \chi(0) = x$$

is smooth and the optimal control law $\mathcal{K}_N^0(x)$ is both smooth and single-valued.

Proof. We proceed with a proof by induction. By definition, $V_0^0(\cdot) = V_f(\cdot)$, which is smooth everywhere, which gives us the base case. Now assume that $V_{N-1}^0(\cdot)$ is smooth in a neighborhood of the origin. Let

$$\check{V}_N(x, u) := \ell(x, u) + V_{N-1}^0(f(x, u))$$

which we distinguish from $V_N(x, \mathbf{u})$ because $V_N(\cdot)$ is a function of a control sequence, whereas $\check{V}_N(\cdot)$ is a function of only the first control action. By a dynamic programming argument, we can show that

$$V_N^0(x) = \min_{u \in \mathbb{U}} \check{V}_N(x, u)$$

Because, by assumption, $V_{N-1}^0(\cdot)$ is smooth in a neighborhood of the origin, there exists some $\rho > 0$ such that

$$V_{N-1}^0(x) \leq \rho |x|^2$$

for all x in that neighborhood of the origin. Because $f(\cdot)$ is smooth, it is locally Lipschitz, and there thus exists $L > 0$ such that

$$|f(x, u)| \leq L |(x, u)|$$

for all (x, u) in a neighborhood of the origin. Any optimal control action $u_N^0(x)$ must perform no worse than idle control. Therefore

$$(1/2) |u_N^0(x)|_R^2 \leq V_N^0(x) \leq (1/2) |x|_Q^2 + V_{N-1}^0(f(x, 0)) \leq (1/2) |x|_Q^2 + \rho L^2 |x|^2$$

and thus

$$|u_N^0(x)| \leq \sqrt{(2\rho L^2 + \bar{\lambda}(Q))/(\underline{\lambda}(R))} |x| \quad (3.8)$$

in a neighborhood of the origin regardless of whether $f(\cdot)$ is stabilizable with small controls. Because the origin is in the interior of \mathbb{U} , there exists a neighborhood of the origin for which input constraints are not active.

Where the appropriate derivatives are defined, we have that

$$\frac{\partial \check{V}_N}{\partial u} = u'R + \frac{dV_{N-1}^0}{dx} \Big|_{f(x,u)} \frac{\partial f}{\partial u} \quad (3.9)$$

and

$$\frac{\partial^2 \check{V}_N}{\partial u^2} = R + \left(\frac{\partial f}{\partial u} \right)' \frac{d^2 V_{N-1}^0}{dx^2} \Big|_{f(x,u)} \frac{\partial f}{\partial u} + \frac{dV_{N-1}^0}{dx} \Big|_{f(x,u)} \frac{\partial^2 f}{\partial u^2}, \quad (3.10)$$

in which $\partial^2 f / \partial u^2$ is a $n \times m \times m$ order three tensor being contracted along its n dimension by dV_{N-1}^0 / dx . Clearly we have that $V_N^0(0) = 0$ and $u^0(0) = 0$. Furthermore, by the necessary conditions for an unconstrained minimum of a smooth function, we have that

$$\frac{dV_{N-1}^0}{dx} \Big|_0 = 0 \quad \text{and} \quad \frac{d^2 V_{N-1}^0}{dx^2} \Big|_0 \geq 0.$$

The Hessian (3.10) therefore has its minimum singular value bounded away from zero in a neighborhood of the origin. Because, by assumption, the origin is in the interior of $\mathbb{X} \times \mathbb{U}$, we can apply the implicit function theorem (Dontchev and Rockafellar, 2014, pp.20-22, Theorem 1B.1 and Proposition 1B.5), to obtain a smooth function $\mathcal{K}_N^0(x)$ setting (3.9) equal to zero in a neighborhood of the origin. Furthermore, this localization is unique in that neighborhood of the origin—there exists no $u_N^0(x)$ setting (3.9) to zero other than $\mathcal{K}_N^0(x)$. Therefore, because of (3.8), we have that $\mathcal{K}_N^0(x)$ is the optimal control law as suggested by our notation.

Thus, we have that

$$V_N^0(x) = \check{V}_{N-1}(x, \mathcal{K}_N^0(x))$$

which, as a composition of (locally) smooth functions, is smooth on a neighborhood of the origin. Thus the inductive case is complete. ■

Remark 3.29. Note, while this proposition states that for every $N \geq 0$ there exists a neighborhood of the origin for which $V_N^0(\cdot)$ is smooth, it *does not* state there exists a neighborhood of the origin for which $V_N^0(\cdot)$ is smooth for all $N \geq 0$, i.e., the neighborhoods of smoothness may shrink as N increases.

Although one would think that this proposition could be extended to $f(\cdot)$ and $V_f(\cdot)$ twice continuously differentiable, frustratingly that appears to not be the case. In that case, the implicit function theorem states only that $\mathcal{K}_N^0(\cdot)$ is once differentiable in a neighborhood of the origin, and a second derivative is required to evaluate $d^2V_N^0/dx^2$. Simply requiring, e.g., third derivatives from $f(\cdot)$ and $V_f(\cdot)$ is not sufficient, because we end up with one fewer derivative for $V_N^0(\cdot)$ than we had for $V_{N-1}^0(\cdot)$. One glimmer of hope remains, because a second derivative of $\mathcal{K}_N^0(\cdot)$ is not required to evaluate the second derivative of $V_N^0(\cdot)$ at the origin—it is multiplied by dV_{N-1}^0/dx and, as a result, is canceled out. However, the second derivative must exist in a neighborhood of the origin in order to apply the implicit function theorem, which leaves us with no way to proceed.

Now, we can present a proof that stabilizability of the linearized system is not simply sufficient, but also necessary for a local version of Assumption 3.24 to hold.

Proposition 3.30. *Suppose that $f(\cdot)$ is smooth, $V_f = 0$, $\ell(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$ with $Q, R > 0$, let there not exist any terminal constraint, i.e., $\mathbb{X}_f = \mathbb{X}$, let the origin lie in the interior of $\mathbb{X} \times \mathbb{U}$, and let $V_\infty^0(x) \leq |x|_{P_\infty}^2$ for some $P_\infty > 0$ and all $x \in \mathbb{X}$. Then the linearized system at the origin is stabilizable.*

Proof. Under the conditions for this theorem we can apply Proposition 3.28 to

obtain second derivatives for $V_N^0(\cdot)$ at the origin for all $N \in \mathbb{R}_{\geq 0}$. We begin with the observation

$$V_N^0(x) = \check{V}_{N-1}(x, \mathcal{K}_N^0(x))$$

which allows us to apply the chain rule to calculate $d^2V_N^0/dx^2$. Let $z_N^0 := (x, \mathcal{K}_N^0(x))$

and let $x_N^\oplus := f(z_N^0)$. We have that

$$\frac{dV_N^0}{dx} = x'Q + \mathcal{K}_N^0(x)'R \frac{d\mathcal{K}_N^0}{dx} + \frac{dV_{N-1}^0}{dx} \Big|_{x_N^\oplus} \frac{\partial f}{\partial u} \Big|_{z_N^0} \frac{d\mathcal{K}_N^0}{dx}$$

and that

$$\begin{aligned} \frac{d^2V_N^0}{dx^2} = & Q + \mathcal{K}_N^0(x)'R \frac{d^2\mathcal{K}_N^0}{dx^2} + \frac{dV_{N-1}^0}{dx} \Big|_{x_N^\oplus} \frac{\partial f}{\partial u} \Big|_{z_N^0} \frac{d^2\mathcal{K}_N^0}{dx^2} \\ & + \left(\frac{d\mathcal{K}_N^0}{dx} \right)' \left(R + \left(\frac{\partial f}{\partial u} \Big|_{z_N^0} \right)' \frac{d^2V_{N-1}^0}{dx^2} \Big|_{x_N^\oplus} \frac{\partial f}{\partial u} \Big|_{z_N^0} + \frac{dV_{N-1}^0}{dx} \Big|_{x_N^\oplus} \frac{\partial^2 f}{\partial u^2} \Big|_{z_N^0} \right) \frac{d\mathcal{K}_N^0}{dx}, \end{aligned}$$

in which the third-order tensors $\partial^2 f / \partial u^2$ and $d^2\mathcal{K}_N^0/dx^2$ are being contracted along their appropriate dimensions (if $m \neq n$, there is only one dimension for which those contractions make sense). Fortunately, those tensors need not trouble us for long. We have that $\mathcal{K}_N^0(0) = 0$, $x_N^\oplus = 0$, and $dV_{N-1}^0/dx|_0 = 0$, and therefore the terms involving the third-order tensors vanish when $d^2V_N^0/dx^2$ is evaluated at the origin. Let

$$\frac{\partial f}{\partial x} \Big|_0 := A \quad \frac{\partial f}{\partial u} \Big|_0 := B \quad \frac{d\mathcal{K}_N^0}{dx} \Big|_0 := K_N \quad \Pi_N := \frac{d^2V_N^0}{dx^2} \Big|_0.$$

We then have that

$$\Pi_N = Q + K_N'(R + B'P_{N-1}B)K_N \tag{3.11}$$

We would like an explicit expression for K_N . Fortunately the implicit function theorem gives us an explicit expression for $d\mathcal{K}_N^0/dx$ in a neighborhood of the origin:

$$\begin{aligned}
\frac{d\mathcal{K}_N^0}{dx} &= \left(\frac{\partial^2 \check{V}_N}{\partial u^2} \Big|_{z_N^0} \right)^{-1} \frac{\partial^2 \check{V}_N}{\partial u \partial x} \Big|_{z_N^0} \\
&= \left(R + \left(\frac{\partial f}{\partial u} \Big|_{z_N^0} \right)' \frac{d^2 V_{N-1}^0}{dx^2} \Big|_{x_N^\oplus} \frac{\partial f}{\partial u} \Big|_{z_N^0} + \frac{dV_{N-1}^0}{dx} \Big|_{x_N^\oplus} \frac{\partial^2 f}{\partial u^2} \Big|_{z_N^0} \right)^{-1} \\
&\quad \times \left(\left(\frac{\partial f}{\partial u} \Big|_{z_N^0} \right)' \frac{d^2 V_{N-1}^0}{dx^2} \Big|_{x_N^\oplus} \frac{\partial f}{\partial x} \Big|_{z_N^0} + \frac{dV_{N-1}^0}{dx} \Big|_{x_N^\oplus} \frac{\partial^2 f}{\partial u \partial x} \Big|_{z_N^0} \right).
\end{aligned}$$

Once again, when this expression is evaluated at the origin, the third-order tensor terms vanish and we are left with

$$K_N = (R + B' \Pi_{N-1} B)^{-1} B' \Pi_{N-1} A$$

which is clearly recognizable as the expression for the gain matrix from the LQR (note—the expression we have here differs from that in (Rawlings et al., 2017, Sec. 1.3) because here we iterate forward by increasing the horizon length N , while there they iterate backwards from a predetermined horizon length N). We can substitute this equation into (3.11) in order to obtain

$$\Pi_N = Q + A' \Pi_{N-1} B (R + B' \Pi_{N-1} B)^{-1} B' \Pi_{N-1} A$$

which is clearly recognizable as the Riccati iteration for the LQR. We have that $\{\Pi_N\}$ constitutes a nondecreasing sequence in the space of symmetric positive semidefinite matrices, and it is bounded above by P_∞ because $V_\infty^0(x) \geq V_N^0(x)$ for all $x \in \mathbb{X}$ and $N \geq 0$. Thus it converges to some symmetric positive definite Π_∞ (which, note, may be distinct from P_∞) which serves as a Lyapunov function for the system $x^+ = (A + BK_\infty)x$. Because K_∞ serves as a stabilizing gain for the linearized system, the system (A, B) must be stabilizable. \blacksquare

Remark 3.31. This result probably also holds for $f(\cdot)$ twice continuously differentiable. A possible proof would proceed as that in (Rawlings et al., 2017, Sec. 2.5.5) in reverse. Rather than showing that the linearized controller is locally stabilizing for the nonlinear system, it would show that the *nonlinear* controller is locally stabilizing for the *linearized* system. The direct connection to the Riccati equation is lost, however.

3.5 Practical stability cannot be avoided

As discussed before, the “practical” aspect of the stability guaranteed by Theorem 3.21 is not a limitation of MPC theory, but represents a real phenomenon in control. Here we provide several examples for which no horizon length N is sufficient to guarantee (even semiglobal) asymptotic stability of the origin.

Example: Mixed-power stage cost

This example is adapted from (Grüne and Pannek, 2016, Example 6.38). Consider the scalar linear integrator

$$x^+ = x + u$$

that we attempt to control to the origin using the stage cost

$$\ell(x, u) = x^2 + |u|$$

without stabilizing terminal conditions. ℓ^1 penalties are known to result in either aggressive or idle control (see, e.g., (Rawlings et al., 2017, Ex. 2.19)), and this case

is no different. Obviously, for any optimal sequence, we have that $\text{sgn}(u^0(j|N)) = -\text{sgn}(x)$ for all j, N and

$$\left| \sum_{j=0}^{N-1} u^0(j|N) \right| \leq |x|$$

because there is no point in overshooting the origin for this simple system. Suppose that, for an optimal solution of $\mathbb{P}_N(x)$, we have $u^0(j|N) > 0$ for $j > 0$. Then we could define a new sequence

$$\tilde{\mathbf{u}} = \left(\sum_{j=0}^{N-1} u^0(j|N), 0, \dots \right)$$

that results in the same energy spent on control but results in less tracking error.

Therefore we have that

$$V_N^0(x) = x^2 + |u^0(0|N)| + (N-1)(x + u^0(0|N))^2$$

It can be shown (by, for example, using subdifferential calculus) that the optimal control law that results is

$$\mathcal{K}_N^0(x) = \begin{cases} -1/(2(N-1)) - x & x < -1/(2(N-1)) \\ 0 & x \in [-1/(2(N-1)), 1/(2(N-1))] \\ 1/(2(N-1)) - x & x > 1/(2(N-1)) \end{cases}$$

for $N \geq 2$. The resulting optimal controller brings x into the interval $[-1/(2(N-1)), 1/(2(N-1))]$ where idle control results. ■

This example shows that problems can occur even with linear systems if the stage cost is not well-chosen. It is relatively easy to induce bad behavior with an ℓ^1 penalty in the stage cost. However, this behavior can occur even with quadratic costs for nonlinear systems. Müller and Worthmann (2017) demonstrate, in continuous

time, that no horizon length suffices to stabilize the nonholonomic robot without terminal conditions when a quadratic stage cost is used. In a similar fashion, it can be shown that the discrete cubic integrator (see, e.g., Meadows et al. (1995); Allan and Rawlings (2018a)), cannot be stabilized without terminal conditions when a quadratic stage cost is used.

One feature all these examples share is that the systems integrate some function of the input. One might reasonably think that having the system be an integrator is an essential ingredient for this type of bad behavior. Obviously, for a locally asymptotically stable steady state, steering the system to a neighborhood of the origin is sufficient for asymptotic stability. One might expect, however, the state to be controlled to a (neutrally) stable subspace in a neighborhood of the origin. We thus have the following conjecture:

Conjecture 3.32. *Suppose there exists a neighborhood of the origin that does not contain any unforced orbits or steady states besides the origin. Then Theorem 3.21 guarantees semiglobal asymptotic stability.*

This conjecture is false. The next example shows that finite-horizon MPC with a quadratic cost function but without terminal conditions can spend infinite control energy to keep the system's state in a neighborhood of the origin without ever controlling it to the origin. However, because $V_\infty^0(\cdot)$ is bounded above appropriately, $\mathcal{K}_\infty^0(\cdot)$ does stabilize the origin.

Example: Practical stability with quadratic cost

Consider the scalar system (inspired by (Grimm et al., 2004, Example 2))

$$x^+ = x(x + 1)(1 - u)$$

defined with a state-space $\mathbb{X} = \mathbb{R}$ and input set $\mathbb{U} = \mathbb{R}$. Note that this system is open-loop unstable, but the control law $u = \min(|x|, 1)$ stabilizes the origin with small controls. We attempt to stabilize this system with MPC without terminal constraints. We use the stage cost

$$\ell(x, u) = x^2 + u^2$$

and a terminal cost

$$V_f(x) = \rho x^2$$

for some $\rho \geq 0$. We can write down the OCP for horizon length $N = 1$

$$\min_{u \in [0,1]} V_1(x, u) = x^2 + u^2 + \rho(x(x + 1)(1 - u))^2.$$

Thus, we can set the gradient of $V_1(\cdot)$ to zero in order to find $\mathcal{K}_1(\cdot)$.

$$\begin{aligned} \frac{\partial V_1}{\partial u} &= 2u + 2\rho(x(x + 1)(1 - u)) \cdot -(x(x + 1)) \\ &= 2u - 2\rho(1 - u)(x(x + 1))^2. \end{aligned}$$

By setting this expression to zero and solving for u , we obtain

$$u = \frac{\rho x^2(x + 1)^2}{1 + \rho x^2(x + 1)^2} := \mathcal{K}_1(x, \rho).$$

The resulting closed-loop system has the evolution equation

$$f(x, \mathcal{K}_1(x, \rho)) = \frac{x(x + 1)}{1 + \rho x^2(x + 1)^2} := f_{cl}(x, \rho).$$

Because the closed loop system is smooth, we can attempt to determine whether it is stable using linear stability analysis. Taking a first derivative, we obtain

$$\frac{df_{cl}}{dx} = \frac{(2x + 1)(1 - \rho x^2(x + 1)^2)}{(1 + \rho x^2(x + 1)^2)^2},$$

which, evaluated at zero, results in a value of one regardless of the value of ρ . Therefore linear stability analysis is inconclusive. Because this is a univariate system, however, we can take further derivatives to evaluate its properties. For the second derivative, we obtain

$$\frac{d^2 f_{cl}}{dx^2} = \frac{2(1 - \rho x(x + 1)(5x^2 + 5x + 1))}{(1 + \rho x^2(x + 1)^2)^2} - \frac{4\rho x(x + 1)(2x + 1)^2(1 - \rho x^2(x + 1)^2)}{(1 + \rho x^2(x + 1)^2)^3},$$

which, evaluated at zero, results in a value of two, again, regardless of the value of ρ . Therefore, near $x = 0$, we have that

$$f_{cl}(x, \rho) \approx x + x^2$$

regardless of the value of ρ . As a result, $x = 0$ is a semistable equilibrium (i.e., it attracts sufficiently close $x < 0$ but does not attract any $x > 0$) and no quadratic terminal penalty results in asymptotic stability.

In principle, we could use $V_1^0(\cdot)$ as a terminal penalty in order to determine $\mathcal{K}_2(\cdot)$ and $V_2^0(\cdot)$, and solve the problem with dynamic programming. However, the calculations are messy and probably, for some horizon N , no longer admit a closed-form solution. Even without determining these control laws explicitly, though, we can infer enough about them to demonstrate that the origin is not a stable equilibrium for *all* finite horizon lengths.

First, because of the one dimensional state space and the way in which the control enters the problem, we have a simple characterization of the aggressiveness

of a control law—a control law $\mathcal{K}_1(\cdot)$ is more aggressive than $\mathcal{K}_2(\cdot)$ at x if $\mathcal{K}_1(x) \geq \mathcal{K}_2(x)$. Next, we note that, for two terminal costs $V_{f,1}(\cdot)$ and $V_{f,2}(\cdot)$ such that they are both strictly increasing and $V_{f,1}(x) \geq V_{f,2}(x)$ for all x , we have that $\mathcal{K}_1(x) \geq \mathcal{K}_2(x)$ for all x .

By Proposition 3.28, there is a neighborhood of the origin for which $V_N^0(\cdot)$ is smooth. Therefore, it admits such an upper bound in this neighborhood. Furthermore, we note that for large x the control law $u = 1$ controls the system to the origin for a cost less than $x^2 + 1$, and thus $V^0(x) \leq x^2 + 1$. Therefore, this quadratic upper bound can be extended to all of \mathbb{R} . Thus, for some $\rho > 0$, $\mathcal{K}_1(x, \rho)$ is more aggressive than $\mathcal{K}_N(x, 1)$. Because no $\mathcal{K}_1(x, \rho)$ stabilizes the origin, $\mathcal{K}_N(x, 1)$ does not stabilize the origin for any N .

Although the control policy $\mathcal{K}_f(x) = \min(|x|, 1)$ does stabilize the origin, it does *not* attain a finite infinite-horizon cost for $x \in (0, 1)$ because (as one can show using a variant of the Cauchy condensation test) the closed-loop system does not converge fast enough to the origin. One can show, however, that an even more aggressive control law, $\mathcal{K}_f(x) = \min(|x|^{1/2}, 1)$, attains a finite cost, because the control Lyapunov function $V_f(x) = 4|x|^{1/2}$ is compatible with it and $\ell(\cdot)$. We can extend this CLF to all of \mathbb{R} with an additional x^2 term, $V_f(x) = 4|x|^{1/2} + x^2$, which then constitutes a global upper bound for $V_\infty^0(x)$. Therefore, the infinite-horizon optimal control law is well-defined and globally stabilizes the origin.

3.6 Three methodologies, one winner

We have gone into some depth about setpoint regulation using MPC without stabilizing terminal conditions. After this discussion, a comparison to the other two main design methodologies, use of a terminal equality constraint and a terminal region, is in order. Rawlings et al. (2017) do not compare controllers synthesized with stabilizing terminal conditions to those that do without them in that work, but Mayne (2013) provides a comprehensive defense of stabilizing terminal conditions. On the other hand, Grüne and Pannek (2016) provide a comparison between these two design philosophies in their Section 7.4. The results derived here can supplement these discussions.

3.6.1 Terminal equality constraint

Minimum energy control with a terminal equality constraint was introduced by Kleinman (1970) in order to construct, by use of a definite matrix integral, a stabilizing gain for a linear system that could be used as a starting point for Newton methods of solving the matrix Riccati equation for the LQR. An input weighting matrix was introduced by Thomas (1975) and state error term was introduced by Kwon and Pearson (1978), who proposed using this form of control for time-varying systems, for which there is no convenient way to solve for the LQR. Kwon et al. (1983) articulate an early form of the trilemma we attempt to solve now. For stability of an optimal controller for a time-varying linear system, one can either use a terminal equality constraint to approximate the infinite-horizon cost from above,

no terminal penalty and approximate it from below, as we did here, or use a (global, in their case) control Lyapunov function as a terminal penalty.

The use of a terminal equality constraint in nonlinear receding horizon control appears to have been independently proposed three times. In an early, unfortunately neglected paper, Chen and Shaw (1982) proposed use of such a constraint for continuous time systems, Keerthi and Gilbert (1988) proposed use of one for discrete time systems, and Mayne and Michalska (1988) proposed use of one for continuous time systems. Meadows and Rawlings (1993) were the first to clearly articulate why using a terminal equality constraint is sufficient for stability—it chooses a (typically suboptimal) infinite horizon sequence of inputs that attains a finite cost, namely, one that steers the system to the origin in N moves and remains zero thereafter.

Traditionally, Lyapunov functions were required to be continuous, and, in continuous time, differentiable. This requirement bedeviled these early researchers, because terminal equality constraints can result in discontinuous optimal cost functions. Meadows et al. (1995) demonstrated that continuity of the Lyapunov function is unnecessary for nominal stability so long as it admits a \mathcal{K}_∞ function upper bound (which, in turn, implies that it is continuous at the origin). However, as pointed out by Grimm et al. (2004), terminal equality constraints can result in closed-loop systems that are nominally stable but have zero robustness.

In this formulation of MPC, the problem $\mathbb{P}_N(x)$ is augmented with an additional constraint

$$\chi(N) = 0 \quad \text{or} \quad \phi(N; x, \boldsymbol{\mu}) = 0$$

The resulting MPC controller stabilizes the origin with any horizon length for

which $\mathbb{P}_N(x)$ is feasible. There are many disadvantages to this approach, however. First, $\mathbb{P}_N(x)$ may not be feasible, and the resulting controller may be undefined. Second, requiring controllability to the origin is a much stronger assumption than stabilizability to the origin, because it excludes, e.g., linear systems with uncontrollable (but stable) modes. Third, as pointed out by Grimm et al. (2004), the resulting closed-loop system may not be robustly stable. Finally, the addition of another constraint to the OCP, especially for nonlinear programs (NLPs), may make it harder to solve numerically.

The last point merits further discussion. For gradient-based local solvers (such as IPOPT), constraint qualifications are important to ensuring good behavior. The most popular and easiest to understand is the linear independence constraint qualification (LICQ). When applied to this problem, it requires the matrix $\partial\phi_N/\partial\mathbf{u}$ to be full row rank at the optimal solution (suppose for simplicity that the origin is in the interior of $\mathbb{X} \times \mathbb{U}$). Because of its special structure, we can obtain a recursive formula for this matrix using the chain rule.

$$\frac{\partial\phi_N}{\partial\mathbf{u}}\Big|_{x,\mathbf{u}} = \left[\frac{\partial f}{\partial x}\Big|_{x^{(N-1)}, u^{(N-1)}} \times \frac{\partial\phi_{N-1}}{\partial\mathbf{u}}\Big|_{x,\mathbf{u}}, \quad \frac{\partial f}{\partial u}\Big|_{x^{(N-1)}, u^{(N-1)}} \right]$$

The most important place for LICQ is the origin. If we expect the MPC regulator to stabilize the origin with small controls, then as x approaches zero u also approaches zero. If we evaluate $\partial\phi_N/\partial\mathbf{u}$ there and define $\partial f/\partial x|_0 := A$ and $\partial f/\partial u|_0 := B$, we have that

$$\frac{\partial\phi_N}{\partial\mathbf{u}}\Big|_{0,0} = [A^{N-1}B, A^{N-2}B, \dots, B]$$

which is a rearrangement of the controllability matrix for the system linearized at the origin. Therefore, $\mathbb{P}_N(\cdot)$ satisfies LICQ at the origin if and only if the linearized system is controllable with index of controllability less than N . Although a small perturbation in x and \mathbf{u} may result in the rank condition being satisfied, the optimal control problem becomes increasingly ill-conditioned as the controller drives x to the origin. If we take LICQ as a requirement⁶ for good behavior of gradient-based nonlinear solvers, then having (A, B) controllable is also a requirement. Fortunately, satisfaction of LICQ ensures a degree of robustness in the resulting controller (Yang et al., 2015, Theorem 14). However, if (A, B) is controllable, the more powerful design methodology of a terminal region can be used.

3.6.2 Terminal cost and terminal region

The terminal cost/terminal region formulation of MPC has its origins in the dual mode MPC proposed by Michalska and Mayne (1993). In that paper, the terminal equality constraint is replaced with a terminal region constraint in order to confer robustness to numerical roundoff error. In the terminal region, the MPC controller is turned off, replaced with a controller synthesized for the linearization of the system at the origin. Two related ideas were proposed in Rawlings and Muske (1993) and Alamir and Bornard (1995). In the former, the controller is required to terminate in the stable subspace of a linear system, where the infinite-horizon cost

⁶The Mangasarian-Fromovitz constraint qualification (MFCQ), a more general constraint qualification, reduces to LICQ in the case of equality constraints. The constant rank constraint qualification (CRCQ) may permit the use of a rank-deficient A when deadbeat modes exist. A detailed treatment of these constraint qualifications in the context of optimal control can be found in Yang et al. (2015).

of idle control is used as a terminal penalty. In the latter, this concept is extended to nonlinear MPC. Alamir and Bornard (1995) also provides an interesting precedent for MPC without stabilizing terminal conditions. Because there is typically no easy representation of a “stable subspace” for nonlinear systems, and calculating the infinite-horizon cost of idle control may be difficult if there were, they use a longer prediction horizon than control horizon, but prove stability for a large, but finite, prediction horizon. These ideas were finally synthesized into their modern form in De Nicolao et al. (1996) and Chen and Allgöwer (1998). Mayne et al. (2000) abstract away implementation details in those works in order to articulate the essentials for theory: a terminal region \mathbb{X}_f , a terminal control law $\mathcal{K}_f(\cdot)$ under which the terminal region is positive invariant, and a terminal cost function $V_f(\cdot)$ which is a control Lyapunov function that achieves a decrease in the stage cost $\ell(\cdot)$.

