

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Towards a Turnkey Model Predictive Controller: Identification, Application, and Theory

Permalink

<https://escholarship.org/uc/item/9q0155dd>

Author

Kuntz, Steven J.

Publication Date

2024

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

**Towards a Turnkey Model Predictive Controller:
Identification, Application, and Theory**

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

Doctor of Philosophy
in
Chemical Engineering

by

Steven J. Kuntz

Committee in charge:

Professor James B. Rawlings, Chair
Professor Andrew Teel
Professor Sho Takatori
Professor Michael F. Doherty

December 2024

The Dissertation of Steven J. Kuntz is approved.

Professor Andrew Teel

Professor Sho Takatori

Professor Michael F. Doherty

Professor James B. Rawlings, Committee Chair

December 2024

Towards a Turnkey Model Predictive Controller: Identification, Application, and Theory

Copyright © 2024

by

Steven J. Kuntz

To my parents.

Acknowledgments

I have had the good fortune, throughout my Ph.D., to be surrounded by an impressive group of colleagues. My advisor, James B. Rawlings, has provided excellent guidance and advice throughout my time in his research group at UCSB. I have appreciated the freedom to investigate many disparate topics, through which I was able to explore most of the standard stack of process controls problems as well as technical aspects of the linear algebra, probability, and optimization theory used in these problems. I will be the first to admit that my research practices often suggest more “madness” than “method.” But I will be forever grateful for Jim’s continuing patience as my research practices matured, and the exploration paid off.

My first real experience with the Rawlings group was through a presentation by Jim on Doug Allan’s research that I attended during advisor selection. The first thing Jim told me was (paraphrasing): “You won’t understand any of this now, but in a year, you’ll understand half of it, and in two, you’ll understand all of it.” That talk turned out to be the most complicated theory result I had and have ever encountered in controls research: a converse Lyapunov theorem for state estimator stability. The first thing Doug said after it ended was (paraphrasing): “Well, Steven is never going to attend a group meeting again.” I don’t think either of them were quite right. The vast majority of the presentation was lost on me, but I attended the remaining group meetings for the quarter and found the bits I was able to pick up interesting and exciting. It motivated me to self-teach myself linear systems theory over the coming winter break so I could take Andrew Teel’s nonlinear control course in the winter quarter. I owe thanks to Doug, as well as Koty McAllister, for having some of the most interesting theory presentations and discussions. Early exposure to their theory work was indispensable for developing Part III of this dissertation. Travis Arnold overlapped with me for a short time, but I inherited a stack of difficult linear algebra problems that became an extremely long technical report, written over two winter breaks of my Ph.D. I owe much of my linear algebra knowledge to him. Chris Kuo-LeBlanc, Davide Mannini, and Titus Quah have been great friends and colleagues throughout my Ph.D. I will always cherish climbing sessions with Chris and Titus. I wish them, as well as Prithvi Dake, the best of luck in their future endeavors. I would also like to thank Evan Pretti and Timothy Quah for being great friends and sharing with me their interesting research problems throughout our graduate student experience.

Thanks are also due to Andrew Teel for serving on my committee and for teaching the excellent nonlinear controls course mentioned above. I also thank Sho Takatori and Michael F. Doherty for serving on my committee. Special thanks goes to Steve Miller and Jim Downs of Eastman Chemical Company for guiding me through an internship working (literally) at the center of a chemical plant. Part II of this dissertation would not have half the relevance without Steve diligently working to find a system to apply Turnkey MPC to and Jim advising me through the process to get our experiment approved (or putting his seal of approval on it, or winking and nudging the operations staff, I will leave that to interpretation). I also owe thanks to Melissa Peterson for mentoring my visit and helping me implement my methods in Eastman’s control software, and the rest of the Advanced Controls Team and other staff at Eastman who were nothing but kind and hospitable for my three month visit. Moritz Diehl from the University of Stuttgart in Germany, as well as his research group, deserve a special

thanks. I immensely enjoyed assisting with the short (and chaotic) MPC course he and Jim hosted over spring break. I always learn something new when I talk to him and his group members.

To Paulina, you deserve the most credit for making my Ph.D. experience enjoyable. You put up with long hours, frequent trips, and incomprehensible rants on real analysis. You always supported me every step of the way and ensured I slept and went outside during the final stretch. To my parents, I am forever grateful for your love and support, always providing an outlet for my interests in all things scientific and mathematical, even from a young age. I doubt I would have realized this endeavor if not for you supporting my interests every step of the way, even (and especially) when it became incomprehensible.

Steven J. Kuntz
Santa Barbara, CA
December 2024

Curriculum Vitæ

Steven J. Kuntz

Education

2019–2024 University of California, Santa Barbara (UCSB) – Santa Barbara, CA
Ph.D. in Chemical Engineering (Expected)

2015–2019 University of Delaware (UD) – Newark, DE
Bachelors of Chemical Engineering

Publications

In progress S. J. Kuntz and J. B. Rawlings. Offset-free model predictive control: stability under plant-model mismatch. *IEEE Trans. Auto. Cont.*, 2025. In progress

Under review S. J. Kuntz and J. B. Rawlings. Beyond inherent robustness: strong stability of MPC despite plant-model mismatch. *IEEE Trans. Auto. Cont.*, 2024c. Submitted 11/22/2024

Under review S. J. Kuntz and J. B. Rawlings. Maximum Likelihood Identification of Linear Models with Integrating Disturbances for Offset-Free Control. *IEEE Trans. Auto. Cont.*, 2024a. Submitted 6/5/2024, Revised 11/4/2024

S. J. Kuntz and J. B. Rawlings. On linear multivariate regression with singular error covariances. Technical Report 2023–02, TWCCC Technical Report, December 2023b. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2023-02.pdf>

S. J. Kuntz and J. B. Rawlings. On the unified theory of linear Gaussian estimation: solution methods, applications, and extensions. Technical Report 2023–01, TWCCC Technical Report, December 2023a. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2023-01.pdf>

S. J. Kuntz, J. J. Downs, S. M. Miller, and J. B. Rawlings. An industrial case study on the combined identification and offset-free control of a chemical process. *Comput. Chem. Eng.*, 179, 2023. doi: <https://doi.org/10.1016/j.compchemeng.2023.108429>

S. J. Kuntz and J. B. Rawlings. Maximum likelihood estimation of linear disturbance models for offset-free model predictive control. In *American Control Conference*, pages 3961–3966, Atlanta, GA, June 8–10, 2022

Research

Sep. 2019–Aug. 2024 Ph.D. Candidate, University of California, Santa Barbara
Mentor: Professor James B. Rawlings

Mar. 2022–June 2022 Advanced Controls Intern, Eastman Chemical Company
Mentors: Stephen M. Miller, Ph.D.; James J. Downs, Ph.D.

Aug. 2017–May 2019 Senior Thesis Research, University of Delaware
Mentor: Professor Dionisios Vlachos

Awards

2023 Selected for CAST Division Plenary Session (AIChE)

2023 Schlinger Fellowship for Excellence in Chemical Engineering Research
(UCSB)

Abstract

Towards a Turnkey Model Predictive Controller: Identification, Application, and Theory

by

Steven J. Kuntz

An out-of-the-box model predictive control (MPC) algorithm, or a “turnkey” model predictive controller has long been a dream of both academics and practitioners. MPC practice currently includes time-consuming and ad hoc tuning steps to achieve adequate performance in the face of persistent disturbances and plant-model mismatch. In this dissertation, we present progress towards developing a turnkey model predictive controller by developing identification methods suitable for out-of-the-box MPC implementations, applying those identification methods to the offset-free control of real-world systems, and developing the theory of the stability of MPC under plant-model mismatch.

In the first part of this dissertation, we propose algorithms for identifying plant and disturbance models. Maximum likelihood (ML) estimation methods are applied directly and in a nested fashion to identify complete plant and disturbance models. For the direct methods, high-level design constraints are imposed on the resulting offset-free controller through eigenvalue constraints on the modeled system matrices. For the nested methods, we present simple algorithms with closed-form solutions that can easily be implemented by practitioners.

In the second part of this dissertation, we apply identification methods to the offset-free control of two real-world systems: a benchmark temperature controller (TCLab), and an industrial-scale chemical reactor. Both case studies showcase the ability of the identification algorithms to produce models adequate for out-of-the-box MPC designs with guaranteed offset-free performance. The industrial application also demonstrates an outsize real-world benefit for adopting a turnkey approach, where we report a 38% improvement in setpoint

tracking performance compared to an existing hand-tuned controller.

In the third and final part of this dissertation, we investigate the theoretical properties of offset-free MPC subject to plant-model mismatch. We first investigate the offset-free performance of linear offset-free MPC for control of nonlinear plants. We then investigate stability of standard MPC under mismatch when the plant and model steady states are fixed and aligned. Finally, we investigate the offset-free performance of nonlinear offset-free MPC, with and without plant-model mismatch.

Contents

Curriculum Vitae	vii
Abstract	ix
List of Figures	xiv
List of Tables	xvii
1 Introduction	1
1.1 System identification	2
1.2 Model predictive control	5
1.3 Towards a turnkey model predictive controller	8
1.4 Notation and basic definitions	11
2 Offset-free model predictive control	15
2.1 Linear systems	15
2.2 Linear model predictive control	17
2.3 Linear offset-free model predictive control	21
2.4 Nonlinear offset-free model predictive control	26
Part I Identification	31
3 Constrained maximum likelihood identification for offset-free control	32
3.1 Maximum likelihood estimation	34
3.2 Problem statement	35
3.3 Algorithm outline	44
3.4 Eigenvalue constraints	52
3.5 Cholesky substitution and elimination	57
3.6 Summary and discussion	63
3.A Additional proofs	67

4	Maximum likelihood estimation of disturbance models	81
4.1	Literature review	82
4.2	Simple methods for disturbance model identification	85
4.3	Closed-loop subspace identification	95
 Part II Application		103
5	Case studies in combined identification and offset-free control	104
5.1	Systems of interest	104
5.2	Maximum likelihood identification	107
5.3	An industrial case study on the combined identification and offset-free control of a chemical process	118
 Part III Theory		130
6	Indirect methods for linear control of nonlinear systems	131
6.1	Lyapunov's indirect method	132
6.2	General indirect methods	137
6.3	Unconstrained linear control of nonlinear systems	143
6.4	Other considerations	156
6.A	Derivatives and \mathcal{K} -functions	158
7	Stability of model predictive control despite plant-model mismatch	160
7.1	Problem statement	163
7.2	Robust and strong stability	171
7.3	Inherent robustness of MPC	174
7.4	Stability of MPC despite mismatch	175
7.5	Examples	185
7.6	Conclusion	192
7.A	Additional proofs	194
7.B	Additional examples details	208
8	Stability of offset-free MPC despite plant-model mismatch	220
8.1	Problem statement	222
8.2	Robust stability for tracking and estimation	234
8.3	Nominal offset-free performance	241
8.4	Offset-free performance without mismatch	243
8.5	Offset-free MPC under mismatch	249
8.6	Examples	255
8.7	Conclusions	268
8.A	Proofs of robust estimation and tracking stability	270

8.B	Proofs of offset-free MPC stability	275
8.C	Steady-state target problem assumptions	295
9	Conclusion	302
9.1	Summary	302
9.2	Future work	304
	Bibliography	307

List of Figures

1.1	Process control heierarchy.	2
1.2	Setpoint tracking MPC.	5
1.3	Offset-free MPC.	7
1.4	Conceptual diagram of a turnkey model predictive controller.	9
4.1	Disturbance model identification.	86
4.2	Closed-loop regularized subspace identification.	96
5.1	Benchmark temperature Control Laboratory (TCLab) Park et al. (2020).	105
5.2	Schematic of the DMT reactor and MPC control strategy.	106
5.3	TCLab identification data	109
5.4	TCLab models open-loop and closed-loop (filter) eigenvalues.	110
5.5	TCLab setpoint tracking tests.	111
5.6	TCLab disturbance rejection tests.	111
5.7	TCLab test performance	112
5.8	TCLab identification index data	113
5.9	Training data for the Eastman reactor	114
5.10	Eastman reactor models open-loop and closed-loop (filter) eigenvalues.	115
5.11	Test performance for the Eastman reactor models	116
5.12	Eastman reactor models' closed-loop (filter) response to the eigenvector corresponding to the fastest eigenvalue.	117
5.13	Eastman reactor models simulated closed-loop test performance.	117
5.14	Disturbance estimates and driving noise estimates for the Eastman reactor models with the MPC variables and raw sensor data	119
5.15	Training data and model fits for the Eastman reactor models with the MPC variables and raw sensor data	120
5.16	Disturbance estimators and driving noise for the unregularized and regularized Eastman reactor models	121
5.17	Training data and model fits for the unregularized and regularized Eastman reactor models	122

5.18	Step responses of the (left) unregularized and (right) regularized models compared to the step responses of the previous MPC model.	123
5.19	Kalman filtered disturbances from the closed-loop tests of the (left) previous and (right) new models using the MPC variables as measurements.	125
5.20	Closed-loop performance model of the previous and new models using the MPC variables as measurements	126
5.21	Closed-loop comparison of the (left) previous and (right) new models using the MPC variables.	127
5.22	Closed-loop comparison of the (left) previous and (right) new models using the raw process variables as measurements.	128
7.1	Cost difference contours for the MPC of $x^+ = x + (1 + \theta)u$	170
7.2	Closed-loop trajectories and cost differences for the MPC of $x^+ = x + (1 + \theta)u$	170
7.3	Cost difference contours for the MPC of $x^+ = \sigma(x + (1 + \theta)u)$	186
7.4	Closed-loop trajectories and cost differences for the MPC of $x^+ = \sigma(x + (1 + \theta)u)$	187
7.5	Cost difference contours for the MPC of $x^+ = x + (1/2)\gamma(x) + (1 + \theta)u$. . .	188
7.6	Closed-loop trajectories and cost differences for the MPC of $x^+ = x + (1/2)\gamma(x) + (1 + \theta)u$	189
7.7	Simulated closed-loop trajectories for the MPC of a simple pendulum	191
	(a) Unmodeled dynamics	191
	(b) Overestimated gain	191
	(c) Underestimated gain	191
	(d) Breakdown cases	191
7.8	Open-loop dynamics, control law, terminal cost difference, optimal value function, and cost difference curves for the MPC of $x^+ = \sigma(x + (1 + \theta)u)$	210
7.9	Open-loop dynamics, control law, terminal cost difference, optimal value function, and cost difference curves for the MPC of $x^+ = x + \gamma(x)/2 + (1 + \theta)u$	215
8.1	Example systems.	256
	(a) Simple pendulum	256
	(b) Continuous stirred-tank reactor	256
8.2	Closed-loop trajectories for the offset-free MPC and tracking MPC of the simple pendulum.	262
	(a) No mismatch: $(w_p)_1 \equiv 0$ and $(w_p)_2 \equiv 0$	262
	(b) Mismatch: $(w_p)_1 \equiv 1$ and $(w_p)_2 \equiv 2$	262
	(c) Noise and mismatch: $(w_p)_3^+ = (w_p)_3 + (\Delta w_p)_3$, $(\Delta w_p)_3 \sim N(0, 10^{-2})$, and $(w_p)_4 \sim N(0, 10^{-4})$	262
	(d) Oscillating disturbance and mismatch: $(w_p)_3(k) = 1 - \cos(\frac{2\pi k}{50})$ and $r_{sp}(k) \equiv \pi$	262
8.3	Nominal steady states for the CSTR (8.65).	264
8.4	Closed-loop trajectories for the offset-free MPC and tracking MPC of the CSTR.	267
	(a) No mismatch: $(w_p)_1 \equiv 0$ and $(w_p)_2 \equiv 0$	267
	(b) Mismatch: $(w_p)_1 \equiv -0.05$ and $(w_p)_2 \equiv -0.05$	267