There are two conventional candidates for terminal regions. For linear quadratic MPC, the natural candidate is the region for which the solution of the LQR is feasible. This choice of terminal region leads to the performance of the infinite-horizon controller where $\mathbb{P}_N(\cdot)$ is feasible (see (Rawlings et al., 2017, Sec 2.5.4) for details). For nonlinear MPC, the conventional choice has been to linearize the system and use the LQR solution as a terminal control law and some sublevel set of a modified LQR cost function as a terminal region (see (Rawlings et al., 2017, Sec 2.5.5) for details).

This methodology has numerous advantages over the terminal equality constraint. As with the terminal equality constraint, whenever $\mathbb{P}_N(\cdot)$ is feasible, the controller is stabilizing. However, satisfaction of inclusion in a region is an easier constraint for a NLP solver to satisfy than termination at a point. Furthermore,

in the absence of hard state constraints, if a sublevel set of a control Lyapunov function is used as the terminal region, such a controller is inherently robust (Yu et al., 2014; Allan et al., 2017).

3.6.3 Comparison to MPC without terminal conditions

Here we assume that we are using MPC with a nonlinear model with state and input dimensions sufficiently large that the resulting OCP cannot be solved with a global nonlinear solver like BARON. For this class of problems, gradient-based solvers like IPOPT are currently the most competitive at solving MPC OCPs. Requiring the system's linearization to be controllable at the origin is a requirement in order for a terminal equality constraint to behave well with this class of solvers. However, if we have a stabilizable linearization, we can design a terminal region using well-established techniques.⁷ So the remaining decision is whether or not a terminal region is worthwhile.

As we have seen, when a terminal region that is a sublevel set of a CLF is used, the resulting MPC is inherently robust. Because, in the absence of state and terminal constraints, $V_N^0(\cdot)$ is continuous, so is MPC without stabilizing terminal conditions.

The conditions in which we can easily design a terminal region (a stabilizable

⁷The guide to terminal region design in (Rawlings et al., 2017, Sec. 2.4.5) is not quite constructive. In particular, no method is given to shrink the terminal region such that the terminal control law always satisfies input constraints (and state constraints, but we exclude those here), and an upper bound for a second-derivative term is required on a neighborhood of the origin. For a set of linear inequalities, the first problem is tractable. It is much less obvious how to get an upper bound for this second derivative term without gridding over the state space, a strategy that is not infallible and suffers from the curse of dimensionality. I was unaware of this gap when I began this work, and do not presently have the time to do literature review to see if it has been filled.

linearization) are the same as those for which we can expect to avoid practical stability for MPC without stabilizing terminal conditions (Proposition 3.30), at least when the origin is in the interior of \mathbb{U} (this requirement can probably be relaxed if a different proof technique is used). In the event that we cannot design a terminal control law and terminal region designed using the LQR, we (probably) cannot stabilize the origin with a quadratic cost. In that case, the local design problem of a terminal control law and terminal region seems easier to solve than the global design problem of a (locally) exponentially-decreasing stage cost.

One advantage of excluding stabilizing terminal constraints is that $\mathcal{K}_N^0(x)$ is defined for all $x \in \mathbb{X}$, while if a terminal region is included, there are almost always states for which $\mathcal{K}_N^0(x)$ is not defined. However, when one uses a terminal region, one has the assurance that, so long as $\mathbb{P}_N(x)$ is feasible, the resulting controller is stabilizing, while that is not the case when a terminal region is excluded. In my view, the decisive point in favor of including a terminal region is an issue that has not been mentioned yet. A tacit assumption in the stability results for MPC without stabilizing terminal conditions is that a *global* solution to $\mathbb{P}_N(\cdot)$ is obtained. Gradient-based solvers, on the other hand, can guarantee only that a *local* solution is obtained. MPC using local solutions is a special case of suboptimal MPC (see Scokaert et al. (1999)).

Suboptimal MPC has been studied in contexts both with (Allan et al., 2017) and without (Grüne and Pannek, 2010) terminal constraints. When the terminal region is included, there exists a clear criteria about when a suboptimal solution is adequate—it satisfies the terminal constraint. When it is absent, however, the picture is much

murkier. Grüne and Pannek propose requiring that the newly generated solution satisfy a sufficient decrease condition from the previous solution. There are several problems with this approach. First, as Grüne and Pannek note, not even all stage costs that satisfy the conditions under which Theorem 3.26 hold can satisfy this sufficient decrease condition, and the stage cost may have to be redesigned as a result. Another issue is when disturbances occur within the system. In that case, the sufficient decrease condition cannot be satisfied because the system has been disturbed. It is unclear how to redesign this decrease condition to take disturbances into account. The formulation with a terminal region, by contrast, is robust to disturbances.

3.6.4 *Conclusions*

We have spent quite a bit of time on a problem that, for all intents and purposes, is solved. Chapter 6 in Grüne and Pannek (2016) gives a very detailed treatment of this problem, and although it might be inaccessible to beginners, it develops tight bounds for the horizon length necessary for MPC to be stabilizing. My final verdict, however, is that it is usually better to include a terminal region. However, there are several good reasons we have spent so much time on this problem.

The first reason is that it was not clear upon beginning the investigation that the terminal region method would prove to be superior. There is not, to my knowledge, any equivalent of Proposition 3.30 in the literature, or any awareness that we can expect Theorem 3.26 to hold with a quadratic stage cost only when the linearized system is stabilizable. If this requirement is disregarded, then practical stability

might result. If this is not permissible, designing a stage cost and terminal region seems like an easier solution than trying to design a stage cost that exponentially decreases.

The second reason, though, is that while the method including a terminal region seems superior when regulating to a steady state, more complicated control problems can be considered. In the next chapter, we consider tracking an arbitrary reference trajectory. While terminal regions are possible (albeit cumbersome) to design for periodic reference trajectories (see (Rawlings et al., 2017, Sec. 2.5.6) for details), for aperiodic references it is impossible to do so using current techniques. In that case, a terminal equality constraint or sufficiently long horizon are the only techniques that remain (one of the motivations for introducing terminal equality constraints in Kwon and Pearson (1978) was to stabilize time-varying systems, for which the solution to the infinite-horizon LQR could not be computed). Finally, in Chapter 4 we cover state estimation, for which any sort of terminal condition is impossible. Many of the same issues that we faced in design of long-horizon steady state regulation occur in those contexts—in particular, the necessity of a well-designed stage cost if practical stability is not acceptable— and the same analysis tools we used here can be used in these more general contexts.

CHAPTER **4**

OUTPUT TRACKING

Now we pass from the problem of stabilizing a fixed setpoint (the regulator problem) to that of stabilizing a reference trajectory (the tracking or servo problem). Trying to track a reference trajectory is a classic problem in control, and there exists an extensive literature devoted to the problem both in linear and nonlinear contexts. An additional wrinkle is added to the problem when it is desirable to penalize a difference in system outputs rather than a difference in system states. The requirement that the outputs be penalized instead of the states is somewhat artificial, because the system's entire state is still necessary to compute the required control action, but this requirement has been considered since Kalman (1960b) proposed the LQR in the western literature.

Penalizing the system outputs instead of states can be justified for linear systems,

however, because linear system models are frequently created through a black-box identification process, and as a result the system states are mathematical constructs without any straightforward physical interpretation. Nonlinear systems are typically created through grey-box identification schemes in which the model's form is deduced from laws of physics and chemistry, but which have unknown parameters that must be identified. The recent revival of neural nets as a method for nonlinear black-box model identification, though, creates an application for penalization of outputs for nonlinear systems. Even if the behavior of the process outputs are all that matters in a physical application, it is desirable that the (artificial) process state be bounded in order to avoid catastrophic cancellation in floating-point arithmetic.

However, the primary reason that we examine this problem is not hypothetical concerns about what might be desirable in neural net models, but rather because this problem lies directly between conventional MPC, which is well-studied and well-understood, and optimization-based state estimation methods like FIE and MHE, which are less studied and much less understood. In an insightful comment, Rawlings and Ji (2012) made the observation:

Another issue that may hinder many researchers from making quick exploratory forays into state estimation is the inherent problem complexity. The simplest control problem is regulation to the origin, and the study of that simplest problem is still far from complete. There is no serious state estimation problem without an entire sequence of measurements. And the appearance of a sequence of measurements significantly complicates the notation and basic state estimation prob-

lem statements. It's more difficult for newcomers to get a foothold in this subject. It would be as if every control problem were stated as a tracking problem with a time-varying setpoint sequence. If that were how one had to get started in control problems, then state estimation would be about the same complexity and would probably compete better for new researchers' attention.

The situation is even worse than their description, though, because they do not mention the fact that the tracking problem must penalize the system's outputs rather than states. In order to deal with the problems created by penalizing outputs instead of states, we require an i-IOSS Lyapunov function.

4.1 Literature review

The dominate paradigm for output tracking in the literature involves a parallel autonomous system, termed exosystem, that generates the signal that is being tracked. The exosystem state enters into both the system evolution equation and an output error equation. The goal of the controller is that the output error equation converge to zero. An early example of this approach is provided by Isidori and Byrnes (1990) who use feedback linearization, while Krener (1998) uses optimal control by Al'brect's method. A good review of this approach is given by Byrnes and Isidori (2000). An approach to this problem using MPC is proposed by Magni et al. (2001), but it requires the solution of the Francis, Byrnes, and Isidori partial differential equations in order to implement. Another MPC approach is given by

Falugi and Mayne (2013), in which a periodic reference trajectory is first computed and then tracked. Although this paradigm results in a general problem that can handle more than simple reference tracking, we do not develop output tracking in this framework because the link to FIE and MHE is not as apparent. Furthermore, most results of this class require a neutrally stable exosystem. We would prefer to consider tracking to potentially unbounded trajectories, and it is not obvious how they could be encoded by a neutrally stable system.

Instead, we use a direct approach, in which we are given a sequence of inputs and outputs, and use an objective function that minimizes error from those reference sequences. Using a direct approach that minimizes *state* error from a reference sequence is a natural extension of regulation to a setpoint for MPC, and has been studied in the context of nonlinear MPC by Grüne and Pannek (2016, Sec. 3.3) and Rawlings and Risbeck (2017). Often, the reference trajectory is assumed to be feasible, i.e., satisfy the state evolution equation. Limon et al. (2016) consider tracking to a (potentially infeasible) periodic output trajectory using linear MPC with a terminal equality constraint, and Köhler et al. (2019) consider tracking to a (potentially infeasible) periodic *state* trajectory using nonlinear MPC. Conversely, Grimm et al. (2005) and Rawlings et al. (2017, Sec. 2.4.4) both consider using stage costs that are merely *detectable* for nonlinear MPC in order to regulate the system to some fixed steady-state. To my knowledge, however, there is no work in the literature about tracking to a (potentially aperiodic) output trajectory using nonlinear MPC.

4.2 Infinite horizon controller

We consider time-invariant discrete-time systems

$$\begin{aligned} x^+ &= f(x, u) \\ y &= h(x) \end{aligned} \tag{4.1}$$

in which $x \in \mathbb{X}$ is the state, $u \in \mathbb{U}$ is the input, $x^+ \in \mathbb{X}$ is the successor state, and $y \in \mathbb{Y}$ is the output.

We want to stabilize a reference trajectory starting from x_r with the reference control sequence \mathbf{u}_r by penalizing the difference in both inputs and outputs.

Assumption 4.1. Let $\ell : \mathbb{X}^2 \times \mathbb{U}^2 \rightarrow \mathbb{R}_{\geq 0}$. Let there exist $\alpha_y, \alpha_u \in \mathcal{K}_{\infty}$ such that

$$\ell(x_1, x_2, u_1, u_2) \geq \alpha_y(|h(x_1) - h(x_2)|) + \alpha_u(|u_1 - u_2|)$$

for all $x_1, x_2 \in \mathbb{X}$ and $u_1, u_2 \in \mathbb{U}$.

For ease of reference, let $x_r^+ := f(x_r, u_r(0))$, $\mathbf{u}_r^+ := \mathbf{u}_r(1 : \infty)$, $x_r(k) := \phi(k; x_r, \mathbf{u}_r)$, and $\mathbf{u}_r(k) := \mathbf{u}_r(k : \infty)$. We then define an infinite-horizon optimal control problem

$$\begin{aligned} \mathbb{P}_{\infty}(x, x_r, \mathbf{u}_r) : \inf_{\mu \in \mathbb{U}^{\infty}} V_{\infty}(x, x_r, \mu, \mathbf{u}_r) &:= \sum_{k=0}^{\infty} \ell(\chi(k), x_r(k), \mu(k), u_r(k)) \\ \text{such that } \chi^+ &= f(\chi, \mu) \\ \chi(0) &= x. \end{aligned}$$

Before we can discuss the existence of a solution to this problem, we require a stabilizability assumption. Because we have already explored the conditions for which we can guarantee only semiglobal practical stability in Section 3.2, we proceed

immediately to conditions for which we can guarantee global stability with a finite horizon. Therefore, we need to guarantee that the stage cost decays exponentially along a solution, as in Section 3.3.

Our first instinct to state a stabilizability assumption along the lines of Assumption 3.24 might be in terms of

$$\ell^*(x_1, x_2) := \inf_{u_1, u_2} \ell(x_1, x_2, u_1, u_2).$$

Note that we cannot guarantee that the minimum is attained by any input pair u_1, u_2 in this case. Such a definition would be sufficient if we were penalizing states directly and, as such, had a positive definite stage cost. However, we are guaranteed that $\ell^*(x_1, x_2)$ is strictly positive only if $h(x_1) \neq h(x_2)$. Stating stabilizability in terms of $\ell^*(x_1, x_2)$, would require $V_\infty^0(x, x_r, \mathbf{u}_r)$ to be zero for any pair of initial states that have the same output, which is unnecessarily restrictive. To address this issue, we define $\alpha : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\rho_1(|x_1 - x_2|) \leq \alpha(x_1, x_2) \leq \rho_2(|x_1 - x_2|) \quad (4.2)$$

in which both $\rho_1, \rho_2 \in \mathcal{K}_\infty$. Note that $\alpha(\cdot)$ has all the properties of a metric except satisfaction of the triangle inequality, and as such is a semimetric. Use of such a function allows different directions to converge at different rates while retaining the property of exponential convergence as measured by $\alpha(\cdot)$.

Now we can state conditions for which $\mathbb{P}_\infty(x, x_r, \mathbf{u}_r)$ has a solution.

Assumption 4.2 (Continuity). The functions $f : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{X}$, $h : \mathbb{X} \rightarrow \mathbb{R}^p$, and $\ell : \mathbb{X}^2 \times \mathbb{U}^2 \rightarrow \mathbb{R}_{\geq 0}$ are continuous, and the sets \mathbb{X} and \mathbb{U} are closed.

Assumption 4.3 (Incremental stabilizability). For every x , x_r , and \mathbf{u}_r , there exists some control sequence \mathbf{u}^* such that

$$V_\infty(x, x_r, \mathbf{u}^*, \mathbf{u}_r) \leq \bar{c} \alpha(x, x_r)$$

in which $\bar{c} > 0$ is a constant independent of x , x_r , and \mathbf{u}_r .

For any given x , x_r , and \mathbf{u}_r , there exists some \mathbf{u}^* that gives this problem finite cost $\rho \leq \bar{c} \alpha(x, x_r)$. As a result, for any \mathbf{u} such that $V_\infty(x, x_r, \mathbf{u}, \mathbf{u}_r) \leq \rho$, we have that $|u(k) - u_r(k)| \leq \alpha_u^{-1}(\rho)$. Therefore, $V_\infty(\cdot)$ is being optimized over a sequence of compact sets, and thus (Keerthi and Gilbert, 1985, Theorem 1) guarantees that there exists some \mathbf{u}^0 such that

$$V_\infty^0(x, x_r, \mathbf{u}_r) := \inf_{\mathbf{u} \in \mathbb{U}^\infty} V_\infty(x, x_r, \mathbf{u}, \mathbf{u}_r) = V_\infty(x, x_r, \mathbf{u}_r, \mathbf{u}^0),$$

i.e., the minimum is attained with \mathbf{u}^0 . We denote the set of solutions $\mathcal{U}_\infty^0(x, x_r, \mathbf{u}_r)$. Because we are tracking to an output sequence rather than a state sequence, we must discuss a notion of detectability. As discussed in Chapter 2, a natural candidate for nonlinear detectability is incremental input/output-to-state stability (i-IOSS), and this property can be characterized by an i-IOSS Lyapunov function. However, unlike linear systems, we cannot simply assume that the system is detectable. We require some compatibility between the i-IOSS Lyapunov function, the stage cost $\ell(\cdot)$, and the semimetric $\alpha(\cdot)$.

Assumption 4.4 (i-IOSS Lyapunov function). There exists a function $\lambda : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ that satisfies the bounds

$$c_1\alpha(x_1, x_2) \leq \lambda(x_1, x_2) \leq c_2\alpha(x_1, x_2) \quad (4.3)$$

$$\begin{aligned} \lambda(f(x_1, u_1), f(x_2, u_2)) &\leq \lambda(x_1, x_2) - c_3\alpha(x_1, x_2) \\ &\quad + \sigma_u(|u_1 - u_2|) + \sigma_y(|y_1 - y_2|), \end{aligned} \quad (4.4)$$

for some $c_1, c_2, c_3 > 0$, in which $\sigma_u(\cdot), \sigma_y(\cdot)$ come from Assumption 4.1, for all $x_1, x_2 \in \mathbb{X}$ and $u_1, u_2 \in \mathbb{U}$.

One consequence of the form that this assumption takes is that the stage cost $\ell(\cdot)$ constitutes a valid supply rate for λ :

$$\lambda(f(x_1, u_1), f(x_2, u_2)) \leq \lambda(x_1, x_2) - c_3\alpha(x_1, x_2) + \ell(x_1, x_2, u_1, u_2). \quad (4.5)$$

The property that $\ell(\cdot)$ constitutes a valid supply rate for $\lambda(\cdot)$ is the key to proceeding from behavior about the input sequence \mathbf{u} and the output sequence \mathbf{y} to behavior about the state sequence \mathbf{x} . Grimm et al. (2005) skipped the intermediate step of assuming a relationship between $\ell(\cdot)$ and the behavior of the output y , instead using a (non-incremental) version of (4.5) directly as an assumption of detectability. We go through this intermediate step in order to make the relationship between this detectability assumption and the sorts of detectability assumptions used in the optimization-based state estimation literature more explicit.

We can now guarantee the convergence of \mathbf{x} to \mathbf{x}_r under $\mathcal{K}_\infty^0(\cdot)$. Unlike in Chapter 3, $V_\infty^0(x, x_r, \mathbf{u}_r)$ cannot serve as an (incremental parametric) Lyapunov function. We do not have that $x \neq x_r$ implies $V_\infty^0(x, x_r, \mathbf{u}_r) > 0$, and the Bellman equation

$$V_\infty^0(x, x_r, \mathbf{u}_r) = V_\infty^0(x^+, x_r^+, \mathbf{u}_r^+) + \ell(x, x_r, u^0, u_r), \quad (4.6)$$

in which $u^0 \in \mathcal{K}_\infty^0(x, x_r, \mathbf{u}_r)$, does not guarantee a strict cost decrease. However, we can use $\lambda(\cdot)$ in order to produce the (incremental parametric) Lyapunov function we seek. Let

$$Y_\infty(x, x_r, \mathbf{u}_r) := V_\infty^0(x, x_r, \mathbf{u}_r) + \lambda(x, x_r)$$

From (4.3) and Assumption 4.3, we immediately have

$$c_1\alpha(x_1, x_2) \leq Y_\infty(x, x_r, \mathbf{u}_r) \leq \bar{c}_2\alpha(x_1, x_2), \quad (4.7)$$

in which $\bar{c}_2 := c_2 + \bar{c}$. We can then combine (4.5) and (4.6) to obtain

$$\begin{aligned} Y_\infty(x^+, x_r^+, \mathbf{u}_r^+) &= V_\infty^0(x^+, x_r^+, \mathbf{u}_r^+) + \lambda(x^+, x_r^+) \\ &\leq V_\infty^0(x, x_r, \mathbf{u}_r) - \ell(x, x_r, u^0, u_r) + \lambda(x, x_r) \\ &\quad - c_3\alpha(x, x_r) + \ell(x, x_r, u^0, u_r) \\ &= Y_\infty(x, x_r, \mathbf{u}_r) - c_3\alpha(x, x_r), \end{aligned} \quad (4.8)$$

which is the strict descent condition we seek. We can then guarantee that the reference trajectory \mathbf{x}_r is stabilized.

Proposition 4.5. *Suppose that Assumptions 4.1 to 4.4 hold. Then there exists $\beta \in \mathcal{KL}$ such that*

$$|x(k) - x_r(k)| \leq \beta(|x(0) - x_r(0)|, k),$$

when the system (4.1) is controlled with the control law $u \in \mathcal{K}_\infty^0(x, x_r, \mathbf{u}_r)$, for all $x, x_r \in \mathbb{X}$ and $\mathbf{u}_r \in \mathbb{U}^\infty$.

Proof. Let $\eta := 1 - c_3/\bar{c}_2 \in (0, 1)$. By combining (4.7) and (4.8), we obtain

$$Y_\infty(x^+, x_r^+, \mathbf{u}_r^+) \leq \eta Y_\infty(x, x_r, \mathbf{u}_r).$$

We can apply this equation repeatedly to obtain

$$Y_\infty(x(k), x_r(k), \mathbf{u}_r(k)) \leq \eta^k Y_\infty(x(0), x_r(0), \mathbf{u}_r).$$

We can then apply (4.7) to obtain

$$\alpha(x(k), x_r(k)) \leq \bar{c}_2/c_1 \eta^k \alpha(x(0), x_r(0))$$

and then apply (4.2) to obtain

$$|x(k) - x_r(k)| \leq \rho_1^{-1}(\bar{c}_2/c_1 \eta^k \rho_2(|x(0) - x_r(0)|))$$

which is a bound of the form sought. ■

4.3 Stability with a finite horizon

Again, as in Chapter 3, we wish to achieve stabilization with a finite horizon. For the tracking problem, however, our options are more limited than before. Terminal region design in particular becomes rather cumbersome. There are methods for designing terminal regions for periodic trajectories (see (Rawlings et al., 2017, Sec. 2.5.6) for details), but there is not, to my knowledge any general method applicable for arbitrary reference trajectories.¹ A terminal equality constraint would suffice to guarantee stability, but in order to guarantee reasonable behavior for gradient-based solvers, the time-varying linear system that results from linearization about \mathbf{x}_r and \mathbf{u}_r must be uniformly controllable.

In order to guarantee stability for stabilizable systems, we analyze MPC without stabilizing terminal conditions. We thus define the finite-horizon OCP

¹One could imagine using the method Kwon and Pearson (1978) proposed for time-varying linear systems using a terminal equality constraint to create a locally stabilizing control law about \mathbf{x}_r and \mathbf{u}_r . While this method may succeed in increasing the region for which the OCP is feasible, it does not avoid the requirement that the linearized system is uniformly controllable about \mathbf{x}_r and \mathbf{u}_r .

$$\mathbb{P}_N(x, x_r, \mathbf{u}_r) : \min_{\mu \in \mathbb{U}^\infty} V_\infty(x, x_r, \mu, \mathbf{u}_r) := \sum_{k=0}^{N-1} \ell(\chi(k), x_r(k), \mu(k), \mathbf{u}_r(k))$$

such that $\chi^+ = f(\chi, \mu)$
 $\chi(0) = x.$

A solution to this optimization problem always exists under Assumption 4.2, but it may not be unique. We denote the set of solutions $\mathcal{U}_N^0(x, x_r, \mathbf{u}_r)$, and denote the control law induced by it as $\mathcal{K}_N^0(x, x_r, \mathbf{u}_r)$. We denote the j^{th} state and input of the optimal control problem with horizon length N by $x^0(j|N)$ and $u^0(j|N)$, respectively.

We use a two-part argument to establish the global asymptotic stability of MPC with a finite control horizon. We first show that the sequence of functions $(V_N^0(x, x_r, \mathbf{u}_r))$ converges to $V_\infty^0(x, x_r, \mathbf{u}_r)$ exponentially by use of a Q function. This result on convergence allows us to find an upper bound on the cost increase from the stages neglected by the truncated horizon. We then show that, for a sufficiently long horizon, this cost increase can be overcome by the cost decrease from the first stage cost.

In a situation similar to that with $V_\infty^0(\cdot)$, however, the function

$$Z(j|k) := V_\infty^0(x, x_r, \mathbf{u}_r) - \sum_{i=0}^{j-1} \ell(x^0(i|k), x_r(i), u^0(i|k), \mathbf{u}_r(i))$$

defined for $j \in \mathbb{I}_{0:k-1}$, in which we suppress the dependence of $Z(\cdot)$ on x , x_r , and \mathbf{u}_r

for brevity, cannot serve as a Q function.² Like $V_\infty^0(\cdot)$, it is not necessarily positive

²Note that $Z(j|k)$ does, in principle, depend on the choice of $\mathbf{u}_k^0 \in \mathcal{U}_N^0(x, x_r, \mathbf{u}_r)$, but neither $Z(0|k)$ nor $Z(k|k)$ do. The upper and lower bounds derived for $Z(\cdot)$ (and subsequently $Q(\cdot)$) hold irrespective of the selection of \mathbf{u}_k^0 .

definite and there does not exist a strict descent condition. We can, however, derive some useful bounds for its behavior. It is useful to abbreviate the partial sum

$$V^0(j|k) := \sum_{i=0}^{j-1} \ell(x^0(i|k), x_r(i), u^0(i|k), u_r(i)).$$

While $Z(j|k)$ is not positive-definite, it is nonnegative, because

$$V^0(j|k) \leq V_k^0(x, x_r, \mathbf{u}_r) \leq V_\infty^0(x, x_r, \mathbf{u}_r).$$

Furthermore, we can use \mathbf{u}_k^0 as a feasible control sequence for $i \in \mathbb{I}_{0:j-1}$ and reoptimize the tail to obtain an upper bound for $V_\infty^0(\cdot)$. As a result, we have that

$$V_\infty^0(x, x_r, \mathbf{u}_r) \leq V^0(j|k) + V_\infty^0(x^0(j|k), x_r(j), \mathbf{u}_r(j)) \leq V^0(j|k) + \bar{c}\alpha(x^0(j|k), x_r(j)),$$

in which the last step follows from Assumption 4.3. Finally, because

$$V^0(j+1|k) = V^0(j|k) + \ell(x^0(j|k), x_r(j), u^0(j|k), u_r(j))$$

for $j \in \mathbb{I}_{0:k-1}$, we get a trivial cost decrease condition. Summarizing, we have that

$$0 \leq Z(j|k) \leq \bar{c}\alpha(x^0(j|k), x_r(j)) \quad \forall j \in \mathbb{I}_{0:k} \quad (4.9)$$

$$Z(j+1|k) \leq Z(j|k) - \ell(x^0(j|k), x_r(j), u^0(j|k), u_r(j)) \quad \forall j \in \mathbb{I}_{0:k-1}. \quad (4.10)$$

Like the situation with $V_\infty^0(\cdot)$, however, $\lambda(\cdot)$ provides the solution to turn $Z(j|k)$ into a Q function. Let

$$Q(j|k) := Z(j|k) + \lambda(x^0(j|k), x_r(j))$$

for $j \in \mathbb{I}_{0:k}$. We immediately have that

$$c_1\alpha(x^0(j|k), x_r(j)) \leq Q(j|k) \leq \bar{c}_2\alpha(x^0(j|k), x_r(j)), \quad (4.11)$$

in which we recall that $\bar{c}_2 := \bar{c} + c_2$, by using (4.3) and (4.9). We can obtain a descent condition by applying (4.4) and (4.10) to obtain

$$\begin{aligned}
Q(j+1|k) &= Z(j+1|k) + \lambda(x^0(j|k), x_r(j)) & (4.12) \\
&\leq Z(j|k) - \ell(x^0(j|k), x_r(j), u^0(j|k), u_r(j)) + \lambda(x^0(j|k), x_r(j)) \\
&\quad - c_3 \alpha (x^0(j|k), x_r(j)) + \ell(x^0(j|k), x_r(j), u^0(j|k), u_r(j)) \\
&= Q(j|k) - c_3 \alpha (x^0(j|k), x_r(j)),
\end{aligned}$$

for $j \in \mathbb{I}_{0:k-1}$. Finally, we note that

$$Q(0|k) := V_\infty^0(x, x_r, \mathbf{u}_r) + \lambda(x, x_r) = Y_\infty(x, x_r, \mathbf{u}_r)$$

These properties of $Q(\cdot)$ allow us to show that $V_k^0(\cdot)$ converges to $V_\infty^0(\cdot)$ exponentially.

Proposition 4.6. *Under Assumptions 4.1 to 4.4, there exists $\eta \in (0, 1)$ such that*

$$V_\infty^0(x, x_r, \mathbf{u}_r) - V_k^0(x, x_r, \mathbf{u}_r) \leq \eta^k Y_\infty(x, x_r, \mathbf{u}_r) \quad (4.13)$$

for all $x, x_r \in \mathbb{X}$, $\mathbf{u}_r \in \mathbb{U}^\infty$, and $k \in \mathbb{I}_{\geq 0}$.

Proof. Let $\eta := 1 - c_3/\bar{c}_2 \in (0, 1)$, as in Proposition 4.5. We can combine (4.11) and (4.12) to obtain

$$\begin{aligned}
Q(j+1; x, x_r, \mathbf{u}_k^0, \mathbf{u}_r) &\leq Q(j; x, x_r, \mathbf{u}_k^0, \mathbf{u}_r) - \frac{c_3}{\bar{c}_2} Q(j; x, x_r, \mathbf{u}_k^0, \mathbf{u}_r) \\
&= \eta Q(j; x, x_r, \mathbf{u}_k^0, \mathbf{u}_r).
\end{aligned}$$

This bound can be applied recursively to obtain

$$Q(k|k) \leq \eta^k Q(0|k) = \eta^k Y_\infty(x, x_r, \mathbf{u}_r).$$

Finally, by definition of $Q(\cdot)$ and nonnegativity of $\lambda(\cdot)$, we have that

$$V_\infty^0(x, x_r, \mathbf{u}_r) - V_k^0(x, x_r, \mathbf{u}_r) \leq Q(k|k) \leq \eta^k Y_\infty(x, x_r, \mathbf{u}_r)$$

which is the sought-after bound. ■

Note that the i-IOSS Lyapunov function $\lambda(\cdot)$ is an integral part of guaranteeing a convergence rate. It essentially acts as a “timer” for how long it takes for a signal to appear in the output. The longer time goes on, the smaller the signal that can appear.

In a situation similar to the infinite-horizon case, we define a Lyapunov function candidate

$$Y_N(x, x_r, \mathbf{u}_r) := V_N^0(x, x_r, \mathbf{u}_r) + \lambda(x, x_r).$$

By noting that $Y_N(x, x_r, \mathbf{u}_r) \leq Y_\infty(x, x_r, \mathbf{u}_r)$, we can apply (4.3) and (4.7) to obtain

$$c_1 \alpha(x, x_r) \leq Y_N(x, x_r, \mathbf{u}_r) \leq \bar{c}_2 \alpha(x, x_r). \quad (4.14)$$

For a descent condition, more work is necessary. By the principle of optimality, we have that

$$V_{N-1}^0(x^+, x_r^+, \mathbf{u}_r^+) = V_N^0(x, x_r, \mathbf{u}_r) - \ell(x, u^0, x_r, u_r),$$

in which $x^+ = f(x, u^0)$, for any $u^0 \in \mathcal{K}_N(x, x_r, \mathbf{u}_r)$. We can rearrange (4.13) to obtain

$$\begin{aligned} V_\infty^0(x^+, x_r^+, \mathbf{u}_r^+) &\leq \frac{1}{1 - \eta^{N-1}} V_{N-1}^0(x^+, x_r^+, \mathbf{u}_r^+) + \frac{\eta^{N-1}}{1 - \eta^{N-1}} \lambda(x^+, x_r^+) \\ &= \frac{1}{1 - \eta^{N-1}} V_N^0(x, x_r, \mathbf{u}_r) - \frac{1}{1 - \eta^{N-1}} \ell(x, x_r, u, u_r) \\ &\quad + \frac{\eta^{N-1}}{1 - \eta^{N-1}} \lambda(x^+, x_r^+) \end{aligned}$$

for $N \geq 1$. By monotonicity of the value function, we have that

$$V_N^0(x^+, x_r^+, \mathbf{u}_r^+) \leq V_\infty^0(x^+, x_r^+, \mathbf{u}_r^+).$$

Thus we have that

$$\begin{aligned}
Y_N(x^+, x_r^+, \mathbf{u}_r^+) &:= V_N^0(x^+, x_r^+, \mathbf{u}_r^+) + \Lambda(x^+, x_r^+) \\
&\leq V_\infty^0(x^+, x_r^+, \mathbf{u}_r^+) + \Lambda(x^+, x_r^+) \\
&\leq \frac{1}{1 - \eta^{N-1}} V_N^0(x, x_r, \mathbf{u}_r) - \frac{1}{1 - \eta^{N-1}} \ell(x, x_r, u, u_r) \\
&\quad + \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \Lambda(x^+, x_r^+).
\end{aligned}$$

We can apply (4.5) to obtain

$$\begin{aligned}
Y_N(x^+, x_r^+, \mathbf{u}_r^+) &\leq \frac{1}{1 - \eta^{N-1}} V_N^0(x, x_r, \mathbf{u}_r) - \frac{1}{1 - \eta^{N-1}} \ell(x, x_r, u, u_r) \\
&\quad + \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \left(\Lambda(x, x_r) - c_3 \alpha(x, x_r) + \ell(x, x_r, u, u_r) \right).
\end{aligned}$$

Finally, because

$$V_N^0(x, x_r, \mathbf{u}_r) \geq \ell(x, x_r, u, u_r)$$

for all $N \geq 1$, we have that

$$\begin{aligned}
Y_N(x^+, x_r^+, \mathbf{u}_r^+) &\leq \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} V_N^0(x, x_r, \mathbf{u}_r) - \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \ell(x, x_r, u, u_r) \\
&\quad + \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \left(\Lambda(x, x_r) - c_3 \alpha(x, x_r) + \ell(x, x_r, u, u_r) \right) \\
&= \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \left(Y_N(x, x_r, \mathbf{u}_r) - c_3 \alpha(x, x_r) \right). \tag{4.15}
\end{aligned}$$

So $Y_N(\cdot)$ loses a term proportional to $\alpha(x, x_r)$, but grows by a factor that decreases with horizon length. From this equation, it is straightforward to find a horizon length sufficiently long to guarantee that the reference trajectory \mathbf{x}_r is stabilized with a finite horizon.

Theorem 4.7. *There exists a horizon length $T \in \mathbb{N}_{\geq 0}$ such that, for all horizons $N \geq T$, states $x \in \mathbb{X}$, and reference trajectories $x_r \in \mathbb{X}$ and $\mathbf{u}_r \in \mathbb{U}^\infty$ MPC with horizon length N stabilizes the reference trajectory, i.e., there exists $\beta_N \in \mathcal{KL}$ independent of x , x_r , and \mathbf{u}_r such that*

$$|x(k) - x_r(k)| \leq \beta_N(|x - x_r|, k)$$

for all $k \in \mathbb{N}_{\geq 0}$.

Proof. We can combine (4.14) and (4.15) to obtain

$$\begin{aligned} Y_N(x^+, x_r^+, \mathbf{u}_r^+) &\leq \frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \left(Y_N(x, x_r, \mathbf{u}_r) - \frac{c_3}{\bar{c}_2} Y_N(x, x_r, \mathbf{u}_r) \right) \\ &= \eta \left(\frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \right) Y_N(x, x_r, \mathbf{u}_r) \end{aligned}$$

in which we recall that $\eta := 1 - c_3/\bar{c}_2$. Because $\eta \in (0, 1)$, there exists some $T \geq 2$ such that

$$\lambda(N) := \eta \left(\frac{1 + \eta^{N-1}}{1 - \eta^{N-1}} \right) \in (0, 1).$$

for all $N \geq T$. For such N , $Y_N(\cdot)$ is an exponential-decrease incremental parametric Lyapunov function. We can repeatedly apply the cost-decrease condition to obtain

$$Y_N(x(k), x_r(k), \mathbf{u}_r(k)) \leq \lambda(N)^k Y_N(x, x_r, \mathbf{u}_r)$$

and apply (4.14) to find

$$\alpha(x(k), x_r(k)) \leq \frac{\bar{c}_2}{c_1} \lambda(N)^k \alpha(x, x_r)$$

i.e., $\alpha(\cdot, \cdot)$ is an exponentially decreasing measure of distance between x and x_r . We can then apply (4.2) to find

$$|x(k) - x_r(k)| \leq \rho_1^{-1} \left(\frac{\bar{c}_2}{c_1} \lambda(N)^k \rho_2(|x - x_r|) \right) := \beta_N(|x - x_r|, k)$$

in which $\beta_N \in \mathcal{KL}$ for all $N \geq T$, which is a bound of the form sought. ■

4.4 Conclusions

We have shown that MPC can stabilize an arbitrary feasible trajectory with a finite horizon but without stabilizing terminal conditions. In principle, we could analyze tracking problems like this one as a special case of Economic MPC. However, presently, results for economic MPC without stabilizing terminal conditions are in flux. The results in Grüne and Pannek (2016, Sec. 8.6) guarantee only practical stability. This shortcoming was overcome by Zanon and Faulwasser (2018) by use of a linear terminal penalty, but the two-step proof first guarantees practical stability, then shows that the control law is locally equivalent to a stabilizing linear controller. They also show that, so long as the stage cost has a nonzero gradient at the steady state in question, the linear terminal penalty is necessary for stability. This gradient term is not necessary in the case of output tracking, because if $\ell(\cdot)$ is differentiable, its derivative when $x = x_r$ and $u = u_r$ must be zero. Whether the techniques presented in this chapter can be extended to economic MPC is an interesting topic for future research.

One factor that simplified this problem is that, once on the reference trajectory, the controller can stay on it for zero cost. In this respect, the problem is different than the estimation problem. In the estimation problem, we assume that there is a real sequence of states and disturbances that generate the measured outputs, but nevertheless it costs the estimator to use those sequences in the optimal estimation

problem. In that respect, the estimation problem is more similar to a tracking problem in which the reference trajectory is infeasible, i.e., one in which $x_r^+ \neq f(x_r, u_r)$. This problem is studied using the framework of economic MPC by Köhler et al. (2019) penalizing a difference of states, rather than outputs. Treating each variation of the problem going from regulation to the origin with positive definite stage cost, to regulation with a detectable stage cost, to tracking with a detectable stage cost, to tracking with an infeasible trajectory, to finally the estimation problem, would weary the reader.³ So in the next chapter, we proceed directly to the state estimation problem.

³And I would be in graduate school forever.

CHAPTER 5

STATE ESTIMATION

Finally, we come to the central focus of this work—optimization-based state estimation. Discovering that Q functions are such a useful tool for analyzing MPC without stabilizing terminal conditions was a welcome surprise. However, the original purpose for which they were devised is to serve as a Lyapunov-like function for the state estimation problem. In the last two chapters, we have followed a pattern of first showing that the infinite-horizon problem is well-defined, then showing that the finite-horizon cost functions converge to this infinite-horizon cost function in a regular way by way of a Q function, then finally showing that there exists a horizon long enough that the finite-horizon cost function can serve as a Lyapunov function for the closed-loop system. In this chapter, however, we necessarily follow a different pattern.

In full information estimation (FIE), there is no equivalent to the closed-loop evolution of a system under a control law. The behavior of the sequence of open-loop finite-horizon cost functions is what matters for the estimation problem. We touched on the robust stability of MPC only tangentially, but, when considering the inherent robustness of MPC, disturbances enter only in the closed-loop evolution of the system. In the estimation problem, however, these disturbances enter in the open-loop sequence of optimal cost functions. As a result, there is typically no sequence of decision variables that attain a finite cost for an infinite-horizon problem formulated with the disturbed output sequence. The insight that permits us to get around this problem is that the FIE problem at some time depends only on the disturbances that have occurred up to that time. Therefore, we can formulate a modified infinite-horizon problem using the sequence of outputs that would result if the disturbances stopped at that time.

In moving horizon estimation (MHE), by contrast, there is both open-loop and closed-loop behavior. However, at present, it is not clear what sort of function would serve as a Lyapunov function. As a result, we must rely on contraction-mapping instead, and can obtain a robust stability result only for exponentially *i*-IOSS systems. Furthermore, because MHE is more internally complicated for reasons that receive a more detailed treatment later, the proofs of closed-loop behavior is more complex than that of MPC.

5.1 Literature review

Modern optimization-based state estimation has its origins in Kalman filtering. Precursors to the Kalman filter, such as the algorithm proposed by Swerling (1959), were based on recursive least squares estimation. Least squares problems are numerically attractive, and can be rationalized as giving the maximum likelihood estimate assuming data is corrupted with Gaussian noise. Kalman (1960a), by contrast, derived the filtering algorithm that bears his name through orthogonal projections. This method of derivation reveals that the Kalman filter does not just simply produce the maximum likelihood estimate so long as the process and measurement noises are normal, but in fact is the best linear unbiased estimator for *any* distributions of process and measurement noise (Anderson and Moore, 1981, Ch. 5). Nevertheless, the least squares formulation is both easier to understand for the uninitiated and easier to extend to nonlinear systems. Ho (1962) made the link between Kalman's work and earlier least squares work explicit, and maximum likelihood estimation problems were proposed for nonlinear continuous time systems by Bryson and Frazier (1963) and nonlinear discrete time systems by Cox (1964).

MHE was developed in several different lines of research. For linear unconstrained systems, a limited memory filter was proposed by Jazwinski (1970, Sec. 7.10) in order to avoid filter divergence in the case of model-system mismatch. In a result dual to the terminal equality constrained minimum energy control derived by Kleinman (1970), Thomas (1975) proposed using an observer gain derived from a moving horizon problem for unconstrained linear systems in order to reduce the number

of tuning parameters needed for Kalman filtering. Motivated by the success of MPC as a control method, several papers in the chemical engineering literature (Jang et al., 1986; Bequette, 1991; Liebman et al., 1992) that propose nonlinear least squares estimators were published. These works considered only measurement noise. Constrained linear MHE with both process and measurement noise was proposed by Muske and Rawlings (1993). Nonlinear MHE with both process and measurement noise was proposed by Robertson et al. (1996), and a probabilistic interpretation for state constraints was given. In an apparently independent thread of research in the automatic control community, Michalska and Mayne (1992) proposed a least squares estimator, partially inspired by their results in receding horizon control, and partially by the Newton-based observer of Grizzle and Moraal (1990). Zimmer (1994) also independently proposed such a least squares estimator.

Once these optimization-based state estimators were proposed, research about the conditions under which they were stable began. Muske et al. (1993) showed that constrained FIE was stable for detectable linear systems, and Rao et al. (2001) showed that constrained MHE was stable for observable linear systems. Michalska and Mayne (1992) showed that unconstrained MHE with no process noise and zero prior weighting was stable for nonlinear systems under an observability assumption. Muske and Rawlings (1995) showed that FIE with both process and measurement noise was stable for observable nonlinear systems. Rao et al. (2003) showed that constrained nonlinear MHE was stable for observable systems.

Once the nominal stability of FIE and MHE was established, research turned to characterizing the conditions under which stability would be robust. Unfor-

Unfortunately, the precise formulation of how robust stability should be characterized is a somewhat subtle issue. Many early papers formulated robust stability in a totally inadequate way, and their results must be parsed carefully in order to see exactly what kind of robust stability is guaranteed. Rawlings and Ji (2012, Def. 9) proposed an inadequate definition of robust global asymptotic stability (RGAS) that was unfortunately adopted by much of the optimization-based state estimation literature. The chief inadequacy is that the \mathcal{KL} function decay rate for the initial error and the \mathcal{K} function asymptotic gains are allowed to depend on the particular initial condition and disturbance sequences. Such a definition does not even guarantee local stability in the classic sense. The proof of robust stability of FIE they provided does, however, guarantee local stability, but not a uniform decay rate function. Another, lesser, inadequacy is that this definition attempts to use an asymptotic-gain definition of input-to-state stability as a definition of robust stability. However, for reasons discussed in Chapter 2, such a definition does not imply estimator convergence for convergent disturbances.

Another feature of this early MHE literature is that these \mathcal{K} and \mathcal{KL} functions depend implicitly on the horizon length, and are allowed to increase without upper bound if the horizon length increases. As we saw in Chapters 3 and 4, for MPC, a longer horizon produces a better approximation of the infinite-horizon optimal cost function, and as a result results in a tighter upper bound in the proof of nominal stability. However, these results do not preclude the scenario where a shorter horizon for MHE may be preferred to a longer one for *stability* reasons. In MPC, a longer horizon is always preferable for robust stability—the motive for shortening

the horizon is to reduce the computation burden.

With these disclaimers, we examine the early literature for the robust stability of both FIE and MHE. Rao et al. (2003, Prop. 3.7) made an attempt to address the robust stability of MHE, but an incorrect proof is used.¹ Alessandri et al. (2008) show, under an observability assumption, that a formulation of MHE with only measurement disturbances is robustly stable on compact sets,² but the upper bounds generated get worse with increasing horizon length. Because the process disturbance is not directly estimated, the divergence between the actual and estimated states as horizon length increases is inherent in this formulation of MHE; it is not from a deficiency in the theory. Rawlings and Mayne (2009, Ch. 4) made a number of important contributions to the theory of MHE and FIE, most notably the introduction of i-IOSS as a nonlinear detectability condition to the optimization-based estimation literature. Attempts at proving the robust stability of FIE for detectable systems and the robust stability of MHE with zero prior weighting for observable systems are present, but are marred by technical issues and significant gaps. Rawlings and Ji (2012) fill some of these gaps and propose an ISS-like definition of RGAS, but the issues with this definition have been noted above. They do prove uniform robust stability, but not uniform convergence.

Ji et al. (2016) attempt to show that FIE is robustly stable in the presence of

¹Rao et al. (2003, Prop. 3.7) assume that both the true initial state and estimated initial state reside in the same compact set. The proposition is stated as if this holds only when MHE is initialized at the time $k = 0$, but is used in a fashion as if this holds at the start of the horizon for every MHE problem thereafter.

²The robustness result (Alessandri et al., 2008, Theorem 1) is not formulated particularly well, but I believe the result could be modified to satisfy an ISS-like definition of robustness without much effort.

bounded disturbances by modifying the cost function. They normalize the sum of stage costs by the horizon length and add a term penalizing the maximum disturbance size (max term). By assuming special forms for the functions that define i-IOSS and the stage cost, and that they satisfy a certain relationship, it is shown that FIE is RGAS. However, they note that they cannot prove that this form of FIE converges for convergent disturbances. This argument is extended by Hu et al. (2015) to a broader class of systems, and certain conditions under which this robust stability can be extended to MHE with such a max term is established by Hu (2017). Furthermore, Hu (2017) is able to establish the convergence of this MHE formulation with convergent disturbances.³ Concurrently, Müller (2017) established convergence of MHE with *and without* a max term by assuming a particular form for both the stage cost and the functions defining i-IOSS. Allan and Rawlings (2019a) streamlined the results for MHE without a max term in (Müller, 2017), revealed that the detectability assumption used by Müller (2017) is equivalent to semiglobal exponential i-IOSS, and integrated additional insights from Hu (2017). However, the results by Müller (2017), Hu (2017), and Allan and Rawlings (2019a) have assumptions about compositions of \mathcal{K} functions that are difficult to interpret, much less verify. Furthermore, the results without a max term have robust stability bounds that become worse with increased horizon length.

Knüfer and Müller (2018) proposed a fading memory formulation of both FIE and MHE for exponentially i-IOSS systems. While this formulation does result in

³Hu et al. (2015) and Hu (2017) claim to also show that FIE with a max term converges to zero so long as knowledge that disturbances converge is used in the optimization problem. However, because there is no slowest-converging disturbance, FIE cannot be constrained to optimize over the set of convergent disturbance sequences, which is required for their proof.

robust stability bounds that become better with increased horizon length, it requires the stage cost satisfy the triangle inequality. Allan and Rawlings (2019b) introduced the notion of a Q function to analyze the nominal stability of FIE for detectable systems. Although this result is ostensibly weaker than its many predecessors in the literature ((Rawlings and Mayne, 2009, Thm. 4.8), (Rawlings and Ji, 2012, Thm. 12), and (Rawlings et al., 2017, Thm. 4.10)), these predecessors either had major gaps or relied on the inadequate definition of RGAS proposed in (Rawlings and Ji, 2012, Def. 9). Apart from the Q function, the two innovations that made it possible to fill those gaps are the introduction of a stabilizability assumption and the use of an i-IOSS Lyapunov function to characterize detectability. The later is covered in more detail in Chapter 2, but controllability or stabilizability assumptions are commonplace in the Kalman filtering literature (e.g., (Kalman, 1960a) and (Anderson and Moore, 1981)), but had not been explicitly used in the FIE or MHE literature.⁴ Allan and Rawlings (2021) demonstrated the robust exponential stability of both FIE and MHE for exponentially i-IOSS systems and power-law stage costs by extending Q function analysis to FIE with bounded disturbances.

Here, we treat the case of FIE with an asymptotically i-IOSS system. Unfortunately, the current state of theory does not allow us to then apply this result to MHE, for reasons that are discussed further along. Semiglobal practical results are possible, but we instead present a version of the result for exponentially i-IOSS systems published in (Allan and Rawlings, 2021).

⁴The use of an additive state disturbance, as in (Knüfer and Müller, 2018), can however be seen as a (strong) controllability assumption.

5.2 Full information estimation

As before, we consider a discrete time system

$$x^+ = f(x, w) \quad y = h(x) + v \quad (5.1)$$

in which $x \in \mathbb{X}$ is the system state, $w \in \mathbb{W} \subseteq \mathbb{R}^g$ is a process disturbance, $y \in \mathbb{Y}$ is the system output, and $v \in \mathbb{R}^p$ is a measurement disturbance. The input is named differently to represent its different role in the problem. Previously, the input came from the controller. Now, it simply enters the system and we are forced to reconstruct it. In Chapter 4, there is no equivalent to the measurement disturbance v .

In previous chapters, we formulated an infinite-horizon problem first, then approximated it with finite horizon problems. Here, however, it makes the most sense to formulate a finite horizon problem first. To reconstruct the system's state, we pose the FIE problem

$$\begin{aligned} \mathbb{P}_k(\bar{x}, \mathbf{y}(k)) : \quad & \min_{\chi(0) \in \mathbb{X}, \omega \in \mathbb{W}^k, \nu \in (\mathbb{R}^p)^k} V_k(x, \omega, \nu, \bar{x}, \mathbf{y}) := \ell_x(\chi(0), \bar{x}) + \sum_{i=0}^{k-1} \ell(\omega(i), \nu(i)) \\ & \text{subject to} \quad \chi^+ = f(\chi, \omega) \\ & \quad \quad \quad y = h(\chi) + \nu. \end{aligned}$$

Note that while \mathbf{y} here, as in Chapter 4, refers to the sequence of outputs generated by the real system, in $\mathbb{P}_k(\cdot)$ it plays a role similar to \mathbf{y}_r from Chapter 4. The major change in the problem structure here is the addition of the prior weighting term $\ell_x(\cdot)$. In the regulation problem, the initial state is a fixed value known to the

controller. In the estimation problem, the initial state is a decision variable, and $\ell_x(\cdot)$ penalizes the distance from \bar{x} , a prior estimate of the initial state $x(0)$.

We denote the estimated states, state disturbances, and measurement disturbances $\widehat{\mathbf{x}}(k)$, $\widehat{\mathbf{w}}(k)$, and $\widehat{\mathbf{v}}(k)$, respectively. For $j \in \mathbb{I}_{0:k}$ we refer to the j^{th} element of $\widehat{\mathbf{x}}(k)$ as $\hat{x}(j|k)$. For the k^{th} element of that vector, we denote $\hat{x}(k) := \hat{x}(k|k)$ for simplicity. Similarly, we denote the j^{th} estimated state and measurement disturbances as $\hat{w}(j|k)$ and $\hat{v}(j|k)$, respectively, for $j \in \mathbb{I}_{0:k-1}$. Finally, we denote the estimation error $e(j|k) := x(j) - \hat{x}(j|k)$ and the prior error $\bar{e} := x(0) - \bar{x}$.

Remark 5.1. In Chapter 2 we have a disturbance forecast \mathbf{u}_f to account from which \mathbf{w} is a deviation. Similarly, in Chapter 4 both the controller's input sequence and a reference input sequence appear in the stage cost. Here, however, only a single input appears in the stage cost. Traditionally in estimation, w was considered to be a zero-mean Gaussian noise, which leads to a least-squares stage cost. In analogy to more general stage costs in MPC, however, the requirement that $\ell(\cdot)$ be a least-squares cost was relaxed. Because typically there is no requirement that $\ell(\cdot)$ be symmetric, such a disturbance may not be zero-mean, but it remains *zero-mode*, i.e., the most likely disturbance is $w = 0$. One could formulate FIE and MHE in which a disturbance forecast \mathbf{u}_f is given and the distance of w from u_f is penalized, rather than the distance of w from zero. However we restrict our attention to the zero mode case for simplicity.

In order to ensure that $\mathbb{P}_k(\cdot)$ has a solution, we require some assumptions.

Assumption 5.2 (Continuity). The functions $\ell_x : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, $\ell : \mathbb{W} \times \mathbb{R}^p \rightarrow \mathbb{R}_{\geq 0}$,

$f : \mathbb{X} \times \mathbb{W} \rightarrow \mathbb{X}$, and $h : \mathbb{X} \rightarrow \mathbb{Y}$ are continuous, and the sets \mathbb{X} and \mathbb{W} are closed.

Assumption 5.3 (Positive-definite cost functions). There exist $\alpha \in \mathcal{K}_\infty$ such that

$$\begin{aligned} \underline{\alpha}_x(|\chi(0) - \bar{x}|) &\leq \ell_x(\chi(0), \bar{x}) \leq \bar{\alpha}_x(|\chi(0) - \bar{x}|) \\ \underline{\alpha}_\ell(|(\omega, \nu)|) &\leq \ell(\omega, \nu) \leq \bar{\alpha}_\ell(|(\omega, \nu)|) \end{aligned}$$

for all $\chi(0) \in \mathbb{R}^n$ and all $\omega \in \mathbb{R}^g$ and $\nu \in \mathbb{R}^p$.