- (c) Noise and mismatch: $(w_p)_3^+ = (w_p)_3 + (\Delta w_p)_3$, $(\Delta w_p)_3 \sim N(0, 10^{-6})$,
and $(w_p)_4 \sim N(0, 10^{-4})$ 267
- (d) Oscillating setpoint: $r_{sp}(k) = 0.05 \sin(\frac{2\pi k}{90}) + 0.65$ 267

List of Tables

5.1	TCLab model fitting results	109
5.2	Eastman reactor model fitting results	115

Chapter 1

Introduction

Process control is the engineering discipline tasked with active regulation of industrial systems, such as chemical plants, power systems, and building energy systems. The broad goal is to profitably maintain operating conditions while satisfying safety, environmental, and quality constraints. With increasing market competition, growing environmental and safety concerns, and intensifying customer demands, the requirements on process control systems have become stricter over time. To meet these requirements, it is necessary to design process operations in a way that is rigorous and data-driven.

In Figure 1.1, the main aspects of process control are categorized based on the space and time scales on which they operate. The lowest levels are occupied by local, continuously operating control systems. The highest levels consist of large-scale (plant-wide or even enterprise-wide) problems that are infrequently solved. In chemical plants, all of these activities are represented in some form, although some industries may neglect large-scale operations.

The dashed box in Figure 1.1 surrounds operations that are most often solved in a model-based framework. Advanced process control (APC) is a catch-all term for multivariable, dynamic, model-based control. APC tracks setpoints given by the steady-state optimization layer, which optimizes the plant steady state based on a rigorous (physics-based) steady-state model. These two problems form the backbone of day-to-day operations in most chemical plants. Finally, the plant schedules (manufacturing, maintenance, distribution, etc.) are de-

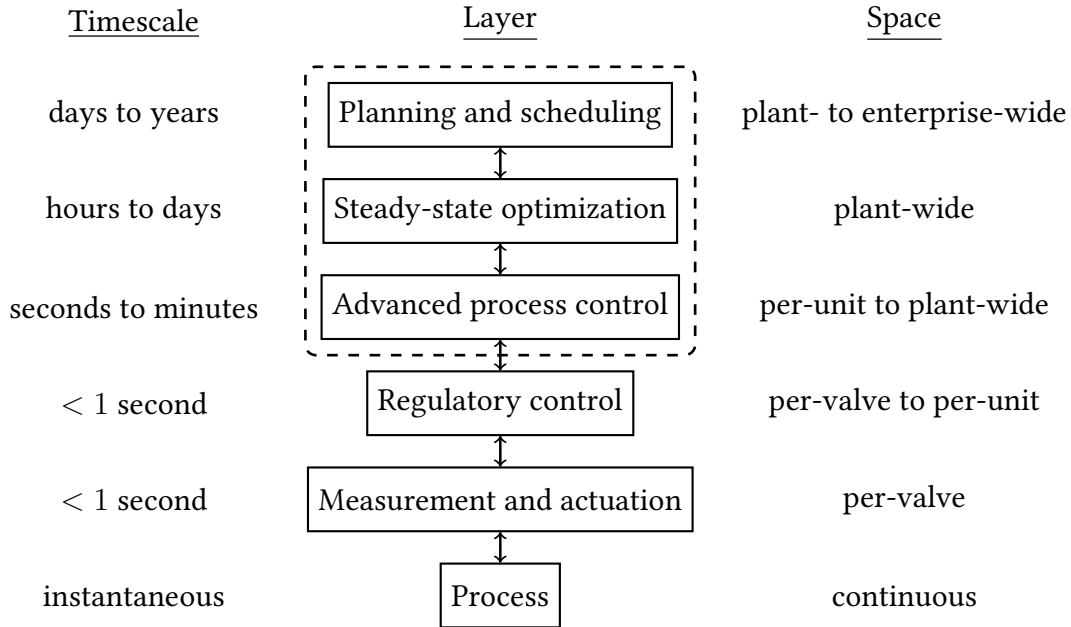


Figure 1.1: Process control heirarchy. The dashed box surrounds the control problem types most often solved in a model-based fashion. Adapted from Seborg et al. (2017).

terminated on an infrequent basis.

1.1 System identification

As many high-level process control operations are model-based, acquiring accurate process models is the paramount concern in designing these systems. Process models derived from conservation laws (mass, energy, components) and physical properties (thermodynamics, kinetics) can be represented as an ordinary differential algebraic equations (DAEs),

$$\frac{dx}{dt}(t) = F(x(t), u(t), w(t)) \quad (1.1a)$$

$$y(t) = h(x(t), u(t), w(t)) \quad (1.1b)$$

where $t \in \mathbb{R}$ is time, $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^{n_u}$ is the input (or actuator), $y \in \mathbb{R}^{n_y}$ is the output (or measurement), and $w \in \mathbb{R}^{n_w}$ is the disturbance or noise vector. However, process control is a largely digital field. Except for measurement, actuation, and a small slice of regulatory control (Figure 1.1), measurements are taken and actuators are updated at fixed discrete intervals. Therefore it is work considering the discretized version of the DAE representation (1.1),

$$x(k+1) = f(x(k), u(k), w(k))$$

$$y(k) = h(x(k), u(k), w(k))$$

where $k \in \mathbb{I}_{\geq 0}$ is the (discrete) sampled time. For ease of notation, the time index may be suppressed,

$$x^+ = f(x, u, w) \tag{1.2a}$$

$$y = h(x, u, w) \tag{1.2b}$$

where $(\cdot)^+$ denotes a forward time shift operator, i.e., $x^+ = x(k+1)$.

The control subfield concerned with producing system models from process data is called *system identification*. This encompasses not only statistical methods, but also system theoretic methods by which process models are extracted from signals (e.g., with linear algebra, Fourier transforms, or curve fitting methods). System identification methods require data to produce a process model. Sometimes an *identification experiment* is required, where the plant (1.2) is perturbed so as to generate a dataset sufficient for identification. Other times, happenstance or historical data can be used to identify the model. In either case, it is a huge advantage to be able to do *closed-loop identification*, that is, identification without ceasing any of the process control operations depicted in Figure 1.1, even the operating layer for which the model is

being identified.

For most control implementations, linear Gaussian state-space (LGSS) models suffice,

$$x^+ = Ax + Bu + w \quad (1.3a)$$

$$y = Cx + Du + v \quad (1.3b)$$

where $w \in \mathbb{R}^n$ is the process noise and $v \in \mathbb{R}^{n_y}$ is the measurement noise. Typically the noise vectors (w, v) are modeled as independent and identically distributed (through time) Gaussian random vectors,

$$\begin{bmatrix} w \\ v \end{bmatrix} \stackrel{\text{iid}}{\sim} N \left(0, \begin{bmatrix} Q_w & S_{wv} \\ S_{wv}^\top & R_v \end{bmatrix} \right). \quad (1.3c)$$

The matrices $(A, B, C, D, Q_w, S_{wv}, R_v)$ are typically black-box coefficients, although they can also be treated as given functions of an unknown vector of model parameters θ . The LGSS model (1.3) is used in a wide variety of control contexts to represent dynamics with process and measurement uncertainty. System identification of LGSS models has a longstanding history of applied use and a large body of literature on its theory (Ljung, 1999; Hannan and Deistler, 2012; Shumway and Stoffer, 2017). Other model types (transfer function, autoregressive) are common in system identification, but do not fit the state-space formalism that has been increasingly adopted by practitioners.