Assumptions 5.2 and 5.3 guarantee that a solution to $\mathbb{P}_k(\cdot)$ exists for all finite k (Rawlings and Ji, 2012). Next, we need stabilizability and detectability assumptions. However, the assumptions that are eventually used to demonstrate robust stability of FIE would appear rather unmotivated at this point in exposition. Therefore, we begin by demonstrating nominal stability of FIE by a process that is a natural sequel to that used in Chapter 4. Afterwards, the problems that occur when we attempt to apply this procedure to the estimation problem with disturbances can help motivate these assumptions.

5.2.1 Nominal Stability of FIE

Consider an output trajectory \mathbf{y} that is generated by a disturbance-free system

$$x^+ = f(x, 0) \quad y = h(x), \quad (5.2)$$

i.e., one in which $\mathbf{w} = \mathbf{0}$ and $\mathbf{v} = \mathbf{0}$. Because we do not consider a disturbance forecast in FIE or MHE, we define a class of state estimators more restrictive than Definition 2.1.

Definition 5.4 (State estimator). A state estimator is a sequence of maps $\psi_k : \mathbb{X} \times (\mathbb{Y})^k \rightarrow \mathbb{X}$ defined for all $k \in \mathbb{I}_{\geq 0}$. The resulting state estimate

$$\hat{x}(k) = \psi_k(\bar{x}, \mathbf{y}(0 : k - 1))$$

is defined for only k .

FIE, of course, gives us smoothed estimates $\hat{x}(j|k)$ in addition to the final estimate $\hat{x}(k)$, but many other methods of state estimation give only the final estimate.

Definition 5.5 (Globally stable estimator). A state estimator (ψ_k) is globally stable if there exists $\beta \in \mathcal{KL}$ such that for all $x(0), \bar{x} \in \mathbb{X}$ and \mathbf{y} generated by (5.2), we have that

$$|x(k) - \hat{x}(k)| \leq \beta(|x(0) - \bar{x}|, k)$$

for all $k \in \mathbb{R}_{\geq 0}$.

Next, in line with the procedures from Chapters 3 and 4, we define an infinite horizon estimation problem

$$\begin{aligned} \mathbb{P}_\infty(\bar{x}, \mathbf{y}) : \quad & \min_{\chi(0) \in \mathbb{X}, \omega \in \mathbb{W}^\infty, \nu \in (\mathbb{R}^p)^\infty} V_\infty(x, \omega, \nu, \bar{x}, \mathbf{y}) := \ell_x(\chi(0), \bar{x}) + \sum_{i=0}^{\infty} \ell(\omega(i), \nu(i)) \\ & \text{subject to } \chi^+ = f(\chi, \omega) \\ & y = h(\chi) + \nu. \end{aligned}$$

Unlike in Chapters 3 and 4, we do not need to assume stabilizability in order to guarantee that this problem has a solution. The choice of $\chi(0) = x(0)$, $\omega = \mathbf{0}$, and $\nu = \mathbf{0}$ is a feasible point with a finite cost $\ell_x(x(0), \bar{x})$. By applying Assumptions 5.2 and 5.3, we note that all feasible sequences that are less costly than this one lie in a compact subset of $\mathbb{X} \times \mathbb{W}^\infty \times (\mathbb{R}^p)^\infty$. Therefore, we can apply (Keerthi and Gilbert,

1985, Thm. 1) to show that the infimum $V_\infty^0(\bar{x}, \mathbf{y})$ is attained by $\hat{x}(0|\infty)$, $\widehat{\mathbf{w}}(\infty)$, and $\widehat{\mathbf{v}}(\infty)$.

In order to proceed further, however, we do need stabilizability and detectability assumptions. As in Chapter 4, we assume the existence of an i-IOSS Lyapunov function. Unlike in Chapter 4, however, we do not assume that it exponentially decreases with respect to a semimetric. Instead, we use a conventional \mathcal{K} function definition.

Assumption 5.6 (i-IOSS Lyapunov function). There exist $\Lambda : \mathbb{X}^n \times \mathbb{X}^n \rightarrow \mathbb{R}_{\geq 0}$, $\alpha_i \in \mathcal{K}_\infty$, and $\sigma_w, \sigma_y \in \mathcal{K}$ such that

$$\alpha_1(|x_1 - x_2|) \leq \Lambda(x_1, x_2) \leq \alpha_2(|x_1 - x_2|) \quad (5.3)$$

$$\begin{aligned} \Lambda(f(x_1, w_1), f(x_2, w_2)) &\leq \Lambda(x_1, x_2) - \alpha_3(|x_1 - x_2|) \\ &\quad + \sigma_w(|w_1 - w_2|) + \sigma_y(|h(x_1) - h(x_2)|) \end{aligned} \quad (5.4)$$

for any $x_1, x_2 \in \mathbb{X}^n$ and $w_1, w_2 \in \mathbb{W}^g$.

Note that here i-IOSS is defined in terms of $h(x_1) - h(x_2)$ rather than $y_1 - y_2$. This change is necessary because of the measurement noise term included in (5.1) that is absent in (4.1). As in Chapter 4, we require a compatibility condition between $\Lambda(\cdot)$ and $\ell(\cdot)$. We note that, in this case of nominal stability, we have that $h(x) = h(\hat{x}) + \hat{v}$ (because $v = 0$), so $\ell(\hat{w}, \hat{v}) = \ell(\hat{w}, h(x) - h(\hat{x}))$.

Assumption 5.7. We have that

$$\ell(\omega, \nu) \geq \sigma_w(|\omega|) + \sigma_y(|\nu|)$$

for all $\omega \in \mathbb{W}$ and $\nu \in \mathbb{R}^p$.

Finally, we require an incremental stabilizability assumption, which we also formulate in terms of general \mathcal{K} functions, rather than in terms of a semimetric as in Chapter 4.

Assumption 5.8. There exists $\check{\alpha} \in \mathcal{K}_\infty$ such that, for every $\check{x}, x_r \in \mathbb{X}$ and $\mathbf{w}_r \in \mathbb{W}^\infty$, there exists some $\check{\mathbf{w}} \in \mathbb{W}^\infty$ such that

$$\sum_{k=0}^{\infty} \ell(\dot{w}(k) - w(k), y(k) - \check{y}(k)) \leq \check{\alpha}(|\check{x} - x_r|), \quad (5.5)$$

in which

$$\begin{aligned} x^+ &= f(x, w) & y &= h(x) \\ \check{x}^+ &= f(\check{x}, \dot{w}) & \check{y} &= h(\check{x}). \end{aligned}$$

Remark 5.9. In (5.5), the reversal of the order of subtraction between the disturbances w and outputs y is intentional.

With these assumptions, we can begin constructing a Q function. Let

$$V^0(j|k) := \ell_x(\hat{x}(0|k), \bar{x}) + \sum_{i=0}^{j-1} \ell(\dot{w}(i|k), \dot{v}(i|k))$$

for all $j \in \mathbb{I}_{0:k}$. Note that $V^0(0|k) = \ell_x(\hat{x}(0|k), \bar{x})$. Next, let

$$Z(j|k) := V_\infty^0(\bar{x}, \mathbf{y}) - V^0(j|k),$$

for all $j \in \mathbb{I}_{0:k}$, in which, as is routine, the dependence of $V_\infty^0(\cdot)$ on \bar{x} and \mathbf{y} is suppressed for brevity. By positivity of $\ell(\cdot)$, we have that $V^0(j|k) \leq V^0(k|k)$ and, because $\hat{x}(0|\infty)$, $\widehat{\mathbf{w}}(\infty)$, and $\widehat{\mathbf{v}}(\infty)$ are feasible in $\mathbb{P}_k(\cdot)$, we have that $V^0(k|k) = V_k^0(\bar{x}, \mathbf{y}) \leq V_\infty^0(\bar{x}, \mathbf{y})$. Thus, $Z(j|k) \geq 0$ for all $j \in \mathbb{I}_{0:k}$ and for all $k \in \mathbb{I}_{\geq 0}$. Furthermore, we have that

$$\begin{aligned}
V_{\infty}^0(\bar{\mathbf{x}}, \mathbf{y}) &\leq V^0(j|k) + \min_{\omega \in \mathbb{W}^{\infty}, \nu \in (\mathbb{R}^p)^{\infty}} \sum_{i=j}^{\infty} \ell(\omega(i), \nu(i)) \\
\text{subject to} \quad &\chi^+ = f(\chi, \omega) \\
&y = h(\chi) + \nu \\
&\chi(j) = \hat{x}(j|k)
\end{aligned} \tag{5.6}$$

because $\hat{x}(0|k)$, $\widehat{\mathbf{w}}(0 : j|k)$, and $\widehat{\mathbf{v}}(0 : j|k)$ are feasible in $\mathbb{P}_{\infty}(\cdot)$. By applying Assumption 5.8 with $\check{x} = \hat{x}(j|k)$, $x_r = x(j)$, and $\mathbf{w}_r = \mathbf{0}$, we obtain $\check{\mathbf{w}} \in \mathbb{W}^{\infty}$ such that

$$\sum_{i=j}^{\infty} \ell(\check{\omega}(i), y(i) - \check{y}(i)) \leq \check{\alpha}(|\hat{x}(j|k) - x(j)|).$$

If we let $\check{\mathbf{v}}(j : \infty) = \mathbf{y}(j : \infty) - \check{\mathbf{y}}(j : \infty)$, we obtain feasible decision variables in (5.6). Therefore, we have that

$$V_{\infty}^0(\bar{\mathbf{x}}, \mathbf{y}) \leq V^0(j|k) + \check{\alpha}(|e(j|k)|).$$

and therefore

$$Z(j|k) \leq \check{\alpha}(|e(j|k)|).$$

Finally, we have that

$$\begin{aligned}
Z(j+1|k) &= V_{\infty}^0(\bar{\mathbf{x}}, \mathbf{y}) - V^0(j+1|k) \\
&= V_{\infty}^0(\bar{\mathbf{x}}, \mathbf{y}) - V^0(j|k) - \ell(\hat{\omega}(j|k), \hat{v}(j|k)) \\
&= Z(j|k) - \ell(\hat{\omega}(j|k), \hat{v}(j|k))
\end{aligned}$$

for all $j \in \mathbb{I}_{0:k-1}$ and all $k \in \mathbb{I}_{\geq 0}$.

As before, we must add the i-IOSS Lyapunov function $\lambda(\cdot)$ to $Z(\cdot)$ to make it positive definite. Thus, let

$$Q(j|k) := Z(j|k) + \lambda(\hat{x}(j|k), x(j))$$

for all $j \in \mathbb{I}_{0:k}$ and all $k \in \mathbb{I}_{\geq 0}$. We can then apply (5.3) to obtain.

$$\begin{aligned} \alpha_1(|e(j|k)|) &\leq Q(j|k) \leq \alpha_2(|e(j|k)|) + \check{\alpha}(|e(j|k)|) \\ &:= \check{\alpha}_2(|e(j|k)|) \end{aligned}$$

for all $j \in \mathbb{I}_{0:k}$ and all $k \in \mathbb{I}_{\geq 0}$. Furthermore, we can apply (5.4) to obtain

$$\begin{aligned} Q(j+1|k) &= \lambda(\hat{x}(j+1|k), x(j+1)) + Z(j+1|k) \\ &\leq \lambda(\hat{x}(j|k), x(j)) - \alpha_3(|e(j|k)|) + \sigma_w(|\hat{w}(j|k) - 0|) \\ &\quad + \sigma_y(|h(\hat{x}(j|k)) - h(x(j))|) + Z(j|k) - \ell(\hat{w}(j|k), \hat{v}(j|k)) \\ &= Q(j|k) - \alpha_3(|e(j|k)|) \\ &\quad + \sigma_w(|\hat{w}(j|k)|) + \sigma_y(|\hat{v}(j|k)|) - \ell(\hat{w}(j|k), \hat{v}(j|k)) \end{aligned}$$

for all $j \in \mathbb{I}_{0:k-1}$ and all $k \in \mathbb{I}_{\geq 0}$. We can then apply Assumption 5.7 to obtain

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(|e(j|k)|)$$

for all $j \in \mathbb{I}_{0:k-1}$ and all $k \in \mathbb{I}_{\geq 0}$. To summarize, we have that

$$\alpha_1(|e(j|k)|) \leq Q(j|k) \leq \check{\alpha}_2(|e(j|k)|) \quad \forall j \in \mathbb{I}_{0:k} \quad (5.7)$$

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(|e(j|k)|) \quad \forall j \in \mathbb{I}_{0:k-1}. \quad (5.8)$$

We can easily make the standard argument, combining (5.7) with (5.8) to obtain

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(\check{\alpha}_2^{-1}(Q(j|k)))$$

for all $j \in \mathbb{I}_{0:k-1}$. By a standard construction, as in (Rawlings et al., 2017, Theorem B.15), there exists $\sigma \in \mathcal{K}_\infty$ such that $\sigma(s) \leq s$ for all $s > 0$ and

$$Q(j|k) - \alpha_3(\check{\alpha}_2^{-1}(Q(j|k))) \leq \sigma(Q(j|k))$$

for all $j \in \mathbb{I}_{0:k-1}$. As a result, we have that

$$Q(k|k) \leq \sigma^k(Q(0|k)) \quad (5.9)$$

for all $k \in \mathbb{I}_{\geq 0}$. In earlier chapters, this would be the point where we switched from analyzing the behavior of open-loop optimal cost functions using a Q function to analyzing the behavior of closed-loop systems using a Lyapunov function. However, in the pure estimation problem, the loop is never closed. In FIE, the estimator does not even have an internal state that is changing. Therefore, we proceed differently.

We can combine (5.7) and (5.9) to obtain

$$\alpha_1(|e(k|k)|) \leq \sigma^k \circ \check{\alpha}_2(|e(0|k)|) = \sigma^k \circ \check{\alpha}_2(|x(0) - \hat{x}(0|k)|).$$

However, this does not give us global stability, because $\hat{x}(0|k)$ is estimated again at every time k . We need to obtain an upper bound on $Q(0|k)$ in terms of \bar{e} . Recall that

$$Q(0|k) = Z(0|k) + \lambda(\hat{x}(0|k), x(0)) = V_{\infty}^0(\bar{x}, \mathbf{y}) - \ell_x(\hat{x}(0|k), \bar{x}) + \lambda(x(0), \hat{x}(0|k)).$$

Because $\ell_x(\cdot)$ is nonnegative and $\lambda(x(0), \hat{x}(0|k)) \leq \alpha_2(|e(0|k)|)$, we have that

$$Q(0|k) \leq V_{\infty}^0(\bar{x}, \mathbf{y}) + \alpha_2(|e(0|k)|). \quad (5.10)$$

We next have that

$$|x(0) - \hat{x}(0|k)| \leq |\bar{x} - \hat{x}(0|k)| + |x(0) - \bar{x}|.$$

Furthermore, we have that

$$\begin{aligned} |\bar{x} - \hat{x}(0|k)| &\leq \underline{\alpha}_x^{-1}(\ell_x(\hat{x}(0|k), \bar{x})) \\ &\leq \underline{\alpha}_x^{-1}(V^0(k|k)) \end{aligned}$$

by Assumption 5.3. Because setting $\chi(0) = x(0)$ generates a feasible sequence with $\widehat{\mathbf{w}} = \mathbf{0}$ and $\widehat{\mathbf{v}} = \mathbf{0}$ we have that

$$V^0(k|k) \leq \ell_x(x(0), \bar{x}) = \bar{\alpha}_x(|\bar{e}|).$$

Combining these equations, we obtain

$$|e(0|k)| \leq |\bar{e}| + \underline{\alpha}_x^{-1} \circ \bar{\alpha}_x(|\bar{e}|).$$

Finally, we have that

$$V_\infty^0(\bar{x}, \mathbf{y}) \leq \bar{\alpha}_x(|\bar{e}|)$$

because the sequence generated by $x(0)$ is feasible. Thus, we can combine these bounds with (5.10) to obtain

$$\begin{aligned} Q(0|k) &\leq \bar{\alpha}_x(|\bar{e}|) + \alpha_2(|\bar{e}| + \underline{\alpha}_x^{-1} \circ \bar{\alpha}_x(|\bar{e}|)) \\ &:= \bar{\alpha}_0(|\bar{e}|) \end{aligned} \tag{5.11}$$

for all $k \in \mathbb{I}_{\geq 0}$. This is a bound on $Q(0|k)$ in terms of only \bar{e} . Finally, we have that

$$|e(k|k)| \leq \alpha_1^{-1}(Q(k|k)) \tag{5.12}$$

for all $k \in \mathbb{I}_{\geq 0}$. Combining (5.9), (5.11), and (5.12), we have that

$$|e(k|k)| \leq \alpha_1^{-1}(\sigma^k(\bar{\alpha}_0(|\bar{x} - x(0)|))) := \beta(|\bar{e}|, k)$$

for all $k \in \mathbb{I}_{\geq 0}$. Because we have that $\sigma(s) < s$ for all $s > 0$, we have that $\beta \in \mathcal{KL}$ and thus FIE is globally asymptotically stable.

Remark 5.10. This argument finds its origin in (Muske et al., 1993, Thms. 2& 3), which first provided a proof that $\lim_{k \rightarrow \infty} V_k^0$ exists for linear constrained FIE. This proof was extended to nonlinear FIE in (Muske and Rawlings, 1995). Each of these works attempts to show convergence of FIE, but only provides proof that $\hat{w}(k-1|k)$ and $\hat{v}(k-1|k)$ (or $\hat{v}(k|k)$ if the last measurement is included) converge to zero. They

then asserted that this fact results in estimator convergence and it is unknown how they would have attempted to complete the proof if pressed for more detail. They do not attempt to show Lyapunov stability of FIE (though Muske et al. (1993) claimed to show that FIE is “asymptotically stable” without ever attempting to show Lyapunov stability).

Some details were filled in by Rao (2000, Prop. 3.3.3). In particular, a proof is provided that, for observable systems, that the convergence of the final N_0 stage costs, in which N_0 is the index of observability, guarantees convergence of $\hat{x}(k|k)$ to $x(k)$. However, he does not show that this fixed-lag sum of stage costs does in fact converge for FIE, showing only that $\ell(\hat{w}(k-1|k), \hat{v}(k-1|k))$ converges to zero. He also makes an attempt to show that FIE is Lyapunov stable, but commits an implicit quantifier shift in his proof.

The next step forward was provided by Rawlings and Mayne (2009, Thm. 4.8), where the argument was extended from observable systems to detectable (i-IOSS) systems, and from nominal stability to robust stability for disturbances that converge sufficiently fast that $\ell(w(k), v(k))$ can be summed. A sound proof of Lyapunov stability of FIE is given, but, again, a proof that $\hat{w}(k-1|k)$ and $\hat{v}(k-1|k)$ converge to zero is supposed to suffice for proof that $\hat{x}(k|k)$ converges to $x(k)$. A better proof was given by Rawlings and Ji (2012), where it is shown that fixed-lag sums of stage costs do converge to zero. However, they require Lyapunov stability for the fixed-point estimates $\hat{x}(j|k)$, but demonstrate it only for the final estimate $\hat{x}(k|k)$. This minor issue is addressed in (Rawlings et al., 2017, Prop. 4.9). A larger issue is the inadequate definition of \mathcal{KL} stability given by (Rawlings and Ji, 2012) that does

not even guarantee Lyapunov stability.

At this point I made my personal contribution to this literature in (Allan and Rawlings, 2019b) by both adding an explicit stabilizability assumption similar to Assumption 5.8 (though (Allan and Rawlings, 2019b, Asm. 11) contains a sign error), and by introducing the notion of the Q function. A function similar to $Z(\cdot)$ was introduced in material ultimately cut from (Rawlings and Mayne, 2009), but because there was not yet any characterization of detectability through a storage function $\Lambda(\cdot)$, it was used in a fashion similar to an N -step Lyapunov function.

None of this history is to blame these past authors for being sloppy, but rather to explain why (Allan and Rawlings, 2019b) was a necessary contribution when the problem had supposedly been resolved in (Rawlings and Ji, 2012) or even in (Muske and Rawlings, 1995). The fact that the state history is reconstructed in each estimation problem is rather confusing to researchers new to the field, and naive application of mathematical intuition calibrated for even a problem as similar as MPC can easily go awry. Short of using proof-checking software, it is impossible to be explicit in every way in mathematics, but a well-formed intuition is needed to know when an explicit proof is necessary and when it is not.

5.2.2 From stability to robustness

Now that we have a proof of nominal stability, it is worth looking at an example about what sort of behavior can happen when there are persistent disturbances.

Consider the system

$$x^+ = e^{-1}x + w \quad y = \min(x, 1) + v.$$

As the state evolution equation is both stable and linear, it is incrementally ISS and thus i-IOSS. Furthermore, because it is Lipschitz with an additive stage cost, it is incrementally stabilizable as well. We then use the stage costs

$$\ell_x(\chi, \bar{x}) = \ln(1 + (\chi - \bar{x})^2) \quad \ell(\omega, \nu) = 4|\omega| + 4|\nu|.$$

While this choice of a prior weighting may raise eyebrows, it satisfies Assumption 5.3, and $\ell(\cdot)$ satisfies Assumption 5.7 with the i-IOSS Lyapunov function

$$\Lambda(x_1, x_2) = \sum_{k=0}^{\infty} |x_1(k) - x_2(k)| \quad x_1^+ = e^{-1}x_1 \quad x_2^+ = e^{-1}x_2.$$

This system satisfies all the assumptions required for nominal stability. Let $x(0) = \bar{x} = 0$, $\mathbf{w} = \mathbf{0}$, and $\mathbf{v} = (1, 1, 1, \dots)$. Using the true disturbance sequence in $\mathbb{P}_k(\cdot)$ results in a cost of $4k$. However, the choice of $\chi(0) = e^{k-1}$, $\omega = \mathbf{0}$, and $\nu = \mathbf{0}$ reproduces $\mathbf{y}(0 : k-1)$ as well, but with a cost of approximately $2(k-1)$. Because the L^1 penalties used in $\ell(\cdot)$, there is no benefit in attempting to use a mixture of noises and initial states in a solution, so in the optimal solution, $\lim_{k \rightarrow \infty} e(0|k) = \infty$. Nevertheless, $e(k|k) = 1$ for all k .

The lesson that this example illustrates is that, with an infinite amount of disturbance “energy” being injected into the system, that can potentially result in extremely large disturbance estimates or prior errors. Ji et al. (2016) explicitly altered the optimization by addition of a max-term in order to prevent enormous disturbances (the problem they formulated could still have enormous initial conditions). Nevertheless, the problem can have bounded error in the estimate $e(k|k)$ if care is taken with the assumptions.

It should be pointed out that this problem with smoothed estimates does not exist with the Kalman filter—smoothed estimates have an exponentially decreasing dependence on future data (Anderson and Moore, 2012, Sec. 7.2), so the amount of error that bounded future noise can induce in filtered estimates is bounded. However, it is unclear what sort of assumptions are necessary to replicate that behavior for nonlinear systems.

5.3 Robust global asymptotic stability of FIE

Recall the definition of robust global asymptotic stability (RGAS) from Chapter 2.

Definition 5.11 (Robust global asymptotic stability). A state estimator (ψ_k) is robustly globally asymptotically stable if there exist $\beta_x, \beta_d \in \mathcal{KL}$ such that

$$|e(k|k)| \leq \beta_x(|\bar{e}|, k) \oplus \max_{j \in \mathbb{I}_{0,k-1}} \beta_d(|(w(j), v(j))|, k - j - 1)$$

for all $k \in \mathbb{I}_{\geq 0}$, $\bar{x} \in \mathbb{X}$, $\mathbf{w} \in \mathbb{W}^\infty$, and $\mathbf{v} \in (\mathbb{R}^p)^\infty$.

Remark 5.12. The disturbances w and v are combined in a single term here for convenience. In general the form of RGAS presented here is equivalent to that presented in Chapter 2 in which the disturbances are separated into different terms.

This definition differs from that proposed in (Rawlings and Ji, 2012) in an important respect beyond simply having the correct quantifier order. The definition used in (Rawlings and Ji, 2012) uses the upper bound

$$|e(k|k)| \leq \beta_x(|\bar{e}|, k) \oplus \max_{j \in \mathbb{I}_{0,k-1}} \gamma_d(|(w(j), v(j))|)$$

with an asymptotic gain $\gamma_d \in \mathcal{K}$. This form of bound was proposed because of its similarity to ISS. However, ISS contains the additional implication of convergence with converging disturbances, because the upper bound can be repeatedly applied at different times. For a state estimator, on the other hand, the initial state plays a special role, and the bound cannot be repeatedly applied. That is how (Ji et al., 2016) and (Hu, 2017) were able to prove RGAS of FIE containing a max-term without the additional implication of convergence in the case of converging disturbances. This new definition does not have this problem.

One of the key insights that has enabled us to proceed thus far is the existence of a well-defined infinite horizon problem with which we can compare the finite-horizon problems. However, when considering the case of persistent disturbances (as opposed to the cases of disturbances that converge sufficiently fast in (Rawlings and Ji, 2012)), there is no obvious candidate infinite horizon problem. There typically does not exist any feasible estimate sequences that attain a finite value in $\mathbb{P}_\infty(\cdot)$.

However, we can generate a sequence of infinite-horizon problems, each containing the first k disturbances, with zero disturbances thereafter. Each of these problems have finite disturbance sequences, and, as a result, have well-defined solutions. As a result, we first define an auxiliary sequence of outputs

$$\begin{aligned}\tilde{x}(0|k) &= x(0) \\ \tilde{x}(j+1|k) &= \begin{cases} f(\tilde{x}(j|k), w) & \text{if } j \in \mathbb{I}_{0:k-1} \\ f(\tilde{x}(j|k), 0) & \text{else} \end{cases} \\ \tilde{y}(j|k) &= \begin{cases} h(\tilde{x}(j|k)) + v & \text{if } j \in \mathbb{I}_{0:k-1} \\ h(\tilde{x}(j|k)) & \text{else} \end{cases}\end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$. Note that $\tilde{\mathbf{y}}(0 : k - 1) = \mathbf{y}(0 : k - 1)$. Now, we can define a sequence of infinite-horizon costs

$$\begin{aligned} \tilde{\mathbb{P}}_{\infty}(k, x, \bar{x}, \mathbf{w}, \mathbf{v}) : \quad \tilde{V}_{\infty}^0(k, x, \bar{x}, \mathbf{w}, \mathbf{v}) := & \min_{\chi(0), \omega, \nu} \ell_x(\chi_0, \bar{x}) + \sum_{i=0}^{\infty} \ell(\omega(i), \nu(i)) \\ & \text{subject to } \chi^+ = f(\chi, \omega) \\ & \tilde{\mathbf{y}}(i|k) = h(\chi(i)) + \nu \\ & \chi(0) \in \mathbb{X}, \omega \in \mathbb{W}^{\infty}, \nu \in (\mathbb{R}^p)^{\infty} \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$. For brevity, we refer to $\tilde{\mathbb{P}}_{\infty}(k, x, \bar{x}, \mathbf{w}, \mathbf{v})$ as $\tilde{\mathbb{P}}_{\infty}(k)$ and $\tilde{V}_{\infty}^0(k, x, \bar{x}, \mathbf{w}, \mathbf{v})$ as $\tilde{V}_{\infty}^0(k)$ when that is unambiguous. By construction of $\tilde{\mathbb{P}}(k)$, we have that $\mathbf{w}(0 : k - 1) \curvearrowright \mathbf{0}$ and $\mathbf{v}(0 : k - 1) \curvearrowright \mathbf{0}$ constitute feasible sequences. Therefore, we have that

$$\tilde{V}_{\infty}^0(k) \leq \ell_x(x(0), \bar{x}) + \sum_{i=0}^{k-1} \ell(w(i), v(i)). \quad (5.13)$$

Because $\tilde{V}_{\infty}^0(k)$ is bounded above, the existence theorem from Keerthi and Gilbert (1985) can be applied, and there exist $\hat{\mathbf{x}}(\infty|k)$, $\hat{\mathbf{w}}(\infty|k)$, and $\hat{\mathbf{v}}(\infty|k)$ for which the infimum is attained.

With an infinite horizon cost function defined, we can now define

$$Z(j|k) := \tilde{V}_{\infty}^0(k) - V^0(j|k) \quad Q(j|k) := Z(j|k) + \Lambda(\hat{x}(j|k), x(j))$$

for all $j \leq k \in \mathbb{I}_{\geq 0}$. We now need to derive equations similar to (5.7), (5.8), and (5.11) in this context. By (5.3), we immediately have that

$$Q(j|k) \geq \alpha_1(|e(j|k)|) \quad (5.14)$$

but the other equations require significantly more work. Before continuing, we recall the fact that, for a \mathcal{K} function $\sigma(\cdot)$, we have that

$$\sigma(a + b) \leq \sigma(\max(2a, 2b)) = \max(\sigma(2a), \sigma(2b)) \leq \sigma(2a) + \sigma(2b) \quad (5.15)$$

for any $a, b \in \mathbb{R}_{\geq 0}$.

We next derive an upper bound for $Q(0|k)$ in terms of \bar{e} , because it helps motivate the form an upper bound for $Q(j|k)$ takes.

Combining Assumption 5.3 and (5.13) gives us the upper bound

$$Z(0|k) \leq \ell_x(x(0), \bar{x}) + \sum_{i=0}^{k-1} \ell(w(i), v(i)) \leq \bar{\alpha}_x(|\bar{e}|) + \sum_{i=0}^{k-1} \bar{\alpha}_\ell(|(w(i), v(i))|) \quad (5.16)$$

which is as much as we can expect, since \mathbf{w} and \mathbf{v} are nonzero. One might think that finding an upper bound for $\lambda(x(0), \hat{x}(0|k))$ would be straightforward. From (5.3), we have that

$$\lambda(x(0), \hat{x}(0|k)) \leq \alpha_2(|e(0|k)|)$$

which is indeed straightforward. However, this upper bound is in terms of the error in the smoothed estimate of the initial state $e(0|k)$ rather than \bar{e} . As illustrated in Section 5.2.2, $e(0|k)$ might grow without bound in the case of bounded disturbances. As in the nominal case, we have that

$$|\hat{x}(0|k) - x(0)| \leq |\hat{x}(0|k) - \bar{x}| + |\bar{x} - x(0)|. \quad (5.17)$$

and that

$$\begin{aligned} |\hat{x}(0|k) - \bar{x}| &\leq \underline{\alpha}_x^{-1}(\ell_x(\hat{x}(0|k) - \bar{x})) \\ &\leq \underline{\alpha}_x^{-1}(V^0(k|k)). \end{aligned}$$

Because $V^0(k|k) \leq \tilde{V}_\infty^0(k)$ we can combine this equation with (5.13) and (5.17) to obtain

$$|\hat{x}(0|k) - x(0)| \leq |\bar{x} - x(0)| + \underline{\alpha}_x^{-1} \left(\bar{\alpha}_x (|x(0) - \bar{x}|) + \sum_{i=0}^{k-1} \bar{\alpha}_\ell (|\boldsymbol{w}(i), \boldsymbol{v}(i)|) \right).$$

As a result, we have that

$$\begin{aligned} \Lambda(\hat{x}(0|k), x(0)) &\leq \alpha_2 (|\hat{x}(0|k) - \bar{x}|) \\ &\leq \alpha_2 \left(|\bar{x} - x(0)| + \underline{\alpha}_x^{-1} \left(\bar{\alpha}_x (|x(0) - \bar{x}|) + \sum_{i=0}^{k-1} \bar{\alpha}_\ell (|\boldsymbol{w}(i), \boldsymbol{v}(i)|) \right) \right) \\ &\leq \alpha_2 (2 |\bar{x} - x(0)|) \\ &\quad + \alpha_2 \circ 2 \underline{\alpha}_x^{-1} \left(\bar{\alpha}_x (|x(0) - \bar{x}|) + \sum_{i=0}^{k-1} \bar{\alpha}_\ell (|\boldsymbol{w}(i), \boldsymbol{v}(i)|) \right) \\ &\leq \alpha_2 (2 |\bar{x} - x(0)|) \\ &\quad + \alpha_2 \circ 2 \underline{\alpha}_x^{-1} \left(\bar{\alpha}_x (|x(0) - \bar{x}|) + \sum_{i=0}^{k-1} \bar{\alpha}_\ell (|\boldsymbol{w}(i), \boldsymbol{v}(i)|) \right), \end{aligned}$$

in which the last step follows from (5.15). The fact that the sum of disturbances occurs within a composition of \mathcal{K} functions is troublesome. As we will see, a growth rate of this upper bound that is linear in k is manageable, but a faster growth rate may not be. If we had that $2\underline{\alpha}_x^{-1}(s) \leq \alpha_2^{-1}(s)$ the two \mathcal{K} functions could be eliminated. That is the case when

$$\alpha_2(s) \leq \underline{\alpha}_x((1/2)s) \tag{5.18}$$

which constitutes a compatibility assumption between $\ell_x(\cdot)$ and $\Lambda(\cdot)$. If (5.18) held, then we would have that

$$\Lambda(\hat{x}(0|k), x(0)) \leq \alpha_2 (2 |\bar{x} - x(0)|) + \bar{\alpha}_x (|x(0) - \bar{x}|) + 2 \sum_{i=0}^{k-1} \bar{\alpha}_\ell (|\boldsymbol{w}(i), \boldsymbol{v}(i)|)$$

and, once combined with (5.16), would yield

$$\begin{aligned}
Q(0|k) &\leq \alpha_2(2|\bar{x} - x(0)|) + 2\bar{\alpha}_x(|x(0) - \bar{x}|) + 2 \sum_{i=0}^{k-1} \bar{\alpha}_\ell(|(w(i), v(i))|) \\
&= \alpha_0(|\bar{x} - x(0)|) + 2 \sum_{i=0}^{k-1} \bar{\alpha}_\ell(|(w(i), v(i))|), \tag{5.19}
\end{aligned}$$

in which $\alpha_0(s) := \alpha_2(2s) + 2\bar{\alpha}_x(s)$.

In addition to a compatibility assumption like (5.18), we need a compatibility assumption like Assumption 5.7 between $\ell(\cdot)$. Unfortunately, Assumption 5.7 is not in of itself adequate. From the properties of \mathbb{P}_k , we have no direct upper bound on $|\hat{w}(j|k) - w(j)|$ and $|h(\hat{x}(j|k)) - h(x(j))| = |v(j) - \hat{v}(j|k)|$, but only on $|\hat{w}(j|k)|$ and $|\hat{v}(j|k)|$. Therefore, we need to use the triangle inequality to create a dissipation condition

$$\begin{aligned}
\Lambda(\hat{x}(j+1|k), x(j+1)) &\leq \Lambda(\hat{x}(j|k), x(j)) - \alpha_3(|e(j|k)|) \\
&\quad + \sigma_w(|\hat{w}(j|k) - w(j)|) + \sigma_y(|v(j) - \hat{v}(j|k)|) \\
&\leq \Lambda(\hat{x}(j|k), x(j)) - \alpha_3(|e(j|k)|) \\
&\quad + \sigma_w(|\hat{w}(j|k)| + |w(j)|) + \sigma_y(|\hat{v}(j|k)| + |v(j)|) \\
&\leq \Lambda(\hat{x}(j|k), x) - \alpha_3(|\hat{x}(j|k) - x(j)|) \\
&\quad + \sigma_w(2|\hat{w}(j|k)|) + \sigma_w(2|w(j)|) \\
&\quad + \sigma_y(2|\hat{v}(j|k)|) + \sigma_y(2|v(j)|), \tag{5.20}
\end{aligned}$$

in which the last step follows from (5.15). We can then combine this equation with the standard descent condition

$$\begin{aligned}
Z(j+1|k) &= \tilde{V}_\infty^0(k) - V^0(j+1|k) \\
&= \tilde{V}_\infty^0(k) - V^0(j|k) - \ell(\hat{w}(j|k), \hat{v}(j|k)) \\
&= Z(j|k) - \ell(\hat{w}(j|k), \hat{v}(j|k)),
\end{aligned}$$

which is unchanged with the addition of persistent disturbances, in order to obtain

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(|e(j|k)| - \ell(\hat{w}(j|k), \hat{v}(j|k))) + \sigma_w(2|\hat{w}(j|k)|) + \sigma_w(2|w(j)|) \\ + \sigma_y(2|\hat{v}(j|k)|) + \sigma_y(2|v(j)|).$$

In this expression, we cannot expect to do anything about the terms involving w and v directly, but we still would like the terms involving \hat{w} and \hat{v} to be eliminated by the stage cost. This could be accomplished if we had that

$$\ell(\omega, \nu) \geq \underline{\alpha}_\ell(|(\omega, \nu)|) \geq \sigma_w(2|\omega|) + \sigma_y(2|\nu|), \quad (5.21)$$

in which case we would have

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(|e(j|k)|) + \sigma_w(2|w(j)|) + \sigma_y(2|v(j)|). \quad (5.22)$$

which is a dissipation inequality like would occur in an ISS Lyapunov function. We summarize (5.18) and (5.21) in an assumption.

Assumption 5.13 (Compatibility of $\ell_x(\cdot)$, $\ell(\cdot)$, and $\lambda(\cdot)$). We have that

$$\underline{\alpha}_x(s/2) \geq \alpha_2(s),$$

in which $\underline{\alpha}_x(\cdot)$ is from Assumption 5.3 and $\alpha_2(\cdot)$ is from Assumption 5.6, for all $s \geq 0$.

Furthermore, we have that

$$\underline{\alpha}_\ell(s) \geq \sigma_w(2s) + \sigma_y(2s),$$

in which $\underline{\alpha}_\ell(\cdot)$ is from Assumption 5.3 and $\sigma_w(\cdot)$ and $\sigma_y(\cdot)$ are from Assumption 5.6, for all $s \in \mathbb{R}_{\geq 0}$.

At this point, we have an upper bound for $Q(0|k)$, a lower bound for $Q(j|k)$, and a dissipation condition. All that remains is to find an upper bound for $Q(j|k)$. The

key to finding such an upper bound, both in earlier chapters and the case of nominal stability, has been a stabilizability assumption. We could apply Assumption 5.8 with $\check{x} = \hat{x}(j|k)$, $x_r = x(j)$, and $\mathbf{w}_r = \mathbf{w}(j : \infty)$ to obtain a sequence $\check{\mathbf{w}} \in \mathbb{W}^\infty$ such that

$$\sum_{i=j}^{\infty} \ell(\check{w}(i) - w(i), y(i) - \check{y}(i)) = \sum_{i=j}^{\infty} \ell(\check{w}(i) - w(i), \check{v}(i) - v(i)) \leq \check{\alpha}(|\hat{x}(j|k) - x(j)|).$$

However, this upper bound cannot be used directly for $\mathbb{P}_\infty(\cdot)$, because the estimation problem penalizes $\ell(\omega, \nu)$, not $\ell(\omega - w, \nu - v)$. Therefore, the stabilizability assumption used needs to be perturbed in a fashion similar to (5.20).

Assumption 5.14 (Perturbed Stabilizability). There exists $\bar{\gamma} \in \mathcal{K}$ such that, for every finite sequences $\mathbf{w}_r(k) \in \mathbb{W}^k$ and $\mathbf{v}_r(k) \in (\mathbb{R}^p)^k$ and any $\check{x}, x_r \in \mathbb{X}$, there exists $\check{\mathbf{w}} \in \mathbb{W}^\infty$ such that

$$\sum_{i=0}^{\infty} \ell(\check{w}(i), \check{v}(i)) \leq \check{\alpha}(|\check{x} - x_r|) + \sum_{i=0}^{k-1} \bar{\gamma}(|(w_r(k), v_r(k))|)$$

in which

$$\begin{aligned} \check{x}^+ &= f(\check{x}, \check{w}) & y_r &= h(\check{x}) + \check{v} \\ x_r^+ &= \begin{cases} f(x_r, w_r) & \text{for } i \in \mathbb{I}_{0:k-1} \\ f(x_r, 0) & \text{for } i \in \mathbb{I}_{k:\infty} \end{cases} & y_r &= \begin{cases} h(x_r) + v_r & \text{for } i \in \mathbb{I}_{0:k-1} \\ h(x_r) & \text{for } i \in \mathbb{I}_{k:\infty} \end{cases} \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$.

We can then use Assumption 5.14 to derive an upper bound for $Z(j|k)$ in terms of both $|e(j|k)|$ and a suitable function of the disturbances sequences \mathbf{w} and \mathbf{v} between times j and k . For $j \in \mathbb{I}_{0:k}$, we can apply Assumption 5.14 to obtain

$$\begin{aligned}
Z(j|k) &= V_\infty^0(k) - V^0(j|k) \\
&\leq V^0(j|k) - V^0(j|k) + \min_{\omega, \nu} \sum_{i=j}^{\infty} \ell(\omega(i), \nu(i)) \\
&\quad \text{subject to } \chi^+ = f(\chi, \omega) \\
&\quad \quad \quad \hat{y}(i|k) = h(\chi) + \nu \\
&\quad \quad \quad \chi(j) = \hat{x}(j|k) \\
&\leq \check{\alpha}(|e(j|k)|) + \sum_{i=j}^{k-1} \bar{\gamma}(|\omega(i), \nu(i)|).
\end{aligned}$$

which can be combined with (5.3) to obtain

$$\begin{aligned}
Q(j|k) &\leq \alpha_2(|e(j|k)|) + \check{\alpha}(|e(j|k)|) + \sum_{i=j}^{k-1} \bar{\gamma}(|\omega(i), \nu(i)|) \\
&:= \check{\alpha}_2(|e(j|k)|) + \sum_{i=j}^{k-1} \bar{\gamma}(|\omega(i), \nu(i)|). \tag{5.23}
\end{aligned}$$

We now have similar upper bounds for both $Q(0|k)$ and $Q(j|k)$. For brevity, it is convenient to abbreviate the summation terms. To that end, let $\bar{\gamma}_\ell(s) := \max(\bar{\gamma}(s), \bar{\alpha}_\ell(s))$, noting that $\bar{\gamma}_\ell \in \mathcal{K}_\infty$, and let

$$E_{j:k-1} := \sum_{i=j}^{k-1} \bar{\gamma}_\ell(|\omega(i), \nu(i)|), \tag{5.24}$$

which represents the amount of “energy” injected into the system by disturbances between times j and k . Note that

$$\sum_{i=j}^{k-1} \bar{\gamma}(|\omega(i), \nu(i)|) \leq E_{j:k-1} \quad \text{and} \quad \sum_{i=j}^{k-1} \bar{\alpha}_\ell(|\omega(i), \nu(i)|) \leq E_{j:k-1}.$$

We now summarize the bounds for $Q(\cdot)$ that we have derived, (5.14), (5.19), (5.22), and (5.23), in a single location.

$$Q(0|k) \leq \alpha_0(|\bar{e}|) + E_{0:k-1} \quad (5.25)$$

$$\alpha_1(|e(j|k)|) \leq Q(j|k) \leq \check{\alpha}_2(|e(j|k)|) + E_{j:k-1} \quad (5.26)$$

$$Q(j+1|k) \leq Q(j|k) - \alpha_3(|e(j|k)|) + \underline{\alpha}_\ell(|(w(j), v(j))|) \quad (5.27)$$

for all $j \in \mathbb{I}_{0:k-1}$. Apart from the upper bound, $Q(j|k)$ appears similar to an ISS Lyapunov function for $e(j|k)$.

5.3.1 Preliminary Results

If $Q(\cdot)$ behaved like a normal ISS Lyapunov function, it could be handled by conventional techniques (e.g., those used by Jiang and Wang (2001)). The total energy term $E_{j:k-1}$ appearing in the upper bounds, however, requires some new techniques to handle. In order to handle it, we need some preliminary results.

Proposition 5.15. *For every $\alpha \in \mathcal{K}_\infty$ such that*

$$0 \leq s - \alpha(s)$$

there exists some $\sigma \in \mathcal{K}_\infty$ such that

$$s - \alpha(s) \leq \sigma(s) < s$$

and that $\gamma(s) := s - \sigma(s)$ is \mathcal{K}_∞ .

Proof. Let

$$\text{semisat}(s) := \begin{cases} 0 & \text{if } s < 0 \\ s & \text{if } s \in [0, 1] \\ 1 & \text{if } s > 1 \end{cases} .$$

We define $\gamma(\cdot)$ as the solution to the ODE

$$\gamma' = (1/2) \text{semisat}(\alpha(s) - \gamma) \quad (5.28)$$

with initial condition $\gamma(0) = 0$. Note that the RHS is continuous in both s and γ , and that, for fixed s , we have that it is Lipschitz continuous in γ . Thus, by the Picard-Lindelöf Theorem (Coddington and Levinson, 1955, p. 12), we have that the ODE has a unique solution $\gamma(s)$ on the interval $[0, s^*)$ for some $s^* > 0$. Furthermore, because $\gamma(s)$ does not leave the domain of definition of the RHS, it can be extended to any $s > 0$ by (Coddington and Levinson, 1955, p. 15, Theorem 4.1), and thus is defined on $[0, \infty)$.

By construction, $\gamma(\cdot)$ is continuous and $\gamma(0) = 0$. Furthermore, because $\gamma'(s, \gamma) \in [0, 1/2]$ we have that $\gamma(\cdot)$ is nondecreasing. Suppose $\gamma(s)$ were bounded above. Then, because $\alpha \in \mathcal{K}_\infty$, there exists some s^* such that $\alpha(s) \geq \gamma(s) + 1$ for all $s > s^*$. Then, by (5.28), we have that $\gamma'(s) \geq (1/2)$ for all $s > s^*$. But then, $\gamma(s)$ is not bounded above—a contradiction. Thus $\lim_{s \rightarrow \infty} \gamma(s) = \infty$.

All that remains to show that $\gamma \in \mathcal{K}_\infty$ is to show that it is strictly increasing. If it failed to be strictly increasing, $\gamma'(\cdot)$ must be zero on some interval $[s_1, s_2]$ of nonzero measure. Because $\gamma'(s) = 0$ if $\alpha(s) = \gamma(s)$, we have that $\gamma(s) \leq \alpha(s)$ for all $s \in \mathbb{R}_{\geq 0}$. Then if there were some such interval, we would have $\alpha(s) = \gamma(s)$ for all $s \in [s_1, s_2]$. But because $\alpha \in \mathcal{K}_\infty$, it is strictly increasing. As a result, we would have

$$\gamma(s_2) = \alpha(s_2) > \alpha(s_1) = \gamma(s_1) = \gamma(s_2) ,$$

which is a contradiction. Thus $\gamma(\cdot)$ is strictly increasing and thus is \mathcal{K}_∞ .

Finally, we need only show that $\sigma(s) := s - \gamma(s)$ is \mathcal{K} . Note that $\gamma'(s) \leq (1/2)$ for all $s \in \mathbb{R}_{\geq 0}$. Thus we have that $\sigma'(s) = 1 - \gamma'(s) \geq 1/2$ for all $s \in \mathbb{R}_{\geq 0}$. Furthermore, $\sigma(0) = 0$ and $\sigma(\cdot)$ is continuous and unbounded above. Thus $\sigma \in \mathcal{K}_\infty$, which completes the proof. ■

Proposition 5.16. For $\sigma \in \mathcal{K}_\infty$ such that $s - \sigma(s) := \alpha(s) \in \mathcal{K}$, the inequality

$$\sigma^k(a + b) \leq \sigma^k(a) + b$$

holds for any $a, b \in \mathbb{R}_{\geq 0}$ and all $k \in \mathbb{N}_{\geq 0}$.

Proof. We proceed by induction. The case for $k = 0$ is trivial. We thus assume that the proposition holds for $\sigma^{k-1}(\cdot)$. By definition of $\alpha(\cdot)$, we have that

$$\sigma^k(a + b) = \sigma^{k-1}(a + b) - \alpha \circ \sigma^{k-1}(a + b).$$

By both applying monotonicity of $\alpha \circ \sigma^{k-1}$ and using the proposition for $\sigma^{k-1}(\cdot)$, we have that

$$\begin{aligned} \sigma^k(a + b) &\leq \sigma^{k-1}(a) + b - \alpha \circ \sigma^{k-1}(a) \\ &= \sigma^k(a) + b \end{aligned}$$

which completes the proof. ■

Proposition 5.17. The function

$$\Phi(k, \mathbf{w}, \mathbf{v}) := \max_{j \in \mathbb{N}_{0:k-1}} \sigma^{k-j-1}(4E_{j:k-1}),$$

in which $\sigma \in \mathcal{K}_\infty$ and $s - \sigma(s) := \alpha(s) \in \mathcal{K}_\infty$, satisfies the bound

$$\Phi(k, \mathbf{w}, \mathbf{v}) \leq \max_{j \in \mathbb{N}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \tilde{\gamma}(|(w(j), v(j))|),$$

in which $\tilde{\sigma}(s) := s - (1/2)\alpha(s) \in \mathcal{K}_\infty$ and $\tilde{\gamma} \in \mathcal{K}_\infty$, for all $k \in \mathbb{N}_{\geq 1}$.

Proof. We can show that the dissipation inequality

$$\Phi(k + 1, \mathbf{w}, \mathbf{v}) \leq \sigma(\Phi(k, \mathbf{w}, \mathbf{v})) + 4\bar{\gamma}_\ell(|(w(k), v(k))|)$$

holds for all $k \in \mathbb{N}_{\geq 1}$. By definition of $\Phi(\cdot)$ and $E_{j:k-1}$, we have that

$$\begin{aligned}\Phi(k+1, \mathbf{w}, \mathbf{v}) &:= \max_{j \in \mathbb{I}_{0:k}} \sigma^{k-j}(4E_{j:k}) \\ &= \max_{j \in \mathbb{I}_{0:k}} \sigma^{k-j}(4E_{j:k-1} + 4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|)).\end{aligned}$$

Then, we can use Proposition 5.16 and the monotonicity of $\sigma(\cdot)$ to obtain

$$\begin{aligned}\Phi(k+1, \mathbf{w}, \mathbf{v}) &\leq \max_{j \in \mathbb{I}_{0:k}} \sigma^{k-j}(4E_{j:k-1} + 4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|)) \\ &= \sigma\left(\max_{j \in \mathbb{I}_{0:k}} \sigma^{k-j-1}(4E_{j:k-1})\right) + 4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|).\end{aligned}\quad (5.29)$$

By again applying the definition of $E_{j:k-1}$, we note that

$$E_{k:k-1} := \sum_{i=k}^{k-1} \bar{\gamma}_\ell(|(\mathbf{w}(i), \mathbf{v}(i))|) = 0.$$

Furthermore, we have that $E_{j:k-1} \geq 0$ for all $j \in \mathbb{I}_{0:k-1}$. Because $k \geq 1$, the set $\mathbb{I}_{0:k-1}$ is not empty. Thus we have that

$$\max_{j \in \mathbb{I}_{0:k}} \sigma^{k-j-1}(4E_{j:k-1}) = \max_{j \in \mathbb{I}_{0:k-1}} \sigma^{k-j-1}(4E_{j:k-1}) = \Phi(k, \mathbf{w}, \mathbf{v}).$$

We can substitute this expression into (5.29) to obtain

$$\Phi(k+1, \mathbf{w}, \mathbf{v}) \leq \sigma(\Phi(k, \mathbf{w}, \mathbf{v})) + 4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|)$$

for all $k \in \mathbb{I}_{\geq 1}$, which is the sought-after inequality.