The phenomenon in which the plant and model are not aligned, especially in the context of a control implementation, is referred to as *plant-model mismatch* or simply *mismatch*. As the identified model is never a perfect representation of the plant, control algorithms have to be designed to handle an appropriate degree of mismatch. As some degree of mismatch is guaranteed, handling mismatch is as necessary a goal as minimizing the mismatch itself.




Figure 1.2: Setpoint tracking MPC. Adapted from Rawlings et al. (2020).

1.2 Model predictive control

For over three decades, model predictive control (MPC) has been the go-to APC method in the chemical process industries (Qin and Badgwell, 2003; Darby and Nikolaou, 2012). MPC is an advanced feedback control method in which an optimal control problem is solved on-line, based on a current state estimate (Rawlings et al., 2020). Since MPC is formulated with optimization problems, it can handle physical and safety constraints and optimize economic objectives, which are key requirements for operating a safe and profitable chemical plant.

For setpoint tracking and output feedback, three basic control blocks are combined, each solving a key problem in the state space:

- **State estimator:** *where are we?* Determine the current state estimate $\hat{x}(k)$ based on past data $(u(0), y(0), \dots, u(k-1), y(k-1))$.
- **Steady-state target problem:** *where should we go?* Determine the steady-state targets (x_s, u_s) that achieve (or get closest to) the setpoints (u_{sp}, y_{sp}) .
- **Regulator:** *how do we get there?* Find a feedback law $u = \kappa(\hat{x}, x_s, u_s)$ that drives the plant to the steady-state targets (x_s, u_s) given the current state estimate $\hat{x}(k)$.

This framework, illustrated in Figure 1.2, enables the translation of process data, constraints, and operating specifications to a feedback law. In the absence of exogenous disturbances and plant-model mismatch, and assuming the setpoints are reachable, the input and output convergence to the supplied setpoints.

1.2.1 Offset-free model predictive control

Noise, mismatch, and process upsets are ever-present realities of plant operations. To enable the MPC to track setpoints in the presence of persistent disturbances and plant-model mismatch, it is common practice, as in classic regulatory control, to remove offset by augmenting the controller with integrators. In *offset-free* MPC, the integrators take the form of uncontrollable integrating modes, called *integrating disturbances*.¹ Offset-free MPC can track setpoints even under significant plant-model mismatch and persistent disturbances, which is crucial for profitability in the modern chemical industry. For the linear case (1.3), the model is augmented as follows:

$$x^+ = Ax + B_d d + Bu + w \quad (1.4a)$$

$$d^+ = d + w_d \quad (1.4b)$$

$$y = Cx + C_d d + Du + v \quad (1.4c)$$

where $d, w_d \in \mathbb{R}^{n_d}$ are the integrating disturbances and disturbance driving noise, respectively. Again, the noise is assumed to be Gaussian,

$$\begin{bmatrix} w^\top & w_d^\top & v^\top \end{bmatrix}^\top \stackrel{\text{iid}}{\sim} N(0, S_d). \quad (1.4d)$$

The offset-free MPC is illustrated in Figure 1.3. The offset-free linear MPC is defined in Chapter 2, but we briefly outline the key modifications to the MPC below.

- **State and disturbance estimator:** Determine the current state *and disturbance* estimates $(\hat{x}(k), \hat{d}(k))$ based on past data $(u(0), y(0), \dots, u(k-1), y(k-1))$.
- **Steady-state target problem:** Determine the steady-state targets (x_s, u_s) that achieve

¹Generally speaking, *uncontrollable modes* are system eigenvalues that cannot be transferred to the origin by any input sequence. In discrete time, an integrating mode is an eigenvalue on the unit circle. Typically, the integrating disturbances are uncontrollable and have repeated, simple eigenvalues of 1.

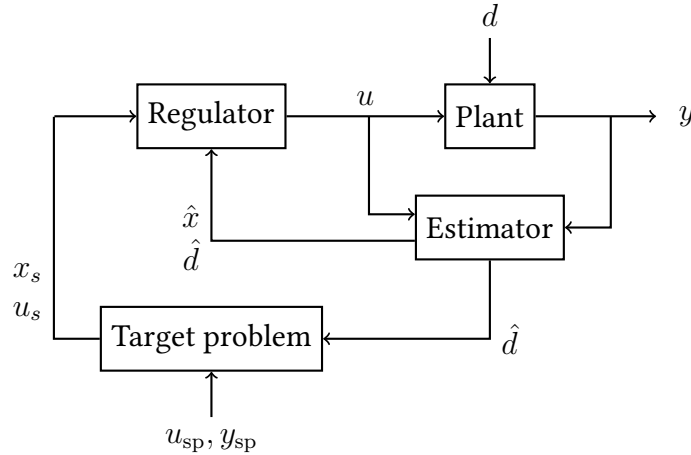


Figure 1.3: Offset-free MPC. Adapted from Rawlings et al. (2020).

(or get closest to) the setpoints $(u_{\text{sp}}, y_{\text{sp}})$, *subject to the current disturbance estimate* $\hat{d}(k)$.

- **Regulator:** Find a feedback law $u = \kappa(\hat{x}, \hat{d}, x_s, u_s)$ that drives the plant to the steady-state targets (x_s, u_s) given the current state *and disturbance* estimates $(\hat{x}(k), \hat{d}(k))$.

1.2.2 Tuning and identification

In any MPC implementation, including the linear offset-free MPC introduced in Chapter 2, a number of design parameters are introduced that must be tuned. For the regulator, there are a number of well-known tuning rules that can be used to design the *cost function weights* to achieve specific closed-loop dynamics for linear systems in the absence of estimator errors (Bryson and Ho, 1975, Chapter 5). Moreover, the model can be modified to formulate rate-of-change penalties on the inputs (Rawlings et al., 2020, Exercise 1.25). For the steady-state target problem, the *steady-state costs* should reflect the relative importance of the setpoints $(u_{\text{sp}}, y_{\text{sp}})$, can be chosen to approximate a steady-state economic optimization problem, or can simply be borrowed from the regulator problem. For the state estimator, the problem of tuning becomes more ambiguous.

While tuning is widely practiced in industry, the Kalman filter is an optimal estimator

and is fully specified by (1.4). Therefore, it is preferable to *identify* the noise model rather than *tune* the state estimator. Nonetheless, both industrial practitioners and academic control researchers hand-tune noise models or state estimators to achieve desirable performance. This can be accomplished with pole placement (Wallace et al., 2012, 2015), covariance matrix selection (Caveness and Downs, 2005; Huang et al., 2010; Petersen et al., 2017), or estimator gain selection (Deenen et al., 2018).

In addition to the tuned parameters, the quality of the LADM (1.4) as a whole determines performance of the offset-free MPC. The plant is necessarily not a member of this class of models, limiting the system identification methods and theory applicable to offset-free MPC implementation. System identification methods for the LADM (1.4) or its disturbance model include autocovariance least squares (ALS) estimation (Odelson et al., 2003), indirect or nested maximum likelihood (ML) estimation of the disturbance model (Kuntz and Rawlings, 2022; Kuntz et al., 2023), and direct ML identification of the complete model (Zagrobelny and Rawlings, 2015; Simpson et al., 2023; Kuntz and Rawlings, 2024a). All but the direct ML methods require identification of a disturbance-free model to which the integrating disturbance states are augmented. To the best of our knowledge, only Kuntz and Rawlings (2022); Kuntz et al. (2023); Kuntz and Rawlings (2024a) integrate plant and disturbance identification on real-world (not simulated) process data, although Simpson et al. (2023) includes a simulated example of LADM identification on a temperature control application.