With this inequality, a typical ISS-Lyapunov argument produces a convolution maximization upper bound. Recall that $s - \sigma(s) := \alpha(s) \in \mathcal{K}_\infty$. We thus have that

$$\Phi(k+1, \mathbf{w}, \mathbf{v}) \leq \Phi(k, \mathbf{w}, \mathbf{v}) - \alpha(\Phi(k, \mathbf{w}, \mathbf{v})) + 4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|)$$

Suppose that $4\bar{\gamma}_\ell(|(\mathbf{w}(k), \mathbf{v}(k))|) \leq (1/2)\alpha(\Phi(k, \mathbf{w}, \mathbf{v}))$. Then we have that

$$\Phi(k+1, \mathbf{w}, \mathbf{v}) \leq \Phi(k, \mathbf{w}, \mathbf{v}) - (1/2)\alpha(\Phi(k, \mathbf{w}, \mathbf{v})) = \tilde{\sigma}(\Phi(k, \mathbf{w}, \mathbf{v}))$$

in which $\tilde{\sigma}(s) := s - (1/2)\alpha(s)$. Note that $\tilde{\sigma}(s) = 1/2(\sigma(s) + s)$, and, as a result, $\tilde{\sigma} \in \mathcal{K}_\infty$. Suppose then that $4\bar{\gamma}_\ell(|(w(k), v(k))|) > (1/2)\alpha(\Phi(k, \mathbf{w}, \mathbf{v}))$. Then we have that

$$\begin{aligned}\Phi(k+1, \mathbf{w}, \mathbf{v}) &\leq \alpha^{-1} \circ 8\bar{\gamma}_\ell(|(w(k), v(k))|) + 4\bar{\gamma}_\ell(|(w(k), v(k))|) \\ &= \tilde{\gamma}(|(w(k), v(k))|)\end{aligned}$$

in which $\tilde{\gamma}(s) := \alpha^{-1} \circ 8\bar{\gamma}_\ell(s) + 4\bar{\gamma}_\ell(s)$ is a \mathcal{K} function. No matter which bound holds, we have that

$$\Phi(k+1, \mathbf{w}, \mathbf{v}) \leq \tilde{\sigma}(\Phi(k, \mathbf{w}, \mathbf{v})) \oplus \tilde{\gamma}(|(w(k), v(k))|),$$

and this expression can be applied recursively to produce

$$\Phi(k, \mathbf{w}, \mathbf{v}) \leq \max_{j \in \mathbb{1}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \tilde{\gamma}(|(w(j), v(j))|).$$

The initial condition $\Phi(1, \mathbf{w}, \mathbf{v}) = 4\bar{\gamma}_\ell(|(w(0), v(0))|)$ is suppressed in this expression because $\tilde{\gamma}(s) \geq 4\bar{\gamma}_\ell(s)$ for all $s \in \mathbb{R}_{\geq 0}$. ■

5.3.2 Main result

Theorem 5.18. *If Assumptions 5.2, 5.3, 5.6, 5.13 and 5.14 are satisfied then full information estimation is RGAS.*

Proof. The proof strategy is to deal with cases based on how large the estimate error is. If the error is “large” compared to the size of the disturbances at time j , we are guaranteed some cost decrease in $Q(j+1|k)$. If, on the other hand, the error is “small”, we can use the upper bound to guarantee that $Q(j+1|k)$ remains “small”. Consider two measures of how “small” $e(j|k)$ is, given by

$$\alpha_3(|e(j|k)|) \leq 2\underline{\alpha}_\ell(|(w(j), v(j))|) \quad (5.30)$$

$$\check{\alpha}_2(|e(j|k)|) \leq E_{j:k-1} . \quad (5.31)$$

Suppose at least one of these equations holds for some $j \in \mathbb{I}_{0:k-1}$. Let $j^* \in \mathbb{I}_{0:k-1}$ be the greatest index for which one of these inequalities holds. For $j \in \mathbb{I}_{j^*+1:k-1}$, we have that

$$E_{j:k-1} < \check{\alpha}_2(|e(j|k)|)$$

and, as a result, have that

$$\begin{aligned} Q(j|k) &\leq \check{\alpha}_2(|e(j|k)|) + E_{j:k-1} \\ Q(j|k) &\leq 2\check{\alpha}_2(|e(j|k)|) . \end{aligned}$$

Furthermore, we have that

$$\underline{\alpha}_\ell(|(w(j), v(j))|) < (1/2)\alpha_3(|e(j|k)|)$$

and thus

$$\begin{aligned} Q(j+1|k) &\leq Q(j|k) - \alpha_3(|e(j|k)|) + \underline{\alpha}_\ell(|(w(j), v(j))|) \\ &\leq Q(j|k) - (1/2)\alpha_3(|e(j|k)|) \\ &\leq Q(j|k) - (1/2)\alpha_3 \circ (1/2)\check{\alpha}_2^{-1}(Q(j|k)) \\ &:= \tilde{\sigma}(Q(j|k)) . \end{aligned}$$

Note that $\tilde{\sigma}(s) < s$ for all $s \in \mathbb{R}_{\geq 0}$. By Proposition 5.15, we can find $\sigma \in \mathcal{K}$ such that $\tilde{\sigma}(s) \leq \sigma(s) < s$ and $s - \sigma(s) \in \mathcal{K}_\infty$. As a result, we have that

$$Q(j+1|k) \leq \sigma(Q(j|k))$$

for all $j \in \mathbb{I}_{j^*+1:k-1}$. By recursively applying this equation, we obtain

$$Q(k|k) \leq \sigma^{k-j^*-1}(Q(j^* + 1|k)) \quad (5.32)$$

Suppose that only (5.30) holds at j^* . Then we have that

$$|e(j^*|k)| \leq \alpha_3^{-1} \circ 2\underline{\alpha}_\ell(|(w(j^*), v(j^*))|) .$$

We can combine this with (5.27) to obtain

$$\begin{aligned} Q(j^* + 1|k) &\leq Q(j^*|k) - \alpha_3(|e(j^*|k)|) + \alpha_\ell(|(w(j^*), v(j^*))|) \\ &\leq \alpha_2(|e(j^*|k)|) + \alpha_\ell(|(w(j^*), v(j^*))|) \\ &\leq \alpha_2 \circ \alpha_3^{-1} \circ 2\underline{\alpha}_\ell(|(w(j^*), v(j^*))|) + \alpha_\ell(|(w(j^*), v(j^*))|) \\ &:= \omega_1(|(w(j^*), v(j^*))|) . \end{aligned}$$

Now suppose that (5.31) holds, irrespective of whether (5.30) holds. Then we have that

$$Q(j^*|k) \leq 2E_{j^*:k}$$

and thus

$$\begin{aligned} Q(j^* + 1|k) &\leq Q(j^*|k) - \alpha_3(|e(j^*|k)|) + \alpha_\ell(|(w(j^*), v(j^*))|) \\ &\leq 2E_{j^*:k} + \alpha_\ell(|(w(j^*), v(j^*))|) . \end{aligned}$$

Regardless of which inequalities hold at j^* , we have that

$$\begin{aligned} Q(j^* + 1|k) &\leq \omega_1(|(w(j^*), v(j^*))|) \oplus \left(2E_{j^*:k} + \alpha_\ell(|(w(j^*), v(j^*))|) \right) \\ &\leq \omega_1(|(w(j^*), v(j^*))|) \oplus 4E_{j^*:k} \oplus 2\alpha_\ell(|(w(j^*), v(j^*))|) \\ &= \omega_2(|(w(j^*), v(j^*))|) \oplus 4E_{j^*:k} , \end{aligned}$$

in which $\omega_2(s) := \omega_1(s) \oplus 2\alpha_\ell(s)$. We can then apply (5.32) to obtain

$$\begin{aligned} Q(k|k) &\leq \sigma^{k-j^*-1} \left(\omega_2(|(w(j^*), v(j^*))|) \oplus 4E_{j^*:k} \right) \\ &= \sigma^{k-j^*-1} \circ \omega_2(|(w(j^*), v(j^*))|) \oplus \sigma^{k-j^*-1}(4E_{j^*:k}) \end{aligned}$$

for all $k \in \mathbb{N}_{\geq 0}$ such that j^* is well-defined. Because j^* is a function of k , it is not obvious how this upper bound changes as a function of k . We have that both

$$\begin{aligned} \sigma^{k-j^*-1} \circ \omega_2(|(w(j^*), v(j^*)))|) &\leq \max_{j \in \mathbb{N}_{0:k-1}} \sigma^{k-j-1} \circ \omega_2(|(w(j), v(j))|) \quad \text{and} \\ \sigma^{k-j^*-1}(4E_{j^*:k}) &\leq \max_{j \in \mathbb{N}_{0:k-1}} \sigma^{k-j-1}(4E_{j:k-1}) := \Phi(k, \mathbf{w}, \mathbf{v}), \end{aligned} \quad (5.33)$$

because, by definition $j^* \in \mathbb{N}_{0:k-1}$, and thus $j = j^*$ is feasible in this maximization.

Recall that $\sigma(s) < s$ for all $s > 0$; in fact $\alpha(s) := s - \sigma(s)$ is a \mathcal{K}_∞ function. As a result, $\sigma^k(s)$ is a \mathcal{KL} function, and the first bound is a convolution-like expression in which disturbances in the distant past are discounted more than those in the recent past. The second bound is more difficult to interpret. In the maximization, there is a trade-off between the number of terms in the sum $E_{j:k-1}$ and the number of times $\sigma(\cdot)$ is iterated: the more terms the summation has, the more time the summation is discounted by $\sigma(\cdot)$. By Proposition 5.17, we have that

$$\Phi(k, \mathbf{w}, \mathbf{v}) \leq \max_{j \in \mathbb{N}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \tilde{\gamma}(|(w(j), v(j))|),$$

in which $\tilde{\sigma}, \tilde{\gamma} \in \mathcal{K}_\infty$ and $s - \tilde{\sigma}(s) = (1/2)\alpha(s) \in \mathcal{K}_\infty$. We thus have that

$$Q(k|k) \leq \max_{j \in \mathbb{N}_{0:k-1}} \sigma^{k-j-1} \circ \omega_3(|(w(j), v(j))|) \oplus \max_{j \in \mathbb{N}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \tilde{\gamma}(|(w(j), v(j))|).$$

Let $\omega(s) := \omega_2(s) \oplus \tilde{\gamma}(s)$. Because $\tilde{\sigma}(s) \geq \sigma(s)$ for all $s \in \mathbb{R}_{\geq 0}$, we have that

$$Q(k|k) \leq \max_{j \in \mathbb{N}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \omega(|(w(j), v(j))|)$$

for all $k \in \mathbb{N}_{\geq 0}$ such that j^* is defined.

Now suppose that neither (5.30) or (5.31) holds for all $j \in \mathbb{N}_{0:k-1}$. Then we have that

$$Q(k|k) \leq \sigma^k(Q(0|k)).$$

We can combine this equation with (5.19) to obtain

$$\begin{aligned} Q(k|k) &\leq \sigma^k \left(\alpha_0(|\bar{x} - x(0)|) + 2E_{0:k-1} \right) \\ &\leq \sigma^k \circ 2\alpha_0(|\bar{x} - x(0)|) \oplus \sigma^k(4E_{0:k-1}) \\ &\leq \sigma^k \circ 2\alpha_0(|\bar{x} - x(0)|) \oplus \Phi(k, \mathbf{w}, \mathbf{v}). \end{aligned}$$

By the same argument made in the case for which j^* was defined, we have that

$$\Phi(k, \mathbf{w}, \mathbf{v}) \leq \max_{j \in \mathbb{I}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \omega(|(w(j), v(j))|).$$

Thus, whether or not $j^*(k)$ is defined, and irrespective of its value, we have that

$$Q(k|k) \leq \sigma^k \circ 2\alpha_0(|\bar{x} - x(0)|) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \tilde{\sigma}^{k-j-1} \circ \omega(|(w(j), v(j))|).$$

Finally, we have that $\alpha_1(|e(k|k)|) \leq Q(k|k)$, and thus

$$|e(k|k)| \leq \alpha_1^{-1} \circ \sigma^k \circ 2\alpha_0(|\bar{x} - x(0)|) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \alpha_1^{-1} \circ \tilde{\sigma}^{k-j-1} \circ \omega(|(w(j), v(j))|).$$

Let $\beta_x(s, k) := \alpha_1^{-1} \circ \sigma^k \circ 2\alpha_0(s)$ and $\beta_d(s, k) := \alpha_1^{-1} \circ \tilde{\sigma}^k \circ \omega(s)$, noting that $\beta_x, \beta_d \in \mathcal{KL}$. We thus have that

$$|e(k|k)| \leq \beta_x(|\bar{x} - x(0)|, k) \oplus \max_{j \in \mathbb{I}_{0:k-1}} \beta_d(|(w(j), v(j))|, k - j - 1)$$

and thus FIE is RGAS. ■

Corollary 5.19 ((Allan and Rawlings, 2021, Prop. 3.11)). *If the sequence $((w(k), v(k)))$ converges to zero, then $e(k|k)$ converges to zero.*

5.3.3 Discussion

With this result in hand, a more detailed discussion of both its assumptions and its form is warranted. Assumptions 5.2 and 5.3 are uncontroversial, being present in

every result on optimization-based state estimation. The assumption that the system is i-IOSS is also standard in the optimization-based state estimation literature. Assuming the existence of an i-IOSS Lyapunov function as in Assumption 5.6, however, is new to this literature. The i-IOSS converse theorem provided in (Allan et al., 2021, Theorem 3.2) shows that every i-IOSS system admits an i-IOSS Lyapunov function, and while the assumption is new, it is not any more restrictive than previous assumptions in the literature. Furthermore, as is discussed in Chapter 2, (Allan et al., 2021, Prop. 2.5) shows that *any* system that admits a robustly stable estimator is i-IOSS, so no weaker notion of detectability is admissible.⁵ It is fair to argue that i-IOSS is an abstract property that is not particularly easy to check, but the best way to deal with that is to develop necessary conditions to check whether a system fails to be i-IOSS or sufficient conditions that can guarantee i-IOSS in, for example, a region.

Assumptions 5.13 and 5.14, on the other hand, are both new to the literature and rather abstract. We examine Assumption 5.13 first. Note that if $\lambda(\cdot)$ is an *i-IOSS* Lyapunov function, then, for any $c > 0$, $c\lambda(\cdot)$ is also an *i-IOSS* Lyapunov function. Assumption 5.13 can thus be reduced to a comparison of the limits of those \mathcal{K} functions as $s \rightarrow 0$ and $s \rightarrow \infty$, i.e., whether

$$\lim_{s \rightarrow 0} \frac{\alpha_x(s/2)}{\alpha_2(s)} > 0 \quad \text{and} \quad \lim_{s \rightarrow \infty} \frac{\alpha_x(s/2)}{\alpha_2(s)} > 0$$

⁵With one minor caveat. As is discussed in Remark 5.1, in the standard formulation of FIE and MHE the process disturbance is assumed to be zero-mode. i-IOSS is necessary for a robustly stable estimator such that, when given a nonzero input forecast \mathbf{u}_f , the estimation error is proportional to the forecasting error $\|\mathbf{u} - \mathbf{u}_f\| = \|\mathbf{w}\|$. Therefore, a slightly weaker property than i-IOSS that gives a special role to $\mathbf{u} = \mathbf{0}$ might produce Theorem 5.18, but this property is more complex conceptually than i-IOSS and not particularly interesting so it is not discussed further.

and

$$\lim_{s \rightarrow 0} \frac{\underline{\alpha}_\ell(s/2)}{\sigma_u(s) + \sigma_y(s)} > 0 \quad \text{and} \quad \lim_{s \rightarrow \infty} \frac{\underline{\alpha}_\ell(s/2)}{\sigma_u(s) + \sigma_y(s)} > 0.$$

This sort of requirement is reminiscent of Proposition 3.20. However, with extra effort, this result can be improved. In (Allan et al., 2021, Corollary 5.4), it is shown that for any convex $\rho \in \mathcal{K}_\infty$, we have that $\rho \circ \lambda(\cdot)$ is an i-IOSS Lyapunov function. As a result, the limits as $s \rightarrow 0$ can be smoothed out, leaving the requirements that

$$\lim_{s \rightarrow \infty} \frac{\alpha_x(s/2)}{\alpha_2(s)} > 0 \quad \text{and} \quad \lim_{s \rightarrow \infty} \frac{\underline{\alpha}_\ell(s/2)}{\sigma_u(s) + \sigma_y(s)} > 0.$$

We note, however, that an assumption like this one is used to show global stability in Chapter 4. The question of why or whether an assumption like this is necessary, then, is also unresolved in the case of tracking or even regulation. In (Grimm et al., 2005), which studies the case of regulation with a detectable stage cost, a compatibility assumption is necessary to establish asymptotic stability, and if such an assumption is absent, the most that existing theory can guarantee is practical stability. Absent a counterexample, however, it is not clear whether the assumption is necessary or whether existing theory is limited.

Assumption 5.14 is a perturbed stabilizability assumption. The two questions that need to be addressed are: why should we need a stabilizability assumption at all, and, if we do need one, why it should have this form? In Kalman filtering, stabilizability is a somewhat pervasive assumption. However, it is known that stabilizability is not a necessary condition for the Kalman filter to be stable. So long as there are no uncontrollable modes with eigenvalue on the unit circle,⁶ then the

⁶This statement is distinct from saying there are no uncontrollable marginally stable modes in

steady-state Kalman gain produces a stable observer, and the time-varying Kalman filter converges exponentially to it (Chan et al., 1984, Theorem 4.2). The presence of uncontrollable modes on the unit circle complicate matters, and while the time-varying Kalman filter may still be stable, the rate of convergence may be harmonic (i.e., like $1/k$), not exponential, and the steady-state Kalman gain produces only a marginally stable observer (Chan et al., 1984, Theorem 4.3). Furthermore, while the time-varying filter is mathematically stable, it may not be numerically stable (Bitmead et al., 1985).

Most results in nonlinear optimization-based state estimation, with the notable exception of my own work (Allan and Rawlings, 2019b), have not used an explicit stabilizability assumption. However, the frequent use of process models with additive process noise, i.e., $f(x, w) = f(x) + w$, clouds the picture to a considerable extent. For a short period following the publication of (Müller, 2017), analysis of MHE leapfrogged that of FIE by the ability to prove robust stability directly without a max-term. One of the reasons for that state of affairs was that MHE has an additional degree of freedom in the ability to move the initial state at the beginning of the moving horizon, while FIE can move the initial state only at the time $k = 0$. However, these analyses of MHE required a careful choice of the prior weighting. It had to be sufficiently small to make moving the smoothed estimate to the real state attractive in terms of optimality. However, it had to be sufficiently large in order to penalize increasing the magnitude of smoothing error in the

the system. If the single-step state transition matrix A is defective, then such a mode can still be unstable.

“unobservable subspace”⁷ At first glance, that may seem like a natural requirement. However, the choice of the prior weighting cannot be made independently of the horizon length N . First the horizon length must be chosen, and only then can the prior weighting be chosen (see, for example, the careful quantifier order in (Allan and Rawlings, 2019a, Theorem 1)). The end result is that, given a functional implementation of MHE, one could not extend the horizon length and guarantee that the implementation would remain functional.

What stabilizability guarantees is that, modulo disturbances, if the estimate error $e(j|k)$ ever becomes small, then it remains small at future times. The only way to prove something about optimization-based state estimation is by either proving that bad behavior is suboptimal, or, the contrapositive, all sufficiently-close-to-optimal solutions are well-behaved. Absent a stabilizability assumption, there are few solutions with which to compare the optimal solution at a time k . There are the sequences that actually underlie the output sequence \mathbf{y} , \mathbf{x} , \mathbf{w} , and \mathbf{v} , and there are the optimal solutions from $\mathbb{P}_j(\cdot)$ for $j \in \llbracket 0:k-1$. The former is always a useful comparison, but the latter do not take into account all the data that is taken into account at time k , and the most that can be guaranteed is a cost that exponentially diverges with $k - j$, except in the trivial case in which $f(\cdot)$ is a contraction map. Stabilizability permits us to stitch solutions together—to follow one solution for a while, then transition to another solution with a bounded cost. The benefits of such an assumption from the analysis it enables is obvious. The bottom line, however, is that stabilizability still seems like an unnecessary assumption, and there should be

⁷Although nonlinear analogs to observable and unobservable subspaces have been proposed, I know of no general way to put a general i-IOSS system into such a form

a method to avoid using it. However, it is not a particularly strong assumption, and it is likely that implementations of FIE or MHE that are not stabilizable may have numerical issues due to round-off errors accumulating in unstable “modes” of the nonlinear system.

Next, there is the question, if we have resigned ourselves to making a stabilizability assumption, why use one like Assumption 5.14 and not Assumption 5.8? The straightforward answer is not intellectually satisfying: it allows us to prove the result we want. The derivation of (5.26) shows why the “normal” incremental stabilizability assumption Assumption 5.8 is not suitable for the task at hand. As is shown in the next section, for the special case of *exponential* stabilizability, Assumption 5.14 can be derived from Assumption 5.8 in a straightforward fashion. In the asymptotic case, however, the most that I have been able to accomplish is the production of a semiglobal version of Assumption 5.14 from Assumption 5.8. Because we have no bound on the size of the smoothed estimates $e(j|k)$, however, this semiglobal result is not particularly useful. It may be that this situation is another case of several properties that happen to coincide in the case of linear time-invariant systems that are for nonlinear systems.

Finally, we revisit the example from Section 5.2.2. It satisfied the assumptions necessary for nominal stability, but does it satisfy those necessary for robust stability? As it turns out, it does not. Assumption 5.13 requires that there exist an i-IOSS Lyapunov function $\lambda(\cdot)$ such that

$$\underline{\alpha}_x(s/2) \geq \alpha_2(s)$$

for all $s \in \mathbb{R}_{\geq 0}$. However, it can be shown that there exists no such $\lambda(\cdot)$ with an

upper bound $\alpha_2(s)$ that has an asymptotic growth rate less than or equal to $\ln(s^2 + 1)$. Then we might be able to hope that the smoothed estimates $e(j|k)$ behave better for a system satisfying Assumption 5.13 than they do in that example. However, I know of no argument that guarantees good behavior.

5.4 Robust global exponential stability of FIE

Many of the above results can be tightened and streamlined in the special case of exponential stability and detectability. In particular, Assumption 5.14 can be derived from Assumption 5.8 in a straightforward and satisfactory manner, Assumption 5.13 results only in the requirement that $\ell(\cdot)$ admit power-law lower and upper bounds, and Propositions 5.15 to 5.17 are reduced to a simple fact about exponential functions. The most important reason to consider the exponential case, however, is that we can show that MHE is a stable estimator with a sufficiently long horizon. These results are presented in detail in (Allan and Rawlings, 2021), so here only the major highlights are presented.

5.4.1 Assumptions

We do not require a continuity assumption stronger than Assumption 5.2. We do require special forms for the upper and lower bounds of $\ell_x(\cdot)$ and $\ell(\cdot)$, as well as an exponential i-IOSS Lyapunov function and an exponential stabilizability assumption.

Assumption 5.20 (Power-law cost function bounds). There exist $\underline{c}_x, \bar{c}_x, \underline{c}_\ell, \bar{c}_\ell > 0$ and $\sigma \geq 1$ such that

$$\begin{aligned} \underline{c}_x |\chi(0) - \bar{x}|^\sigma &\leq \ell_x(\chi(0), \bar{x}) \leq \bar{c}_x |\chi(0) - \bar{x}|^\sigma \\ \underline{c}_\ell |(\omega, \nu)|^\sigma &\leq \ell(\omega, \nu) \leq \bar{c}_\ell |(\omega, \nu)|^\sigma \end{aligned}$$

for all $\chi(0) \in \mathbb{R}^n$ and all $\omega \in \mathbb{R}^g$ and $\nu \in \mathbb{R}^p$.

Remark 5.21. This assumption permits the ever-popular least squares stage costs, but *does not* permit mixed power stage costs such as

$$\ell(\omega, \nu) = |(\omega, \nu)|^2 + |(\omega, \nu)|_1$$

that might occur if one took a Lasso approach to MHE.

Assumption 5.22 (i-IOSS Lyapunov function). There exist $\Lambda : \mathbb{X}^n \times \mathbb{X}^n \rightarrow \mathbb{R}_{\geq 0}$ and $c_i > 0$ such that

$$\begin{aligned} c_1 |x_1 - x_2|^\sigma &\leq \Lambda(x_1, x_2) \leq c_2 |x_1 - x_2|^\sigma \\ \Lambda(f(x_1, w_1), f(x_2, w_2)) &\leq \Lambda(x_1, x_2) - c_3 |x_1 - x_2|^\sigma + c_d |(w_1 - w_2, h(x_1) - h(x_2))|^\sigma, \end{aligned}$$

in which $\sigma \geq 1$ comes from Assumption 5.20, for any $x_1, x_2 \in \mathbb{X}^n$ and $w_1, w_2 \in \mathbb{W}^g$.

Remark 5.23. We do not make any explicit assumption like Assumption 5.13 because the stage cost has a lower bound proportional to s^σ and the dissipation rate is proportional to s^σ . Furthermore, we have that $(s/2)^\sigma = (1/2)^\sigma s^\sigma$, so the factors of two inside the \mathcal{K} function composition do not pose any problem for power costs. As a result, we can always rescale $\Lambda(\cdot)$ such that the dissipation inequality (5.22) is satisfied.

Assumption 5.24 (Incremental exponential stabilizability). There exists $\bar{c} > 0$ such that, for every $\check{x}, x_r \in \mathbb{X}$ and $\mathbf{w}_r \in \mathbb{W}^\infty$, there exists some $\check{\mathbf{w}} \in \mathbb{W}^\infty$ such that

$$\sum_{k=0}^{\infty} \ell(\check{w}(k) - w(k), y(k) - \check{y}(k)) \leq \bar{c} |\check{x} - x_r|^\sigma,$$

in which

$$\begin{aligned} x^+ &= f(x, w) & y &= h(x) \\ \check{x}^+ &= f(\check{x}, \check{w}) & \check{y} &= h(\check{x}) \end{aligned}$$

and $\sigma \geq 1$ comes from Assumption 5.20.

A perturbed stabilizability assumption like Assumption 5.14 can be derived from Assumption 5.24 in a straightforward manner. We first require a proposition about distributing a power over summations.

Proposition 5.25 ((Allan and Rawlings, 2021, Prop. 3.8)). *Let $x_i \in \mathbb{R}_{\geq 0}$. Then we have that*

$$\left(\sum_{i=1}^n x_i \right)^\sigma \leq n^{\sigma-1} \sum_{i=1}^n x_i^\sigma$$

for any $\sigma \geq 1$.

Proposition 5.26 ((Allan and Rawlings, 2021, Prop. 3.9)). *If Assumption 5.24 holds, then there exist $\check{c}_x, \check{c}_d > 0$ such that, for every finite sequences $\mathbf{w}_r(k) \in \mathbb{W}^k$ and $\mathbf{v}_r(k) \in (\mathbb{R}^p)^k$ and any $\check{x}, x_r \in \mathbb{X}$, there exists $\check{\mathbf{w}} \in \mathbb{W}^\infty$ such that*

$$\sum_{i=0}^{\infty} \ell(\check{w}(i), \check{v}(i)) \leq \check{c}_x |\check{x} - x_r|^\sigma + \check{c}_d \sum_{i=0}^{k-1} |(w_r(k), v_r(k))|^\sigma$$

in which

$$\begin{aligned}
& \check{x}^+ = f(\check{x}, \check{w}) & y_r &= h(\check{x}) + \check{v} \\
x_r^+ &= \begin{cases} f(x_r, w_r) & \text{for } i \in \mathbb{I}_{0:k-1} \\ f(x_r, 0) & \text{for } i \in \mathbb{I}_{k:\infty} \end{cases} & y_r &= \begin{cases} h(x_r) + v_r & \text{for } i \in \mathbb{I}_{0:k-1} \\ h(x_r) & \text{for } i \in \mathbb{I}_{k:\infty} \end{cases}
\end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$.

Proof sketch. There are two key steps in this proof. The first is to note that, because $\mathbf{w}_r(k)$ and $\mathbf{v}_r(k)$ are finite, we can apply Assumption 5.24 to obtain disturbance sequences $\check{\mathbf{w}}$ and $\check{\mathbf{v}}$ such that

$$\sum_{j=0}^{k-1} \ell(\check{w}(j) - w(j), \check{v}(j) - v(j)) + \sum_{j=k}^{\infty} \ell(\check{w}(j), \check{v}(j)) \leq \bar{c} |\check{x} - x_r|^\sigma,$$

which leaves us with only a finite number of terms that are in an unsuitable form.

For an individual term, we can apply Assumption 5.20 to obtain

$$\ell(\check{w}(i) - w(i), \check{v}(i) - v(i)) \geq \underline{c}_\ell |(\check{w}(i) - w(i), \check{v}(i) - v(i))|^\sigma$$

to bring the stage cost into a form where we can apply Proposition 5.25 to obtain

$$\begin{aligned}
\underline{c}_\ell |(\check{w}(i), \check{v}(i))|^\sigma &= \bar{c}_\ell |(\check{w}(i) - w(i), \check{v}(i) - v(i)) + (w(i), v(i))|^\sigma \\
&\leq 2^{\sigma-1} |(\check{w}(i) - w(i), \check{v}(i) - v(i))|^\sigma + 2^{\sigma-1} |(w(i), v(i))|^\sigma.
\end{aligned}$$

Further algebraic manipulation and application of Assumption 5.20 produces a bound of the form sought. ■

Remark 5.27. The absence of a relaxed triangle inequality like Proposition 5.25 is what causes difficulties in attempting to deduce Assumption 5.14 from Assumption 5.8.

5.4.2 Robust exponential stability

Using these assumptions, we can deduce equations analogous to (5.25)–(5.27). First, we define the amount of disturbance “energy” injected into the system between times j and k , $E_{j:k-1}$, in a slightly different fashion:

$$E_{j:k-1} := \sum_{i=j}^{k-1} |(\omega(i), \nu(i))|^\sigma,$$

which differs from (5.24) in that there is no need to consolidate the constants \check{c}_d and \bar{c}_ℓ into a single constant because they can be pulled outside of the summation. Then we can obtain a Q function with bounds of the form

$$\begin{aligned} Q(0|k) &\leq K_{0,x} |\bar{e}|^\sigma + K_{0,E} E_{0:k-1} \\ K_1 |e(j|k)|^\sigma &\leq Q(j|k) \leq K_{2,x} |e(j|k)|^\sigma + K_{2,E} E_{j:k-1} \\ Q(j+1|k) &\leq Q(j|k) - K_{3,x} |e(j|k)|^\sigma + K_{3,d} |(\omega(j), \nu(j))|^\sigma \end{aligned}$$

in which $K_i > 0$ and $\sigma \geq 1$ is from Assumption 5.20. We can now prove that, under these assumptions, FIE is robustly globally exponentially stable (RGES).

Theorem 5.28 ((Allan and Rawlings, 2021, Thm 3.15)). *If Assumptions 5.2, 5.20, 5.22 and 5.24 hold, then there exists $\lambda \in (0, 1)$ and constants $\gamma_x, \gamma_d > 0$ such that*

$$|e(k|k)| \leq \gamma_x \lambda^k |\bar{e}| + \sum_{j=0}^{k-1} \gamma_d |(\omega(j), \nu(j))| \lambda^{k-j-1}$$

for all $x, \bar{x} \in \mathbb{X}$, $\mathbf{w} \in \mathbb{W}^\infty$, $\mathbf{v} \in (\mathbb{R}^p)^\infty$, and $k \in \mathbb{I}_{\geq 0}$.

Proof sketch. The major difference in the exponential case is that, instead of using Proposition 5.17 in order to transform (5.33) from a bound based on $E_{j:k-1}$ to one based on $|(\omega(j), \nu(j))|$, we can use basic facts about exponential functions. In a manner similar to the proof of Theorem 5.18, we can show

$$Q(k|k) \leq K_{0,x} \eta^k |\bar{e}|^\sigma + \bar{K} \max_{j \in \mathbb{I}_{0:k-1}} \eta^{k-j-1} E_{j:k-1},$$

in which $\bar{K} > 0$ and $\eta \in (0, 1)$. We have that

$$E_{j:k-1} = \sum_{j=0}^{k-1} |(w(j), v(j))|^\sigma \leq (k-j) \|(\mathbf{w}, \mathbf{v})\|_{j:k-1}^\sigma$$

from the definition of $E_{j:k-1}$. If we let $\nu \in (\eta, 1)$, there exists some $C_\nu > 0$ such that

$$t\eta^t \leq C_\nu \nu^t$$

for all $t \in \mathbb{R}_{\geq 0}$. We can then choose some such ν , e.g., $\eta^{1/2}$, to find an upper bound

$$\max_{j \in \mathbb{I}_{0:k-1}} \eta^{k-j-1} E_{j:k-1} \leq C_\nu \max_{j \in \mathbb{I}_{0:k-1}} \nu^{k-j-1} \|(\mathbf{w}, \mathbf{v})\|_{j:k-1}^\sigma.$$

Finally, we have that

$$\max_{j \in \mathbb{I}_{0:k-1}} \nu^{k-j-1} \|(\mathbf{w}, \mathbf{v})\|_{j:k-1}^\sigma = \max_{j \in \mathbb{I}_{0:k-1}} \nu^{k-j-1} |(w(j), v(j))|^\sigma,$$

because if $\|(\mathbf{w}, \mathbf{v})\|_{j:k-1} \neq |(w(j), v(j))|$, then the choice of j will not be optimal because there is another feasible time with the same disturbance size that is less discounted. ■

5.5 Moving horizon estimation

While FIE is a useful theoretical benchmark, the only practical case in which it can be implemented is in the case of the Kalman filter, where all past data can be summarized with the Kalman filter's state estimate and a quadratic arrival cost related to the covariance matrix associated to that estimate. Therefore, as in the case of MPC, we choose a fixed horizon length N and include only the N most

recent measurements in the model, possibly taking into account past data in an approximate way. Define the MHE problem $\hat{\mathbb{P}}_k(\bar{x}, \mathbf{y}(k))$:

$$\begin{aligned} \min_{\chi(k-N), \omega, \nu} \hat{V}_k(x, \omega, \nu, \bar{x}, \mathbf{y}) &:= \Gamma_{k-N}(\chi(k-N), \bar{x}_{k-N}) + \sum_{i=k-N}^{k-1} \ell(\omega(i), \nu(i)) \\ \text{subject to} \quad \chi^+ &= f(\chi, \omega) \\ y &= h(\chi) + \nu \\ \chi(k-N) &\in \mathbb{X}, \omega \in \mathbb{W}^N, \nu \in (\mathbb{R}^p)^N, \end{aligned}$$

in which $\Gamma_{k-N}(\cdot)$ is the prior weighting at time $k-N$ and \bar{x}_{k-N} is an approximate estimate of $x(k-N)$. For times $k \leq N$ before the horizon has been fully populated with measurements, the MHE problem $\hat{\mathbb{P}}_k(\cdot)$ is defined as the FIE problem $\mathbb{P}_k(\cdot)$.

5.5.1 Choice of prior weighting

The key to ensuring that MHE is robustly stable is a good choice of prior weighting. The simplest choice is to use a prior weight of zero. However, in order for this choice to work, the system must be *observable* (Rawlings et al., 2017, Sec. 4.3.1). In order to apply MHE to detectable systems, some sort of positive definite prior must be used. Another important choice of prior is the *arrival cost*

$$\begin{aligned} L_{k-N}(p) &:= \min_{\chi(0), \omega, \nu} V_{k-N}(x, \omega, \nu, \bar{x}, \mathbf{y}) := \ell_x(\chi(0), \bar{x}) + \sum_{i=0}^{k-N-1} \ell(\omega(i), \nu(i)) \\ \text{subject to} \quad \chi^+ &= f(\chi, \omega) \\ y &= h(\chi) + \nu \\ \chi(k-N) &= p \\ \chi(0) &\in \mathbb{X}, \omega \in \mathbb{W}^{k-N}, \nu \in (\mathbb{R}^p)^{k-N}. \end{aligned}$$

The arrival cost $L_{k-N}(\cdot)$ is a result of decomposing the FIE problem $\mathbb{P}_k(\cdot)$ using dynamic programming into a problem containing the last N measurements, which corresponds to the MHE problem, and the rest of the measurements taken before time $k - N$. In order to calculate the arrival cost, one solves an FIE problem with a terminal equality constraint. The benefit of using the arrival cost is that using it as a prior results in the MHE problem being equivalent to the FIE problem, and therefore inheriting the robust stability properties of FIE. However, except in the important case of the Kalman filter, solving MHE using the arrival cost is just as expensive as solving the FIE problem, and therefore there is no reduction in computational complexity by using MHE. As a result, the arrival cost, like FIE, is important as a benchmark in what we would want from a prior weighting $\Gamma_{k-N}(\cdot)$.

The first property we note about the arrival cost is that it attains its minimum at $\hat{x}(k - N|k - N)$. The choice of $\bar{x}_{k-N} = \hat{x}(k - N|k - N)$ is known as using a *filtering update*. One popular strategy for MHE is then to use a local approximation of $\Gamma_{k-N}(\cdot)$ in a neighborhood of $\hat{x}(k - N|k - N)$, whether by using the extended Kalman filter (Rao et al., 2003) or by using the KKT conditions of the optimization problem (Chu et al., 2012). While this process has good theoretical and statistical properties, in practice it can cause undesirable dynamics in the state estimates. To resolve this problem, the choice of using $\bar{x}_{k-N} = \hat{x}(k - N|k - 1)$ was introduced. This choice is called a *smoothing update* (Findeisen, 1997; Rao et al., 2001; Tenny, 2002). The problem with using the smoothing update is that $\hat{x}(k - N|k - 1)$ is based on the measurements $\mathbf{y}(k - N - 1 : k - 2)$, while the MHE problem $\hat{\mathbb{P}}_k(\cdot)$ is based on the measurements $\mathbf{y}(k - N : k - 1)$. This overlap in the data considered introduces

spurious correlations in the MHE problem. To counteract this problem in the case of the Kalman filter, the prior weighting can be updated to remove the correlations in the data in an exact fashion. For the nonlinear case, these Kalman filtering equations can be used for the system linearized around the optimal trajectory.

As the example in Section 5.2.2 illustrates, we cannot guarantee good behavior of smoothed estimates $\hat{x}(k-N|k-1)$. Therefore, we choose a filtering prior weighting.

Assumption 5.29. For all $k \in \mathbb{I}_{\geq 0}$, we have that $\Gamma_k : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is continuous, $\bar{x}_k = \hat{x}(k|k)$, and there exist $\underline{c}_\Gamma, \bar{c}_\Gamma > 0$ such that

$$\underline{c}_\Gamma |\chi - \hat{x}(k|k)|^\sigma \leq \Gamma_k(\chi, \hat{x}(k|k)) \leq \bar{c}_\Gamma |\chi - \hat{x}(k|k)|^\sigma$$

uniformly in k , in which $\sigma \geq 1$ comes from Assumption 5.20.

For the time-varying Kalman filter using the arrival cost as a prior weighting, this assumption is fulfilled if the system is uniformly detectable, which gives us the lower bound, and uniformly *controllable* (not stabilizable), which gives us the upper bound. For linear systems, the uncertainty in uncontrollable modes asymptotically goes to zero and the arrival cost asymptotically diverges in those modes. For nonlinear systems, I know of no clear separation between “controllable” and “uncontrollable” modes the way there is for time-invariant linear systems. As a result, we must treat the entire space \mathbb{X} as “controllable”, that is to say, a region where our certitude is bounded. Furthermore, because the arrival cost $L_k(\cdot)$ is given by a constrained nonlinear optimization problem, there is no guarantee that it is even continuous at $\hat{x}(k|k)$. Therefore, this assumption could exclude using the arrival cost as a prior weighting. Theoretically, this state of affairs is unsatisfactory, but since we are

almost always approximating the arrival cost, limiting the certitude in the prior estimate is an acknowledgment that the approximation is not perfect. Use of the EKF or optimality conditions to approximate the prior weighting may not guarantee these bounds, so a post-processing step to ensure that they are satisfied may be required.

5.5.2 Robust stability of MHE

Theorem 5.30 ((Allan and Rawlings, 2021, Thm. 4.2)). *If Assumptions 5.2, 5.20, 5.22, 5.24 and 5.29 hold, then there exists a horizon length N_0 such that for all horizons $N \geq N_0$ MHE is robustly globally exponentially stable.*

Proof sketch. The key insight to this proof (articulated first by Hu (2017)) is that the solution to $\hat{\mathbb{P}}_k(\cdot)$ with prior weighting $\Gamma_{k-N}(\cdot)$ and $\bar{x}_{k-N} = \hat{x}(k-N|k-N)$ is the same as the solution to $\mathbb{P}_N(\cdot)$ with $\ell_x(\cdot) = \Gamma_{k-N}(\cdot)$ and $\bar{x} = \hat{x}(k-N|k-N)$. With care, we can apply Theorem 5.28 in order to get bounds on $e(k|k)$ based on $e(k-N|k-N)$ and the disturbances \mathbf{w} and \mathbf{v} .

Because MHE and FIE are identical for $k \leq N$, we immediately have that

$$|e(k|k)| \leq \gamma_x \lambda^k |\bar{e}| + \sum_{j=0}^{k-1} \gamma_d |(\mathbf{w}(j), \mathbf{v}(j))| \lambda^{k-j-1} \quad (5.34)$$

for $k \in \mathbb{I}_{0:N}$. For $k > N$, we can use the relationship between the MHE problem and an FIE problem with a different starting time and prior estimate to derive the bound

$$|e(k|k)| \leq \tilde{\gamma}_x \lambda^N |e(k-N|k-N)| + \sum_{j=k-N}^{k-1} \tilde{\gamma}_d |(\mathbf{w}(j), \mathbf{v}(j))| \lambda^{k-j-1}, \quad (5.35)$$

in which $\tilde{\gamma}_x, \tilde{\gamma}_d > 0$ differ from γ_x, γ_d because the upper and lower bounds for $\Gamma_{k-N}(\cdot)$ from Assumption 5.29 are used rather than those for $\ell_x(\cdot)$ from Assumption 5.20. If N_0 is chosen large enough such that $\tilde{\gamma}_x \lambda^{N_0} < 1$, then the error in the prior shrinks over the horizon N modulo disturbances. Robust stability then follows through a tedious but straightforward argument by induction (pioneered by Müller (2017)), in which (5.35) is applied T times until $k - NT \leq N$. Then (5.34) can be used to bound the error in terms of \bar{e} . ■

5.5.3 Discussion

There are two major questions that should be addressed:

1. What about the case of asymptotic detectability or stabilizability?
2. Why are we not using a Q function to analyze MHE?

These issues are related. The fundamental problem is that we have no better way of relating $\hat{x}(0|k)$ with $x(0)$ than (5.17)

$$|\hat{x}(0|k) - x(0)| \leq |\hat{x}(0|k) - \bar{x}| + |\bar{x} - x(0)|.$$

For MHE with a filtering prior, this equation becomes

$$\begin{aligned} |\hat{x}(k - N|k) - x(k - N)| &\leq |\hat{x}(k - N|k) - \hat{x}(k - N|k - N)| \\ &\quad + |\hat{x}(k - N|k - N) - x(k - N)|. \end{aligned}$$

We require a relationship like this to relate $\lambda(\hat{x}(k - N|k), x(k - N))$ with $\lambda(\hat{x}(k - N|k - N), x(k - N))$. The former is needed to form a Q function for $\hat{\mathbb{P}}_k(\cdot)$, the

latter is given by a Q function for the MHE problem $\hat{\mathbb{P}}_{k-N}(\cdot)$. Using the triangle inequality, we end up with a relationship

$$\begin{aligned} \Lambda(\hat{x}(k-N|k), x(k-N)) &\leq \alpha_2(|\hat{x}(k-N|k) - \hat{x}(k-N|k-N)| \\ &\quad + |\hat{x}(k-N|k-N) - x(k-N)|) \\ &\leq \alpha_2(|\hat{x}(k-N|k) - \hat{x}(k-N|k-N)| \\ &\quad + \alpha_1^{-1} \circ \Lambda(\hat{x}(k-N|k-N), x(k-N))). \end{aligned}$$

If $\alpha_1(\cdot)$ and $\alpha_2(\cdot)$ are of the form $c_i |\cdot|^\sigma$, this sort of relationship does not pose a problem—we can apply Proposition 5.25 to split this expression apart. For general \mathcal{K}_∞ functions, we can proceed by using (5.15) to get the expression

$$\begin{aligned} \Lambda(\hat{x}(k-N|k), x(k-N)) &\leq \alpha_2(2|\hat{x}(k-N|k) - \hat{x}(k-N|k-N)|) \\ &\quad \oplus \alpha_2 \circ 2\alpha_1^{-1} \circ \Lambda(\hat{x}(k-N|k-N), x(k-N)) \end{aligned}$$

which is where I can see no path forward. Out of necessity $\alpha_2(s) \geq \alpha_1(s)$ and therefore $\alpha_2^{-1}(s) \leq \alpha_1^{-1}(s)$. No assumption like Assumption 5.13 can solve this problem. Every \mathcal{K}_∞ function $\alpha(\cdot)$ that satisfies the functional equation $\alpha(a \cdot b) = \alpha(a) \cdot \alpha(b)$ has the form s^σ for some $\sigma > 0$.⁸ Given this situation, there is little point in introducing a Q function for the MHE problem. Because of this forced step through $|\cdot|$, we require MHE to contract the estimate error $|e(k|k)|$ over the horizon length N .

⁸Above, we treated the case of $\sigma \geq 1$ to simplify the exposition. In general, the same result holds for $\sigma \in (0, 1)$, but the proofs require appealing to subadditivity rather than Proposition 5.25. Cost functions with bounds of powers less than one are necessarily nondifferentiable at certain points, however, and therefore are of no practical relevance that I can see.

5.6 Conclusions

Several results on the nominal stability of FIE, robust stability of FIE, robust exponential stability of FIE, and robust exponential stability of MHE have been presented. These results are of a broad enough class to cover unaltered least squares objective functions, unlike previous results on robust stability of FIE or MHE such as (Ji et al., 2016) or (Knüfer and Müller, 2018). They also guarantee convergence in the case of converging disturbances. Despite the considerable progress, however, there are still unresolved questions. The phenomenon of the smoothed estimates $e(j|k)$ increasing without bound in the example in Section 5.2.2 has not been, to my knowledge, noted before in the literature. Finding necessary and sufficient conditions to prevent that problem from occurring is an object for future research. With such conditions, it might be possible to prove that MHE with a smoothing prior is a stable robustly stable estimator.

CHAPTER 6

CONCLUSIONS AND FUTURE WORK

6.1 Conclusions

The principle contribution of this thesis has been to introduce a new type of Lyapunov-like function, termed Q function, for the analysis of FIE and MHE. When Q functions were first introduced in (Allan and Rawlings, 2019b), the results were phrased in a way to suggest that other state estimators might admit Q functions of their own. That would make Q functions a method of analyzing state estimation in a similar role to how Lyapunov functions fit in regulation. That paper was drafted before the results applying Q functions to regulation and tracking, which constitute major parts of Chapters 3 and 4, were known to me. Upon further reflection, it seems like they play an important role in analysis of *open-loop optimal*

control problems, in particular, the convergence of finite-horizon optimal control problems towards an infinite-horizon optimal control problem in which there is no difference between open-loop and closed-loop behavior. This new method of analysis has been applied to three problems.

The first, presented in Chapter 3 is one well-studied in the literature: setpoint regulation using MPC without stabilizing terminal constraints. Although well-studied, existing proofs seem complex and inaccessible to new researchers.¹ The central proof that produces a relaxed descent condition parallels that of a standard Lyapunov argument, with which all researchers should be familiar, and the end results qualitatively match those generated by Grüne and Pannek (2016), even if they produce a tighter bound on how long a horizon length is necessary to produce stability. The qualitative behavior is important, however, because it reveals the importance of choosing an appropriate stage cost and that the choice of an inappropriate stage cost might result in practical stability. However, it is argued that, from a practical standpoint, stabilizing terminal conditions are still superior, because the systems for which we can use a quadratic cost, increase the horizon length, and expect good results are precisely those for which we have a good method of designing a terminal region (at least given the other technical conditions of Proposition 3.30). Satisfaction of the terminal constraint is a verifiable condition that guarantees that the (typically local) NLP solver is making sufficient progress to guarantee robust stability.

The second, presented in Chapter 4 is much less studied in the literature. A

¹Or at least, I found them rather inaccessible when new to this area of research.

number of references (e.g., (Falugi and Mayne, 2013; Köhler et al., 2019)) have studied tracking a time-varying reference trajectory with MPC, and several ((Grimm et al., 2005) and (Rawlings et al., 2017, Sec. 2.4.4)) have considered semidefinite but detectable stage costs. To my knowledge, no work in the literature considers both at once. However, one must consider both to create a control problem that even approaches the complexity of the estimation problem. Construction of a Q function and Lyapunov function for the tracking problem, required a new element—an i-IOSS Lyapunov function that is a consequence of a detectability assumption.

In Chapter 2, i-IOSS and nonlinear detectability are explored in some detail. In particular, it is shown that a system admits a robustly stable state estimator only if it is i-IOSS. The converse Lyapunov theorem from (Allan et al., 2021) is related, and the argument from there is specialized to obtain an exponential i-IOSS Lyapunov function for exponentially i-IOSS systems. If the underlying system is Lipschitz, then so is the i-IOSS Lyapunov function.

Finally, Chapter 5 provides a clear payoff for these new methods of analysis. The proof of nominal stability of FIE using a Q function from (Allan and Rawlings, 2019b), which represents a culmination of effort from authors stretching back to Muske and Rawlings (1995), is provided. Furthermore, an extension of this Q function analysis to the problem of robust asymptotic stability of FIE was provided, settling (at least in part) Conjecture 13 from Rawlings and Ji (2012). The proof in the case of asymptotic stability required several nonintuitive assumptions, in particular the perturbed stabilizability assumption Assumption 5.14. Therefore the argument was specialized to the case of exponential detectability and stabilizability, as presented in (Allan

and Rawlings, 2021), for which we have a much more satisfactory characterization. Finally, in this exponential case, the robust stability of MHE can be concluded with moderate assumptions on the prior weighting.

It is worthwhile to summarize the state of optimization-based state estimation now. A system admits a robustly exponentially stable state estimator only if it is exponentially i-IOSS, and, furthermore, if it is exponentially incrementally stabilizable, FIE constitutes such a state estimator and there exists a horizon length sufficiently large that MHE is also robustly exponentially stable. A system admits a robustly asymptotically stable state estimator only if it is i-IOSS, but unfortunately there is no simple statement about for which systems FIE constitutes such a state estimator. Assumption 5.8 constitutes a good definition of asymptotic incremental stabilizability, and can be derived from a more general assumption in the form of Definition 3.1 by an argument similar to that used in Proposition 3.6. Unfortunately, I see no way to derive Assumption 5.14 from Assumption 5.8 except in a semiglobal fashion, which is insufficient for the purposes of a robust stability result like Theorem 5.18. Worse, I do not even see a way to articulate Assumption 5.14 in a way independent of the stage cost $\ell(\cdot)$, which makes it hard to even consider asking about the class of systems that satisfy this assumption.

6.2 Future work

6.2.1 *Is stabilizability really necessary for robust stability of FIE and MHE?*

One of the ways through which this advancement of theory of optimization-based state estimation was possible was the introduction of a stabilizability assumption. However, it does not seem like a necessary assumption, and is in fact not necessary in the case of the time-varying Kalman filter (Chan et al., 1984), although the convergence rate is degraded from exponential to harmonic if there are uncontrollable eigenvalues on the unit circle. It is worth noting, however, that the results for the Kalman filter that do not assume controllability or stabilizability came several decades after the initial results about estimator stability. Perhaps the use of stabilizability here can act as scaffolding for future work, and as more is learned about the estimation problem, it can be discarded as an unnecessary assumption. A more complete discussion about why or whether stabilizability is necessary is included in Section 5.3.3.

6.2.2 *Stability and robustness of MHE for asymptotically i -IOSS systems*

One of the limitations of the work presented here is that MHE is shown to be robustly stable only in the case of exponentially stabilizable and detectable systems. For asymptotically detectable and stabilizable systems, there is still a gap between FIE and MHE. The first gap is that the characterization of perturbed incremental asymptotic stabilizability given in Assumption 5.14 is not independent of the stage cost. Another gap, however, occurs when attempting to find an upper bound in

the estimation error in the filtering update. The proof of Theorem 5.30 relies on finding a horizon length long enough that estimate error is contracted, modulo disturbances. However, in the general case of asymptotic detectability, no such horizon length exists. One could, using the techniques outlined in this work, use a result like Proposition 3.20 to guarantee a *semiglobal practical* contraction given a sufficiently long horizon length. That would result in the state estimate ending up in a bounded neighborhood of the true state, but it would not necessarily converge in the case of convergent disturbances and sufficiently large disturbances would result in the system becoming unstable.

Tied up with this problem is proper design of the prior weighting for MHE. Here the relatively weak Assumption 5.29 has been used, requiring only that a filtering prior be used and that the prior weighting be uniformly bounded above and below. Theoretically, this property is unsatisfactory, however, because in the case of detectable linear systems, the arrival cost is uniformly bounded above if and only if the system is *controllable*. In the case of stabilizable linear systems, certitude about the state estimate in the uncontrollable subspace asymptotically goes to infinity. The fact that this element of the linear theory is excluded indicates that, at present, nonlinear theory is incomplete. Rao et al. (2003) introduced the criteria that the prior weighting be a positive definite *lower bound* to the arrival cost, in parallel to the condition in regulation that the terminal cost be an *upper bound* to the infinite horizon cost. However, it is not enough for the terminal cost to be an upper bound to the infinite horizon cost—it must furthermore be a *control Lyapunov function*. There needs to be some minimum amount of information carried ahead in the MHE

problem to guarantee convergence, and perhaps there is an estimation-equivalent of a CLF that could guarantee this information is carried forward.

However, designing a CLF is tractable in the regulation problem only because the system is actively being controlled to a setpoint. Furthermore, the only systematic method of designing a CLF and terminal region that I am aware of requires the linearization of the system at the setpoint to be stabilizable. Translated to the estimation problem, the system must be controlled to remain in a neighborhood of a setpoint whose linearization is detectable. In that case, it might be possible to prove that the EKF update of the prior retains sufficient information about the past state to ensure convergence. While this may be a good result, I suspect it would hold only in conditions Theorem 5.30 would apply anyway, and thus would do nothing towards ensuring asymptotically stable MHE in the case of asymptotically detectable systems.

6.2.3 *Smoothed estimates with FIE and smoothing prior update for MHE*

A related subject is whether or not a smoothing prior update for MHE could ensure robust stability. Although a filtering update is more statistically well-founded because it relays the most information about measurements before the start of the horizon, in practice it can induce oscillations in the state estimate because there is a closer relationship between $\hat{x}(k|k)$ and $\hat{x}(k-N|k-N)$ than $\hat{x}(k|k)$ and $\hat{x}(k-1|k-1)$. A smoothing update often provides much better performance in practice. However, existing theory is inadequate to analyze the behavior of a smoothing update. At the heart of the problem is that we cannot guarantee good

behavior of the smoothed estimates in FIE. As we can see in Section 5.2.2, there is no guarantee that the smoothed estimates even remain bounded as more information is taken into account in the FIE problem. In order for a smoothing estimate to work, not only must the smoothed estimate error be bounded, but it must be a contraction in some metric (or some generalization of a metric like the estimation-equivalent of a CLF). Careful analysis of using an EKF-like update for the prior in a neighborhood of a detectable setpoint may guarantee such a contraction. Findeisen (1997, Sec. 3.3.7) provides a Lyapunov function for constrained linear quadratic MHE with a smoothing update that might provide a starting point for future research.²

6.2.4 Return to probability and statistics

The Kalman filter was originally posed as a problem in probability and statistics. It produces the conditional mean when Gaussian noises and priors are used, and even in the case of non-Gaussian noise it is the best linear unbiased estimator. However, through the introduction of controllability and observability, it became possible to study it using system-theoretic methods, and the result that it constitutes a robustly stable estimator under those conditions followed. However, once the form of the FIE problem was motivated by maximum likelihood estimation, we left probability and statistics behind and analyzed it solely in system-theoretic terms. While this interpretation is valuable, it is also incomplete.

²Regrettably, I discovered this argument only in the last few months, and regrettably have not been able to afford it the attention it probably deserves. Repeated frustrations at replicating the descent condition in the Lyapunov function for the Kalman filter from (Deyst and Price, 1968) (and repeated in (Jazwinski, 1969, Th. 7.4)) have left me suspicious of accepting the conclusions of this sort of argument without careful inspection of the steps.

For instance, in Section 5.2.2, we noted that a bad choice of prior weighting and stage cost can cause the smoothed estimates of the initial state to diverge. Ideally, we could find conditions under which that phenomenon could not occur. However, that choice of prior weighting and stage cost could be a result of genuine problem statistics. If the prior is distributed by a Cauchy distribution

$$x(0) \sim \frac{1}{\pi(\chi^2 + 1)}$$

and the system noises are distributed according to Laplace distributions

$$w(k) \sim 2 \exp(-4 |\omega|) \quad v(k) \sim 2 \exp(-4 |\nu|)$$

then the FIE problem in Section 5.2.2 results. The fact that these are named distributions is an indication that such a problem may be reasonable in some context. The quotient of two normal distributions has a Cauchy distribution, and the fact that we use Laplace distributions rather than normal distributions for $w(k)$ and $v(k)$ is mostly for convenience—if they were normally distributed, the smoothing error would still diverge, but \widehat{w} and \widehat{v} would be nonzero.

The Cauchy distribution is somewhat notorious among statisticians, however, because it has an undefined mean and infinite variance. As a result, the law of large numbers does not apply to it (Pillai and Meng, 2016). However, different distributions can give rise to a similar maximum-likelihood problem. If we had that

$$x(0) \sim \frac{6\sqrt{3}}{\pi(\chi^2 + 3)^2},$$

which corresponds to Student's t distribution with three degrees of freedom, and

$$w(k) \sim 4 \exp(-8 |\omega|) \quad v(k) \sim 4 \exp(-8 |\nu|),$$

which are Laplace distributions. Now the mean and variance of $x(0)$ are both defined and the smoothing error still diverges.

What are we to make of this phenomenon? A point to reemphasize is that, under suitable conditions, FIE gives the maximum likelihood estimate of the *entire state trajectory* $\mathbf{x}(0 : k)$, which does not necessarily coincide with the sequence of maximum likelihood estimates of individual states. Perhaps the maximum likelihood estimate of $x(0)$ has better properties. Jazwinski (1969, p 157) would go further:

If these densities are Gaussian (linear problems), or just symmetric and unimodal, then their (respective) modes coincide with their (respective) means. In that case, the *maximum likelihood (Bayesian) estimate is the same as the minimum variance estimate* defined in Section 2. We note that the maximum likelihood estimate is of questionable value unless the density function is unimodal and concentrated near the mode.