1.3 Towards a turnkey model predictive controller

We define turnkey MPC as a MPC algorithm that can be deployed out-of-the-box, with no additional tuning required. This extends to the design of the regulator, state and disturbance estimator, target problem, and even on model upkeep, as illustrated in Figure 1.4. Such a MPC implementation has been considered the “holy grail” of industrial MPC design. While

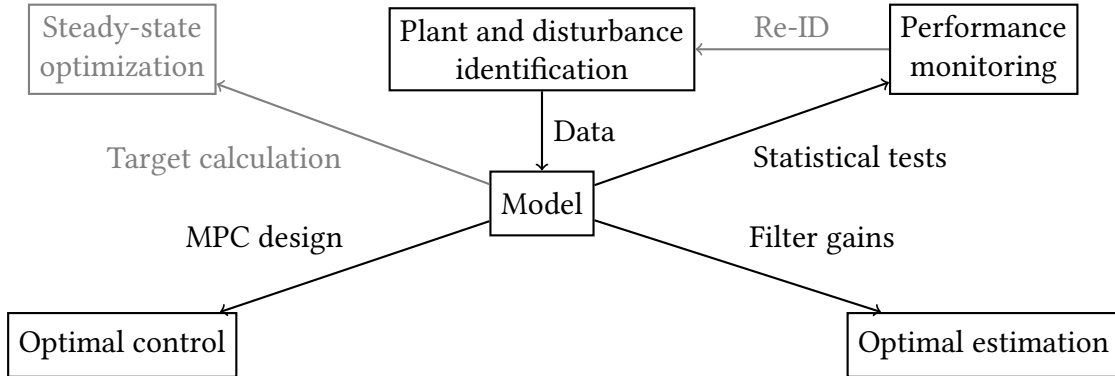


Figure 1.4: Conceptual diagram of a turnkey model predictive controller. Solid black parts of the diagram are discussed in this thesis, whereas gray parts are left for future work.

a few MPC technologies (e.g., adaptive MPC, reinforcement learning MPC) have promised to achieve a turnkey design, none have seen widespread industry adoption, if any adoption at all. While it is difficult, if not impossible, to achieve a turnkey controller with a *blind* implementation (i.e., no process knowledge or data), we can at least hope to implement such controllers when an initial seed of process data is available.

To automatically implement industrially relevant MPCs of the type discussed thus far, it is necessary to develop high-quality system identification methods for the LADM (1.4). We approach this problem in three parts. First, system identification algorithms for (1.4) are proposed. Second, the algorithms are applied to real-world data on industrially relevant control problems. Third, the control-theoretic significance of implementing MPCs with identified models is explored. The remainder of this section is devoted to outlining the thesis and describing how each chapter contributes to the overall goal of developing a turnkey MPC implementation.

Chapter 2: Offset-free MPC. Relevant background information on the offset-free MPC problem is presented: basic linear systems theory, problem formulation, design considerations, and closed-loop properties.

Part I: Identification

Chapter 3: Constrained maximum likelihood identification. A constrained ML identification approach to the identification of the LADM (1.4) is presented. The algorithm enforces constraints that are relevant to control performance and safe operation.

Chapter 4: Maximum likelihood estimator of disturbance models. A method is presented for augmenting standard identification methods with disturbance model identification. This method has closed-form solutions and is thus easily implemented. It also serves as template for generating initial guesses for the algorithm in Chapter 3.

Part II: Application

Chapter 5: Case studies in combined identification and offset-free MPC. This chapter presents a number of case studies on implementing turnkey MPCs with the identification techniques proposed in Chapters 3 and 4.

Part III: Theory

Chapter 6: Linear control of nonlinear systems. The implications of using linear identified models for control of nonlinear systems is investigated. Linear offset-free MPC of nonlinear systems is considered as a special case. These results differ from the standard offset-free theorems of Muske and Badgwell (2002); Pannocchia and Rawlings (2003); Morari and Maeder (2012) in that they directly address closed-loop stability, rather than assume it is achieved.

Chapter 7: Stability of MPC despite plant-model mismatch. The stability (or lack thereof) of MPC with plant-model mismatch is investigated. Only the special case of plant-model mismatch that does not affect the steady state is considered, with the general case considered in Chapter 8. These results differ from standard inherent robustness results De Nicolao

et al. (1996); Scokaert et al. (1997); Grimm et al. (2004); Pannocchia et al. (2011); Allan et al. (2017) in that they consider strict convergence of the state to the origin, rather than convergence to a ball around the origin, despite perturbations to the system. While the MPC is not *generally* stable despite arbitrarily small mismatch, quadratic costs and a mild differentiability requirement are enough to guarantee stability.

Chapter 8: Offset-free performance of MPC. The results of Chapter 7 are extended to offset-free MPC, showing that under similar quadratic cost designs, the same differentiability requirement, and a constraint backoff in the steady-state target problem, offset-free MPC tracks setpoints and rejects disturbances robustly in *the changes to the setpoints and disturbances*. In the absence of changes to the setpoints, and asymptotically constant disturbances and noises, offset-free performance is achieved.

Chapter 9: Conclusion. Finally, the progress towards developing a turnkey MPC is summarized. Future research goals for establishing more “hands-off” MPC implementations are discussed.

1.4 Notation and basic definitions

Sets Denote the integers, nonnegative integers, positive integers, and intervals of integers by \mathbb{I} , $\mathbb{I}_{\geq 0}$, $\mathbb{I}_{> 0}$, and $\mathbb{I}_{a:b} = \{a, a + 1, \dots, b - 1, b\}$, respectively. Denote the set of reals, nonnegative reals, positive reals, and $n \times m$ real matrices by \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\mathbb{R}_{> 0}$, and $\mathbb{R}^{n \times m}$, respectively. Real intervals are denoted by square and round brackets, e.g., $(a, b]$, where a round bracket denotes the limit *is not* included, and a square bracket denotes the limit *is* included. Denote the complex numbers, vectors, and matrices by \mathbb{C} , \mathbb{C}^n , and $\mathbb{C}^{n \times m}$, respectively. For any $z \in \mathbb{C}$,

let \bar{z} denote its complex conjugate, and define the conjugate of a vector or matrix element-wise. Let $\bar{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$ and $\bar{\mathbb{R}}_{\geq 0} := \mathbb{R}_{\geq 0} \cup \{\infty\}$ denote the extended reals and extended nonnegative reals. The complement, interior, closure, and boundary of a set S are denoted S^c , $\text{int}(S)$, $\text{cl}(S)$, and ∂S , respectively. For any function $V : X \rightarrow \bar{\mathbb{R}}$ and $\rho > 0$, we define $\text{lev}_\rho V := \{x \in X \mid V(x) \leq \rho\}$. We say $V : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}_{\geq 0}$ is lower semicontinuous (l.s.c.) if $\text{lev}_\rho V$ is closed for each $\rho \geq 0$.

Matrices and norms We denote by I_n and $0_{m \times n}$ the $n \times n$ identity matrix and $m \times n$ zero matrix, respectively. Subscripts are omitted when the dimensions are clear from context. Denote the vector and matrix transpose and Hermitian by $(\cdot)^\top$ and $(\cdot)^{\text{H}} := \overline{(\cdot)^\top}$. Denote the matrix inverse (for a square matrix) and pseudoinverse (for any matrix) by $(\cdot)^{-1}$ and $(\cdot)^\ddagger$, respectively. The trace and determinant of $A \in \mathbb{R}^{n \times n}$ are denoted $\text{tr}(A)$ and $|A|$, respectively. For any matrix $B \in \mathbb{R}^{n \times m}$, we denote by $\underline{\sigma}(B)$ and $\bar{\sigma}(B)$ the smallest and largest singular values of B , respectively.

Denote the set of $n \times n$ symmetric, positive definite, and positive semidefinite matrices by \mathbb{S}^n , \mathbb{S}_{++}^n , and \mathbb{S}_+^n . We denote the positive semidefinite square root of $Q \succeq 0$ as $Q^{1/2} \succeq 0$, where $Q = (Q^{1/2})^2$. We define the vector 2-norm and Q -weighted norm (seminorm) as $|x| := \sqrt{x^\top x}$ and $|x|_Q := \sqrt{x^\top Q x}$, respectively, for all $x \in \mathbb{R}^n$ and any $Q \succ 0$ ($Q \succeq 0$). We define the induced matrix 2-norm by $\|A\| := \max_{|x|=1} |Ax|$ and note $\|A\| = \bar{\sigma}(A)$ for all $A \in \mathbb{R}^{n \times m}$. We define the closed (open) unit n -ball by $\mathbb{B}^n := \{x \in \mathbb{R}^n \mid |x| \leq 1\}$ ($\text{int}(\mathbb{B}^n) := \{x \in \mathbb{R}^n \mid |x| < 1\}$). Denote the set of lower triangular matrices and lower triangular matrices with positive diagonal entries by \mathbb{L}^n and \mathbb{L}_{++}^n . Recall $M \in \mathbb{R}^{n \times n}$ is positive definite if and only if there exists a unique $L \in \mathbb{L}_{++}^n$, called the *Cholesky factor*, such that $M = LL^\top$. Let $\text{chol} : \mathbb{S}_{++}^n \rightarrow \mathbb{L}_{++}^n$ denote the invertible function that maps a positive definite matrix to its Cholesky factor.

Define the matrix direct sum \oplus and the Kronecker product \otimes by

$$A \otimes B := \begin{bmatrix} A_{11}B & \dots & A_{1m}B \\ \vdots & & \vdots \\ A_{n1}B & \dots & A_{nm}B \end{bmatrix}, \quad A \oplus B := \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}.$$

Define the vectorization operator $\text{vec} : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{nm}$ and symmetric vectorization operator $\text{vecs} : \mathbb{S}^n \rightarrow \mathbb{R}^{n^2}$ by

$$\begin{aligned} \text{vec}(A) &:= [A_{11} \ \dots \ A_{n1} \ A_{12} \ \dots \ A_{n2} \ \dots \ A_{1m} \ \dots \ A_{nm}]^\top, \\ \text{vecs}(M) &:= [M_{11} \ \dots \ M_{n1} \ M_{22} \ \dots \ M_{n2} \ \dots \ M_{nn}]^\top \end{aligned}$$

for each $A \in \mathbb{R}^{n \times m}$ and $M \in \mathbb{S}^n$.

Define the set of eigenvalues of a matrix $A \in \mathbb{R}^{n \times n}$ by

$$\lambda(A) := \{ \lambda \in \mathbb{C} \mid \exists v \neq 0 : Av = \lambda v \} = \{ \lambda \in \mathbb{C} \mid \det(A - \lambda I) = 0 \}.$$

The spectral radius and spectral abscissa are defined as $\rho(A) := \max_{\lambda \in \lambda(A)} |\lambda|$ and $\alpha(A) := \max_{\lambda \in \lambda(A)} \text{Re}(\lambda)$, respectively. We say a matrix A is Schur (Hurwitz) stable if $\rho(A) < 1$ ($\alpha(A) < 0$).

Signals and comparison functions For any signal $a(k)$, we denote, with slight abuse of notation, both infinite and finite sequences in bold font as $\mathbf{a} := (a(0), \dots, a(k))$ or $\mathbf{a} := (a(0), a(1), \dots)$, respectively, where length is specified or implied from context. We define the infinite and length- k signal norm as $\|\mathbf{a}\| := \sup_{k \geq 0} |a(k)|$ and $\|\mathbf{a}\|_{0:k} := \max_{0 \leq i \leq k} |a(i)|$.

Let \mathcal{PD} be the class of functions $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\alpha(0) = 0$ and $\alpha(s) > 0$ for all $s > 0$. Let \mathcal{K} be the class of \mathcal{PD} -functions that are continuous and strictly increasing. Let \mathcal{K}_∞ be the class of \mathcal{K} -functions that are unbounded. Let \mathcal{KL} be the set of functions $\beta :$

$\mathbb{R}_{\geq 0} \times \mathbb{I}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ such that $\beta(\cdot, k) \in \mathcal{K}$, $\beta(r, \cdot)$ is nonincreasing, and $\lim_{i \rightarrow \infty} \beta(r, i) = 0$ for all $(r, k) \in \mathbb{R}_{\geq 0} \times \mathbb{I}_{\geq 0}$. Denote the identity map by $\text{ID}(\cdot) := (\cdot) \in \mathcal{K}_{\infty}$.

Probability and random variables We use \sim as a shorthand for “distributed as” and $\overset{\text{iid}}{\sim}$ as a shorthand for “independent and identically distributed as.” We denote that a random vector x has a Gaussian distribution with mean μ and covariance Σ by $x \sim \text{N}(\mu, \Sigma)$. We denote the expectation by \mathbb{E} and the probability measure by $\text{Pr}[\cdot]$.

Chapter 2

Offset-free model predictive control

This chapter reviews problem formulations and properties of linear and nonlinear offset-free MPC. This chapter contains no new results (except for a few trivial generalizations of standard offset-free MPC theorems) and can safely be skipped by readers already familiar with offset-free MPC.

2.1 Linear systems

Before defining the linear MPC schemes, it is worth reviewing basic facts of linear systems. For simplicity, consider the noise-free system,

$$x^+ = Ax + Bu \tag{2.1a}$$

$$y = Cx + Du. \tag{2.1b}$$

Recall the system (2.1) (or the pair (A, B)) is *controllable* (*stabilizable*) if any state can be brought to the origin with a finite (infinite) sequence of inputs. Similarly, the system (2.1) (or the pair (A, C)) is *observable* (*detectable*) if any unknown initial state can be determined by a finite (infinite) sequence of inputs and outputs, starting from the initial time. The pair (A, B)

is controllable if and only if the following controllability matrix is full row rank,

$$\mathcal{C}_n := [B \ AB \ \dots \ A^{n-1}B], \quad (2.2)$$

and the pair (A, C) is observable if the following observability matrix is full column rank,

$$\mathcal{O}_n := \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}. \quad (2.3)$$

The system (2.1) (or the triple (A, B, C)) is *minimal* if both (A, B) is controllable and (A, C) is observable, or equivalently, the matrix $\mathcal{H}_{nn} := \mathcal{O}_n \mathcal{C}_n$ has rank n . For the system (2.1), there always exists a similarity transformation $\begin{bmatrix} x_c \\ x_{uc} \end{bmatrix} := T_c x$ such that

$$\begin{bmatrix} x_c^+ \\ x_{uc}^+ \end{bmatrix} = \begin{bmatrix} A_c & A_{12} \\ 0 & A_{uc} \end{bmatrix} \begin{bmatrix} x_c \\ x_{uc} \end{bmatrix} + \begin{bmatrix} B_c \\ 0 \end{bmatrix} u, \quad y = [C_c \ C_{uc}] \begin{bmatrix} x_c \\ x_{uc} \end{bmatrix} + Du \quad (2.4)$$

where (A_c, B_c) is controllable (or empty). We say the system (2.4) is in *controllability canonical form*. We say (A, B) is *stabilizable* if A_{uc} is Schur stable. Similarly, there always exists (different) a similarity transformation $\begin{bmatrix} x_o \\ x_{uo} \end{bmatrix} := T x$ such that

$$\begin{bmatrix} x_o^+ \\ x_{uo}^+ \end{bmatrix} = \begin{bmatrix} A_o & 0 \\ A_{21} & A_{uo} \end{bmatrix} \begin{bmatrix} x_o \\ x_{uo} \end{bmatrix} + \begin{bmatrix} B_o \\ B_{uo} \end{bmatrix} u, \quad y = [C_o \ 0] \begin{bmatrix} x_o \\ x_{uo} \end{bmatrix} \quad (2.5)$$

where (A_o, C_o) is observable (or empty). We say the system (2.5) is in *observability canonical form*. We say (A, C) is *detectable* if A_{uo} is Schur stable. Next, we have the following generalization of the classic *Popov-Belevitch-Hautus* test.