In his view, then, the conditional mean is the best estimator, and the mode is something of a distraction. The conditional mean, however, is a problematic estimate as well, because it makes sense as a state estimate *only when the state domain \mathbb{X} is convex*. If \mathbb{X} is not convex, then the state estimate may fall outside of \mathbb{X} .³ Whether the conditional mean suffers from this issue is a question for future research.

³Note that this issue is distinct from (but related to) the issue of the conditional mean occurring in a region of zero probability as outlined in (Rawlings and Bakshi, 2006). My intuition is that, so long as the domain is convex, the conditional mean is the best estimate to use for MPC, even if it occurs in a place with zero probability. However, I know of no method to calculate the conditional mean that does not suffer from the curse of dimensionality. If the state space is nonconvex, then the conditional mean may lay outside of it, and an optimal control problem starting from it cannot even be defined.

6.2.5 Probabilistic interpretation of i-IOSS

A related issue is finding a probabilistic or statistical interpretation of an i-IOSS Lyapunov function. As we have seen, i-IOSS is a fundamental requirement for a nonlinear system to be detectable. Therefore, it seems like i-IOSS should show up in some form in the statistical literature. Unfortunately, I have not had the chance to engage with that literature in a meaningful way. A sort of heuristic interpretation of $\lambda(x_1, x_2)$ for some i-IOSS Lyapunov function $\lambda(\cdot)$ is that $\lambda(x_1, x_2)$ is the amount of “information” required to distinguish states x_1 and x_2 . The inequalities

$$\alpha_1(|x_1 - x_2|) \leq \lambda(x_1, x_2) \leq \alpha_2(|x_1 - x_2|),$$

for $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$, represent that because the states x_1 and x_2 are distinct, some “information” must be required to distinguish them, but the amount of “information” required to distinguish them must be proportional to the distance between the states. The dissipation inequality

$$\lambda(x_1^+, x_2^+) \leq \lambda(x_1, x_2) - \alpha_3(|x_1 - x_2|) + \sigma_w(|w_1 - w_2|) + \sigma_y(|y_1 - y_2|),$$

for $\alpha_3 \in \mathcal{K}_\infty$ and $\sigma_w, \sigma_y \in \mathcal{K}$, represents the fact that the information required to distinguish these two states must asymptotically shrink ($\alpha_3(\cdot)$) in the absence of disturbances that drive the states apart ($\sigma_w(\cdot)$) or “information” distinguishing the states that appears in the output ($\sigma_y(\cdot)$).

“Information” appears in quotation marks because whatever quantity $\lambda(x_1, x_2)$ represents is distinct from Shannon information. Any sort of systematic study of this matter should start with linear systems and Kalman filtering. In a cursory inspection of the literature, I found Duncan (1970) to be the best reference relating

information theory and Kalman filtering. More recent works like Venkat and Weissman (2012) work with scalar systems and outputs, which is inadequate for a full characterization of Kalman filtering. The work by Bucy (1979) on information content in a nonlinear estimation problem may also be of interest to a researcher trying to provide such a characterization of an i-IOSS Lyapunov function.

BIBLIOGRAPHY

- D. Aeyels. Generic Observability of Differentiable Systems. *SIAM J. Cont. Opt.*, 19(5):595–603, 1981. doi: 10.1137/0319037.
- M. Alamir and G. Bornard. Stability of a truncated infinite constrained receding horizon scheme: The general discrete nonlinear case. *Automatica*, 31:1352–1356, 1995.
- A. Alessandri, M. Baglietto, and G. Battistelli. Moving-horizon state estimation for nonlinear discrete-time systems: New stability results and approximation schemes. *Automatica*, 44(7):1753–1765, 2008.
- D. A. Allan and J. B. Rawlings. Inherent robustness of discontinuous MPC: Even (u, u^3) is robust. 6th IFAC Conference on Nonlinear Model Predictive Control NMPC 2018, Madison, WI, August 2018a.
- D. A. Allan and J. B. Rawlings. An input/output-to-state stability converse theorem for closed positive invariant sets. Technical Report 2018–01, TWCCC Technical Report, December 2018b. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2018-01.pdf>.
- D. A. Allan and J. B. Rawlings. Moving horizon estimation. In S. Raković and W. S. Levine, editors, *Handbook of Model Predictive Control*, pages 99 – 124. Springer, 2019a.
- D. A. Allan and J. B. Rawlings. A Lyapunov-like function for full information estimation. In *American Control Conference*, pages 4497–4502, Philadelphia, PA, July 10–12, 2019b.
- D. A. Allan and J. B. Rawlings. Robust stability of full information information. 2021. Submitted.

- D. A. Allan, C. N. Bates, M. J. Risbeck, and J. B. Rawlings. On the inherent robustness of optimal and suboptimal nonlinear MPC. *Sys. Cont. Let.*, 106:68 – 78, 2017. ISSN 0167-6911. doi: <https://doi.org/10.1016/j.sysconle.2017.03.005>. URL <http://www.sciencedirect.com/science/article/pii/S016769111730049X>.
- D. A. Allan, J. B. Rawlings, and A. R. Teel. Nonlinear detectability and incremental input/output-to-state stability. Technical Report 2020-01, TWCCC Technical Report, July 2020. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2020-01.pdf>.
- D. A. Allan, J. B. Rawlings, and A. R. Teel. Nonlinear detectability and incremental input/output-to-state stability. 2021. Submitted.
- B. Anderson and J. Moore. *Optimal Filtering*. Dover Books on Electrical Engineering. Dover Publications, 2012. ISBN 9780486136899. URL <https://books.google.com/books?id=iYMqLQp49UMC>.
- B. D. O. Anderson and J. B. Moore. Detectability and stabilizability of time-varying discrete-time linear systems. *SIAM J. Cont. Opt.*, 19(1):20–32, 1981.
- D. Angeli. A Lyapunov approach to incremental stability properties. *IEEE Trans. Auto. Cont.*, 47(3):410–421, 2002.
- D. Angeli. Further results on incremental input-to-state stability. *IEEE Trans. Auto. Cont.*, 54(6):1386–1391, June 2009. ISSN 0018-9286. doi: 10.1109/TAC.2009.2015561.
- B. W. Bequette. Nonlinear predictive control using multi-rate sampling. *Can. J. Chem. Eng.*, 69:136–143, February 1991.
- R. R. Bitmead, M. R. Gevers, I. R. Petersen, and R. J. Kaye. Monotonicity and stabilizability properties of solutions of the Riccati difference equation: Propositions, lemmas, theorems, fallacious conjectures and counterexamples. *Sys. Cont. Let.*, 5: 309–315, 1985.
- R. W. Brockett. System Theory on Group Manifolds and Coset Spaces. *SIAM J. Cont.*, 10(2):265–284, 1972. doi: 10.1137/0310021.
- A. E. Bryson and M. Frazier. Smoothing for linear and nonlinear dynamic systems. Proc. Optimum Systems Synthesis Conf., U.S. Air Force Tech. Rept. ASD-TDR-063-119, February 1963.

- R. S. Bucy. Information and filtering. *Information Sciences*, 18(2):179–187, 1979. ISSN 0020-0255. doi: [https://doi.org/10.1016/0020-0255\(79\)90014-8](https://doi.org/10.1016/0020-0255(79)90014-8). URL <http://www.sciencedirect.com/science/article/pii/0020025579900148>.
- C. I. Byrnes and A. Isidori. Output regulation for nonlinear systems: an overview. *Int. J. Robust and Nonlinear Control*, 10(5):323–337, 2000.
- C. Cai and A. R. Teel. Input–output-to-state stability for discrete-time systems. *Automatica*, 44(2):326 – 336, 2008. ISSN 0005-1098. doi: <http://dx.doi.org/10.1016/j.automatica.2007.05.022>. <http://www.sciencedirect.com/science/article/pii/S0005109807003056>.
- D. A. Carlson, A. B. Haurie, and A. Leizarowitz. *Infinite Horizon Optimal Control*. Springer Verlag, second edition, 1991.
- S. W. Chan, G. C. Goodwin, and K. S. Sin. Convergence properties of the Riccati difference equation in optimal filtering of nonstabilizable systems. *IEEE Trans. Auto. Cont.*, 29(2):110–118, 1984.
- C. C. Chen and L. Shaw. On receding horizon feedback control. *Automatica*, 16(3): 349–352, 1982.
- H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34(10):1205–1217, 1998.
- E. Chu, A. Keshavarz, D. Gorinevsky, and S. Boyd. Moving horizon estimation for staged QP problems. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 3177–3182, dec 2012. doi: 10.1109/CDC.2012.6425976.
- E. A. Coddington and N. Levinson. *Theory of Ordinary Differential Equations*. McGraw Hill, 1955.
- H. Cox. On the estimation of state variables and parameters for noisy dynamic systems. *IEEE Trans. Auto. Cont.*, 9(1):5–12, 1964.
- G. De Nicolao, L. Magni, and R. Scattolini. On the robustness of receding-horizon control with terminal constraints. *IEEE Trans. Auto. Cont.*, 41(3):451–453, March 1996.
- G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Trans. Auto. Cont.*, 43(7):1030–1036, 1998.

- J. Deyst and C. Price. Conditions for asymptotic stability of the discrete minimum-variance linear estimator. *IEEE Transactions on Automatic Control*, 13(6):702–705, dec 1968. doi: 10.1109/TAC.1968.1099024.
- Z. Dong and D. Angeli. Analysis of economic model predictive control with terminal penalty functions on generalized optimal regimes of operation. *Int. J. Robust and Nonlinear Control*, 28(16):4790–4815, 2018. ISSN 1049-8923. doi: 10.1002/rnc.4283. URL <https://browzine.com/articles/240212433>.
- A. L. Dontchev and R. T. Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer Monographs in Mathematics. Springer New York, 2014. ISBN 9781493910366. doi: 10.1007/978-1-4939-1037-3. URL <https://sites.math.washington.edu/~rtr/papers/rtr234-Implicit2.pdf>.
- T. E. Duncan. On the Calculation of Mutual Information. *SIAM J. Appl. Math.*, 19(1): 215–220, 1970. doi: 10.1137/0119020.
- P. Falugi and D. Q. Mayne. Tracking a periodic reference using nonlinear model predictive control. In *52nd IEEE Conference on Decision and Control*, pages 5096–5100, 2013. doi: 10.1109/CDC.2013.6760689.
- P. K. Findeisen. Moving horizon state estimation of discrete time systems. Master’s thesis, University of Wisconsin–Madison, 1997.
- G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel. Examples when nonlinear model predictive control is nonrobust. *Automatica*, 40:1729–1738, 2004.
- G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel. Model predictive control: For want of a local control Lyapunov function, all is not lost. *IEEE Trans. Auto. Cont.*, 50(5):546–558, 2005.
- J. W. Grizzle and P. E. Moraal. Observer based control of nonlinear discrete-time systems. In *Proceedings of the 29th IEEE Conference on Decision and Control*, pages 760–767, Honolulu, 1990.
- L. Grüne and C. M. Kellett. ISS-Lyapunov functions for discontinuous discrete-time systems. *IEEE Transactions on Automatic Control*, 59(11):3098–3103, Nov 2014. ISSN 0018-9286. doi: 10.1109/TAC.2014.2321667.
- L. Grüne and D. Nešić. Optimization based stabilization of sampled-data nonlinear systems via their approximate discrete-time models. *42(1):98–122*, 2003. ISSN 03630129. doi: 10.1137/S036301290240258X.

- L. Grüne and J. Pannek. Analysis of unconstrained NMPC schemes with incomplete optimization *. *IFAC Proceedings Volumes*, 43(14):238–243, 2010. ISSN 1474-6670. doi: <https://doi.org/10.3182/20100901-3-IT-2016.00124>. URL <http://www.sciencedirect.com/science/article/pii/S147466701536972X>.
- L. Grüne and J. Pannek. *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer, second edition, 2016.
- L. Grune and A. Rantzer. On the Infinite Horizon Performance of Receding Horizon Controllers. 53(9):2100–2111, 2008.
- R. Hermann and A. J. Krener. Nonlinear controllability and observability. *IEEE Trans. Auto. Cont.*, 22(5):728–740, 1977.
- J. P. Hespanha. *Lineary Systems Theory*. Princeton University Press, Princeton and Oxford, second edition, 2018.
- Y. C. Ho. The Method of Least Squares and Optimal Filtering Theory. Technical report, RAND Corporation, 1962. URL <http://www.dtic.mil/docs/citations/AD0286742>.
- W. Hu. Robust Stability of Optimization-based State Estimation Under Bounded Disturbances. *ArXiv e-prints*, Feb 2017.
- W. Hu, L. Xie, and K. You. Optimization-based state estimation under bounded disturbances. In *2015 54th IEEE Conference on Decision and Control CDC*, pages 6597–6602, Dec 2015. doi: 10.1109/CDC.2015.7403258.
- A. Isidori and C. I. Byrnes. Output regulation of nonlinear systems. *IEEE Trans. Auto. Cont.*, 35(2):131–140, February 1990.
- A. Jadbabaie and J. Hauser. On the stability of receding horizon control with a general terminal cost. 50(5):674–678, 2005.
- S.-S. Jang, B. Joseph, and H. Mukai. Comparison of two approaches to on-line parameter and state estimation of nonlinear systems. *Ind. Eng. Chem. Proc. Des. Dev.*, 25:809–814, 1986.
- A. Jazwinski. Adaptive filtering. *Automatica*, 5(4):475–485, 1969.
- A. H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, New York, 1970.

- L. Ji, J. B. Rawlings, W. Hu, A. Wynn, and M. M. Diehl. Robust stability of moving horizon estimation under bounded disturbances. *IEEE Trans. Auto. Cont.*, 61(11): 3509–3514, November 2016.
- Z.-P. Jiang and Y. Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37:857–869, 2001.
- S. J. Julier and J. K. Uhlmann. Unscented filtering and nonlinear estimation. *Proc. IEEE*, 92(3):401–422, March 2004.
- T. Kailath. A view of three decades of linear filtering theory. *IEEE Trans. Inform. Theory*, IT-20(2):146–181, March 1974.
- R. E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME, J. Basic Engineering*, pages 35–45, March 1960a.
- R. E. Kalman. Contributions to the theory of optimal control. *Bull. Soc. Math. Mex.*, 5:102–119, 1960b.
- R. E. Kalman. New methods in Wiener filtering theory. In *Proceedings of the First Symposium on Engineering Applications of Random Function Theory and Probability*, edited by JL Bogdanoff and F. Kozin, John Wiley & Sons, New York, 1963.
- S. S. Keerthi and E. G. Gilbert. An existence theorem for discrete-time infinite-horizon optimal control problems. *IEEE Trans. Auto. Cont.*, 30(9):907–909, September 1985.
- S. S. Keerthi and E. G. Gilbert. Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations. *J. Optim. Theory Appl.*, 57(2):265–293, May 1988.
- C. M. Kellett and A. R. Teel. Discrete-time asymptotic controllability implies smooth control-Lyapunov function. *Sys. Cont. Let.*, 52:349–359, 2004a.
- C. M. Kellett and A. R. Teel. Smooth Lyapunov functions and robustness of stability for difference inclusions. *Sys. Cont. Let.*, 52:395–405, 2004b.
- C. M. Kellett and A. R. Teel. Weak converse Lyapunov theorems and control-Lyapunov functions. *SIAM J. Cont. Opt.*, 42(6), 2004c. ISSN 03630129. URL <http://search.proquest.com/docview/925801102/>.

- H. K. Khalil. High-Gain Observers in Feedback Control: Application to Permanent Magnet Synchronous Motors. *IEEE Ctl. Sys. Mag.*, 37(3):25–41, jun 2017. ISSN 1066-033X. doi: 10.1109/MCS.2017.2674438.
- D. L. Kleinman. An easy way to stabilize a linear constant system. *IEEE Trans. Auto. Cont.*, 15(12):692, December 1970.
- S. Knüfer and M. A. Müller. Robust Global Exponential Stability for Moving Horizon Estimation. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 3477–3482, Dec 2018. doi: 10.1109/CDC.2018.8619617.
- J. Köhler, M. A. Müller, and F. Allgöwer. Nonlinear reference tracking: An economic model predictive control perspective. *IEEE Trans. Auto. Cont.*, 64(1):254–269, 2019.
- A. J. Krener. The existence of optimal regulators. In *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No.98CH36171)*, volume 3, pages 3081–3086 vol.3, 1998.
- M. Krichman, E. D. Sontag, and Y. Wang. Input-output-to-state stability. *SIAM J. Cont. Opt.*, 39(6):1874–1928, 2001.
- W. H. Kwon and A. E. Pearson. On feedback stabilization of time-varying discrete linear systems. *IEEE Trans. Auto. Cont.*, 23(3):479–481, June 1978.
- W. H. Kwon, A. M. Bruckstein, and T. Kailath. Stabilizing state-feedback design via the moving horizon method. *Int. J. Control*, 37(3):631–643, 1983.
- M. J. Liebman, T. F. Edgar, and L. S. Lasdon. Efficient data reconciliation and estimation for dynamic processes using nonlinear programming techniques. *Comput. Chem. Eng.*, 16(10/11):963–986, 1992.
- D. Limon, M. Pereira, D. Munoz De La Pena, T. Alamo, C. N. Jones, and M. N. Zeilinger. MPC for Tracking Periodic References. 61(4):1123–1128, apr 2016. ISSN 00189286. doi: 10.1109/TAC.2015.2461811.
- L. Magni, G. De Nicolao, and R. Scattolini. Output feedback and tracking of nonlinear systems with model predictive control. *Automatica*, 37(10):1601–1607, 2001. ISSN 0005-1098. doi: [https://doi.org/10.1016/S0005-1098\(01\)00102-9](https://doi.org/10.1016/S0005-1098(01)00102-9). URL <http://www.sciencedirect.com/science/article/pii/S0005109801001029>.
- D. Mayne. An apologia for stabilising terminal conditions in model predictive control. *Int. J. Control*, 86(11):2090–2095, 2013.

- D. Q. Mayne and H. Michalska. Receding horizon control of non-linear systems. In *Proceedings of the 27th Conference on Decision and Control*, pages 464–465, December 1988.
- D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- E. S. Meadows and J. B. Rawlings. Receding horizon control with an infinite horizon. In *Proceedings of the 1993 American Control Conference*, San Francisco, California, June 1993. American Automatic Control Council.
- E. S. Meadows, M. A. Henson, J. W. Eaton, and J. B. Rawlings. Receding horizon control and discontinuous state feedback stabilization. *Int. J. Control*, 62(5): 1217–1229, 1995.
- H. Michalska and D. Q. Mayne. Moving Horizon Observers. *IFAC Proceedings Volumes*, 25(13):185–190, 1992. ISSN 1474-6670. doi: [https://doi.org/10.1016/S1474-6670\(17\)52279-X](https://doi.org/10.1016/S1474-6670(17)52279-X). URL <http://www.sciencedirect.com/science/article/pii/S147466701752279X>.
- H. Michalska and D. Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Trans. Auto. Cont.*, 38(11):1623–1633, 1993.
- S. A. Morris. *Topology Without Tears*. 2020. URL <http://www.topologywithouttears.net/>.
- M. A. Müller. Nonlinear moving horizon estimation in the presence of bounded disturbances. *Automatica*, 79:306–314, 2017. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automat.2017.01.033>. URL <http://www.sciencedirect.com/science/article/pii/S0005109817300432>.
- M. A. Müller and K. Worthmann. Quadratic costs do not always work in MPC. *Automatica*, 82:269–277, 2017. ISSN 00051098. doi: 10.1016/j.automat.2017.04.058. URL <http://dx.doi.org/10.1016/j.automat.2017.04.058>.
- K. R. Muske and J. B. Rawlings. Model predictive control with linear models. *AIChE J.*, 39(2):262–287, 1993.
- K. R. Muske and J. B. Rawlings. Nonlinear moving horizon state estimation. In R. Berber, editor, *Methods of Model Based Process Control*, NATO Advanced Study Institute series: E Applied Sciences 293, pages 349–365, Dordrecht, The Netherlands, 1995. Kluwer.

- K. R. Muske, J. B. Rawlings, and J. H. Lee. Receding horizon recursive state estimation. In *Proceedings of the 1993 American Control Conference*, pages 900–904, June 1993.
- A. Pavlov and L. Marconi. Incremental passivity and output regulation. *Sys. Cont. Let.*, 57(5):400–409, 2008. ISSN 0167-6911. doi: <https://doi.org/10.1016/j.sysconle.2007.10.008>. URL <http://www.sciencedirect.com/science/article/pii/S0167691107001417>.
- N. S. Pillai and X.-L. Meng. An unexpected encounter with Cauchy and Lévy. *Ann. Statist.*, 44(5):2089–2097, 2016. doi: [10.1214/15-AOS1407](https://doi.org/10.1214/15-AOS1407). URL <https://doi.org/10.1214/15-AOS1407>.
- E. Polak. *Optimization: Algorithms and Consistent Approximations*. Springer Verlag, New York, 1997. ISBN 0-387-94971-2.
- J. A. Primbs and V. Nevistić. Feasibility and stability of constrained finite receding horizon control. *Automatica*, 36:965–971, 2000.
- R. Rajamani. Observers for nonlinear systems: Introduction to part 1 of the special issue. *IEEE Ctl. Sys. Mag.*, 37(3):22–24, June 2017a. ISSN 1066-033X. doi: [10.1109/MCS.2017.2674400](https://doi.org/10.1109/MCS.2017.2674400).
- R. Rajamani. Observers for nonlinear systems : Part 2: An overview of the special issue. *IEEE Ctl. Sys. Mag.*, 37(4):30–32, Aug 2017b. ISSN 1066-033X. doi: [10.1109/MCS.2017.2696758](https://doi.org/10.1109/MCS.2017.2696758).
- C. V. Rao. *Moving Horizon Strategies for the Constrained Monitoring and Control of Nonlinear Discrete-Time Systems*. PhD thesis, University of Wisconsin–Madison, 2000.
- C. V. Rao, J. B. Rawlings, and J. H. Lee. Constrained linear state estimation – a moving horizon approach. *Automatica*, 37(10):1619–1628, 2001.
- C. V. Rao, J. B. Rawlings, and D. Q. Mayne. Constrained state estimation for nonlinear discrete-time systems: stability and moving horizon approximations. *IEEE Trans. Auto. Cont.*, 48(2):246–258, February 2003.
- J. B. Rawlings and B. R. Bakshi. Particle filtering and moving horizon estimation. *Comput. Chem. Eng.*, 30:1529–1541, 2006.

- J. B. Rawlings and L. Ji. Optimization-based state estimation: Current status and some new results. *J. Proc. Cont.*, 22:1439–1444, 2012.
- J. B. Rawlings and D. Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, Madison, WI, 2009. 668 pages, ISBN 978-0-9759377-0-9.
- J. B. Rawlings and K. R. Muske. Stability of constrained receding horizon control. *IEEE Trans. Auto. Cont.*, 38(10):1512–1516, October 1993.
- J. B. Rawlings and M. J. Risbeck. Model predictive control with discrete actuators: Theory and application. *Automatica*, 78:258–265, 2017.
- J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Design, and Computation*. Nob Hill Publishing, Madison, WI, 2nd edition, 2017. 770 pages, ISBN 978-0-9759377-3-0.
- D. G. Robertson, J. H. Lee, and J. B. Rawlings. A moving horizon-based approach for least-squares state estimation. *AIChE J.*, 42(8):2209–2224, August 1996.
- W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, New York, third edition, 1976.
- H. Santoso, D. Hioe, J. Bao, and P. L. Lee. Operability analysis of nonlinear processes based on incremental dissipativity. *J. Proc. Cont.*, 22(1):156–166, Jan 2012. ISSN 0959-1524. doi: 10.1016/J.JPROCONT.2011.09.006. URL <https://www.sciencedirect.com/science/article/pii/S0959152411001910?via%3Dihub>.
- P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings. Suboptimal model predictive control (feasibility implies stability). *IEEE Trans. Auto. Cont.*, 44(3):648–654, March 1999.
- E. D. Sontag. A Lyapunov-Like Characterization of Asymptotic Controllability. 21(3):462–471, 1983. doi: 10.1137/0321028.
- E. D. Sontag. A concept of local observability. *Sys. Cont. Let.*, 5(1):41–47, 1984. ISSN 0167-6911. doi: [https://doi.org/10.1016/0167-6911\(84\)90007-0](https://doi.org/10.1016/0167-6911(84)90007-0). URL <http://www.sciencedirect.com/science/article/pii/0167691184900070>.
- E. D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, Apr 1989. ISSN 0018-9286. doi: 10.1109/9.28018.

- E. D. Sontag. *Mathematical Control Theory*. Springer-Verlag, New York, second edition, 1998.
- E. D. Sontag and Y. Wang. On the characterization of the input to state stability property. *Sys. Cont. Let.*, 24:351–359, 1995.
- E. D. Sontag and Y. Wang. Output-to-state stability and detectability of nonlinear systems. *Sys. Cont. Let.*, 29:279–290, 1997.
- Springer Verlag GmbH, European Mathematical Society. Semicontinuous function. Website. URL http://encyclopediaofmath.org/index.php?title=Semicontinuous_function&oldid=40145. Accessed June 4, 2020.
- H. J. Sussmann. Existence and uniqueness of minimal realizations of nonlinear systems. *Mathematical systems theory*, 10(1):263–284, 1976. ISSN 1433-0490. doi: 10.1007/BF01683278. URL <https://doi.org/10.1007/BF01683278>.
- P. Swerling. First-Order Error Propagation in a Stagewise Smoothing Procedure for Satellite Observations, 1959.
- M. J. Tenny. *Computational Strategies for Nonlinear Model Predictive Control*. PhD thesis, University of Wisconsin–Madison, October 2002. URL <http://jbrwww.che.wisc.edu/theses/tenny.ps>.
- Y. A. Thomas. Linear quadratic optimal estimation and control with receding horizon. *Electron. Lett.*, 11:19–21, January 1975.
- D. N. Tran, B. S. Rüffer, and C. M. Kellett. Incremental stability properties for discrete-time systems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 477–482, 2016.
- S. E. Tuna, M. J. Messina, and A. R. Teel. Shorter horizons for model predictive control. In *2006 American Control Conference*, 2006.
- K. Venkat and T. Weissman. Pointwise Relations Between Information and Estimation in Gaussian Noise. *IEEE Trans. Inform. Theory*, 58(10):6264–6281, Oct 2012. ISSN 1557-9654. doi: 10.1109/TIT.2012.2206794.
- W. M. Wonham. On a Matrix Riccati Equation of Stochastic Control. *SIAM J. Cont.*, 6(4):681–697, 1968a. doi: 10.1137/0306044.

- W. M. Wonham. On the Separation Theorem of Stochastic Control. *SIAM J. Cont.*, 6(2):312–326, 1968b. doi: 10.1137/0306023.
- X. Yang, D. W. Griffith, and L. T. Biegler. Nonlinear programming properties for stable and robust NMPC. *IFAC-P. Online*, 48(23):388 – 397, 2015. ISSN 2405-8963. doi: <http://dx.doi.org/10.1016/j.ifacol.2015.11.310>. URL <http://www.sciencedirect.com/science/article/pii/S240589631502594X>.
- S. Yu, M. Reble, H. Chen, and F. Allgöwer. Inherent robustness properties of quasi-infinite horizon nonlinear model predictive control. *Automatica*, 50(9):2269 – 2280, 2014. ISSN 0005-1098. doi: <http://dx.doi.org/10.1016/j.automatica.2014.07.014>.
- M. Zanon and T. Faulwasser. Economic MPC without terminal constraints: Gradient-correcting end penalties enforce asymptotic stability. *J. Proc. Cont.*, 63: 1–14, 2018.
- G. Zimmer. State observation by on-line minimization. *Int. J. Control*, 60(4):595–606, 1994.