Lemma 2.1 ((Hespanha, 2018, Thms. 12.3, 14.3)). *The pair (A, B) is controllable (stabilizable) if and only if*

$$\text{rank} \left(\begin{bmatrix} A - \lambda I & B \end{bmatrix} \right) = n \quad (2.6)$$

for all $\lambda \in \mathbb{C}$ ($|\lambda| \geq 1$).

Finally, we note these properties are, in some sense, *duals* of each other, in that (A, B) is controllable (stabilizable) if and only if (A^\top, B^\top) is observable (detectable). Thus Lemma 2.1 can be slightly modified to characterize observable and detectable pairs.

When referring to controllability or observability properties of the LGSS model (1.3) and LADM (1.4), assume we are referring to the noise-free versions unless otherwise stated. Notice that the LADM (1.4) is in controllability canonical form if (A, B) is controllable. In any case, the integrating disturbances are uncontrollable and have eigenvalues on the unit circle, so they are uncontrollable integrating modes.

2.2 Linear model predictive control

The setpoint tracking MPC outlined in Chapter 1 combines three distinct problems that are solved at each time step: estimation, target calculation, and regulation. The goal is firstly to remove offset from a setpoint $r_{\text{sp}} \in \mathbb{R}^{n_r}$ in the *controlled variables* $r = H_y y + H_u u \in \mathbb{R}^{n_r}$ and secondly to minimize the distance from a pair of input-output setpoints $(u_{\text{sp}}, y_{\text{sp}}) \in \mathbb{R}^{n_u + n_y}$. While only offset-free MPC will be used in the case studies of Chapter 5, we summarize the standard MPC described in (Rawlings et al., 2020, Chapter 1) to highlight the key differences between standard and offset-free MPC.

State estimator The goal of the estimator is to determine a state estimate $\hat{x}(k)$ from past data $(u(0), y(0), \dots, u(k-1), y(k-1))$. Typically a linear filter suffices,

$$\hat{x}^+ = A\hat{x} + Bu + K(y - C\hat{x} - Du) \quad (2.7)$$

where $K \in \mathbb{R}^{n_x \times n_y}$ is the filter gain. To determine the filter gain from the LGSS model (1.3), one can solve the *discrete algebraic Riccati equation* (DARE),

$$P = APA^\top + Q_w - (APC^\top + S_{wv})(CPC^\top + R_v)^{-1}(APC^\top + S_{wv})^\top \quad (2.8)$$

and implement the famous *Kalman filter* gain $K := (APC^\top + S_{wv})(CPC^\top + R_v)^{-1}$. We refer the reader to (Kwakernaak and Sivan, 1972; Hespanha, 2018) for a classical treatment of the linear optimal estimation problem and to (Rawlings et al., 2020, pp. 27–46) for a derivation of the optimal filter gain K from least squares theory. Stability of the filter (2.7) is equivalent to stability of the matrix $A - KC$, and is discussed further in Chapter 3. Since the Kalman filter gain K is fully specified by the parameters in the LGSS model (1.3), identification of (1.3) automatically gives the state estimator.

Remark 2.2. In most literature, the direct feedthrough term is not included (i.e., $D = 0$), and the state estimator can be split into two steps: prediction,

$$(\hat{x}^*)^+ = A\hat{x} + Bu + K_p(y - C\hat{x}) \quad (2.9)$$

and filtering,

$$\hat{x} = \hat{x}^* + K_f(y - C\hat{x}^*) \quad (2.10)$$

where \hat{x} and \hat{x}^* are the filtered and predicted estimates, and for the Kalman filter, $K_p := (APA^\top + S_{wv})(CPC^\top + R_v)^{-1}$ and $K_f := PC^\top(CPC^\top + R_v)^{-1}$. This has the advantage of providing feedback from the *current* output to the current input, which could not be accomplished when the output was an explicit function of the input.

Steady-state target problem The goal of the steady-state target problem is find targets (\bar{x}_s, \bar{u}_s) that track a setpoint $r_{\text{sp}} \in \mathbb{R}^{n_r}$ in the controlled variables,

$$r_{\text{sp}} = H_y \bar{y}_s + H_u \bar{u}_s = H_y C \bar{x}_s + (H_y D + H_u) \bar{u}_s \quad (2.11)$$

while maintaining steady state,

$$\bar{x}_s = A \bar{x}_s + B \bar{u}_s \quad (2.12)$$

and satisfying the linear constraints,

$$E \bar{y}_s + F \bar{u}_s = EC \bar{x}_s + (ED + F) \bar{u}_s \leq f. \quad (2.13)$$

This can be accomplished in the following optimization problem

$$\min_{\bar{x}_s, \bar{u}_s} \frac{1}{2} |C \bar{x}_s + D \bar{u}_s - y_{\text{sp}}|_{Q_s}^2 + \frac{1}{2} |\bar{u}_s - u_{\text{sp}}|_{R_s}^2 \quad \text{subject to} \quad (2.11)-(2.13) \quad (2.14)$$

where $(u_{\text{sp}}, y_{\text{sp}})$ are auxilliary setpoints and (Q_s, R_s) are positive definite weighting matrices.

We denote the solutions to (2.14) by (x_s, u_s) .

In practice, the setpoint (2.11) may not be *reachable*, meaning the target problem (2.14) may not have solutions. In these cases, it may suffice to implement (2.11) as a soft constraint, allowing (2.14) to choose the closest reachable setpoint.

Regulator The goal of the regulator is to produce a control law $u = \kappa(\hat{x}, x_s, u_s)$ that drives the system to the steady-state targets (x_s, u_s) given a state estimate \hat{x} . This can be accom-

plished in the following finite horizon optimal control problem

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{u}}} \frac{1}{2} \sum_{i=0}^{N-1} |\tilde{x}_i - x_s|_Q^2 + |\tilde{u}_i - u_s|_R^2 + \frac{1}{2} |\tilde{x}_N - x_s|_{P_f}^2 \quad (2.15a)$$

$$\text{subject to } \tilde{x}_0 = \hat{x}(k), \quad (2.15b)$$

$$\tilde{x}_{i+1} = A\tilde{x}_i + B\tilde{u}_i \quad \forall i \in \mathbb{I}_{0:N-1}, \quad (2.15c)$$

$$EC\tilde{x}_i + (ED + F)\tilde{u}_i \leq f \quad \forall i \in \mathbb{I}_{0:N-1} \quad (2.15d)$$

where $y_s := Cx_s + Du_s$ and $\tilde{\mathbf{x}} := (\tilde{x}_0, \dots, \tilde{x}_N)$ and $\tilde{\mathbf{u}} := (\tilde{u}_0, \dots, \tilde{u}_{N-1})$ are state and input sequences, and (Q, R, P_f) are positive definite weighting matrices. Let $x_i^0(\hat{x}, x_s, u_s)$ and $u_i^0(\hat{x}, x_s, u_s)$ denote solutions to (2.15) as a function of the current state estimate \hat{x} and steady-state targets (x_s, u_s) . The control law is defined as the first input of the optimal input sequence $u = \kappa(\hat{x}, x_s, u_s) := u_0^0(\hat{x}, x_s, u_s)$.

Remark 2.3. If the filter-predictor equations (2.9) and (2.10) of Remark 2.2 are used, and there is cross-covariance (i.e., $S_{wv} \neq 0$), then the regulation problem should account for a nonzero filtered process noise (Jørgensen et al., 2011). Specifically, the first step of (2.20c) should be modified to $\tilde{x}_1 = A\tilde{x}_0 + B\tilde{u}_0 + \hat{w}(k)$ where $\hat{w} = S_{wv}(CPC^\top + R_v)^{-1}(y - C\hat{x}^*)$. Note this correction only applies at the current time step, and future predictions of the noise are zero.

Constraints For the SSTP and regulator we have used general linear constraints on the inputs and outputs $Ey + Fu \leq f$, but in practice *box constraints* are easiest to implement,

$$\underline{y} \leq y \leq \bar{y}, \quad \underline{u} \leq u \leq \bar{u} \quad (2.16)$$

where (\underline{u}, \bar{u}) are the input bounds, (\underline{y}, \bar{y}) are the output bounds. Box constraints can be put in the form $Ey + Fu \leq f$ by defining

$$E := \begin{bmatrix} I \\ -I \\ 0 \\ 0 \end{bmatrix}, \quad F := \begin{bmatrix} 0 \\ 0 \\ I \\ -I \end{bmatrix}, \quad f := \begin{bmatrix} \bar{y} \\ -\underline{y} \\ \bar{u} \\ -\underline{u} \end{bmatrix}. \quad (2.17)$$

For all problems discussed herein, constraints on outputs should be implemented as soft constraints in the optimizer so as to preserve feasibility of the control problem at all times.

2.3 Linear offset-free model predictive control

For the LADM (1.4), the linear MPC formulation from Section 2.2 is generalized as follows.

State and disturbance estimator The filter (2.7) is trivially generalized by lumping the plant states x and integrating disturbances d into a single state vector $x_{\text{aug}} := \begin{bmatrix} x^\top & d^\top \end{bmatrix}^\top$, producing the filter

$$\begin{bmatrix} \hat{x}^+ \\ \hat{d}^+ \end{bmatrix} = \begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{x} \\ \hat{d} \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} u + \begin{bmatrix} K_x \\ K_d \end{bmatrix} \left(y - [C \ C_d] \begin{bmatrix} \hat{x} \\ \hat{d} \end{bmatrix} - Du \right). \quad (2.18)$$

As for the general LGSS model (1.3), the the Kalman filter is the optimal state estimator for LADM (1.4), and can be found by solving a DARE in the system matrices corresponding to the augmented state. Moreover, (2.18) can be split into prediction and filtering steps as outlined in Remark 2.2.

Steady-state target problem Since the integrating disturbances are *uncontrollable*, one cannot expect to choose a disturbance target. Indeed, since they are integrating, one cannot hope to forecast their evolution without taking into account future data. Instead, the distur-

bances are assumed to be constant, and the steady-state targets (x_s, u_s) are computed as a function of the current disturbance estimate $\hat{d}(k)$,

$$\min_{\bar{x}_s, \bar{u}_s} \frac{1}{2} |\bar{u}_s - u_{\text{sp}}|_{R_s}^2 + \frac{1}{2} |C\bar{x}_s + C_d \hat{d}_s + D\bar{u}_s - y_{\text{sp}}|_{Q_s}^2 \quad (2.19a)$$

$$\text{subject to } \bar{x}_s = A\bar{x}_s + B_d \hat{d}(k) + B\bar{u}_s \quad (2.19b)$$

$$r_{\text{sp}} = H_y C \bar{x}_s + H_y C_d \hat{d}(k) + (H_y D + H_u) \bar{u}_s \quad (2.19c)$$

$$EC\bar{x}_s + EC_d \hat{d}(k) + (ED + F) \bar{u}_s \leq f. \quad (2.19d)$$

Denote the solutions to this problem by $(x_s(k), u_s(k)) = (x_s(\hat{d}(k)), u_s(\hat{d}(k)))$.

Regulator As in the steady-state target problem, the regulator is solved under the assumption the disturbances are constant,

$$\min_{\tilde{x}, \tilde{u}} \frac{1}{2} \sum_{i=0}^{N-1} |\tilde{x}_i - x_s(k)|_Q^2 + |\tilde{u}_i - u_s(k)|_R^2 + \frac{1}{2} |\tilde{x}_N - x_s(k)|_{P_f}^2 \quad (2.20a)$$

$$\text{subject to } \tilde{x}_0 = \hat{x}(k), \quad (2.20b)$$

$$\tilde{x}_{i+1} = A\tilde{x}_i + B_d \hat{d}(k) + B\tilde{u}_i \quad \forall i \in \mathbb{I}_{0:N-1}, \quad (2.20c)$$

$$EC\tilde{x}_i + EC_d \hat{d}(k) + (ED + F)\tilde{u}_i \leq f \quad \forall i \in \mathbb{I}_{0:N-1}. \quad (2.20d)$$

where (Q, R, P_f) are positive definite weighting matrices. In the case studies (Chapter 5), an infinite-horizon, rate-of-change-penalized, and box-constrained variant of (2.20) is sometimes

solved:

$$\min_{\substack{\tilde{x}_0, \tilde{x}_1, \dots \\ \tilde{u}_0, \tilde{u}_1, \dots}} \frac{1}{2} \sum_{i=0}^{\infty} |C\tilde{x}_i + C_d\hat{d}(k) + D\tilde{u}_i - y_s(k)|_{Q_y}^2 + |\tilde{u}_i - u_s(k)|_R^2 + |\tilde{u}_i - \tilde{u}_{i-1}|_M^2 \quad (2.21a)$$

$$\text{s.t. } \tilde{x}_0 = \hat{x}(k), \quad (2.21b)$$

$$\tilde{x}_{i+1} = A\tilde{x}_i + B_d\hat{d}(k) + B\tilde{u}_i \quad \forall i \in \mathbb{I}_{\geq 0}, \quad (2.21c)$$

$$\underline{y} \leq C\tilde{x}_i + C_d\hat{d}(k) + D\tilde{u}_i \leq \bar{y} \quad \forall i \in \mathbb{I}_{\geq 0}, \quad (2.21d)$$

$$\underline{u} \leq \tilde{u}_i \leq \bar{u} \quad \forall i \in \mathbb{I}_{\geq 0} \quad (2.21e)$$

where $y_s(k) := Cx_s(k) + C_d\hat{d}(k) + Du_s(k)$, and the weighting matrices Q_y and M must be positive definite and positive semidefinite. In practice, the infinite horizon problem (2.21) is solved as a finite horizon problem, where the horizon length is taken sufficiently large to approximate the infinite horizon controller.

For either problem, let $x_i^0(\hat{x}, \hat{d}, x_s, u_s)$ and $u_i^0(\hat{x}, \hat{d}, x_s, u_s)$ denote the solutions as a function of the current state estimate \hat{x} , disturbance estimate \hat{d} , and steady-state targets (x_s, u_s) . The control law is defined as $u = \kappa(\hat{x}, \hat{d}, x_s, u_s) := u_0^0(\hat{x}, \hat{d}, x_s, u_s)$. In the absence of output constraints, either regulation problem, (2.20) or (2.21), is agnostic to the current disturbance estimate, and the control law is only a function of $(\hat{x}(k) - x_s(k), u_s(k))$. In fact, if neither the inputs nor outputs are constrained, the unconstrained controller has a control law of the form $u = \kappa(\hat{x}, \hat{d}, x_s, u_s) = \kappa(\hat{x} - x_s) + u_s$. For the unconstrained case, the effect of the disturbance enters the controller through the steady-state targets. Finally, we note if the filter-predictor equations as used (c.f. Remark 2.2), the first constraint of (2.20c) and (2.21c) should be modified to $\tilde{x}_1 = A\tilde{x}_0 + B_d\hat{d}(k) + B\tilde{u}_0 + \hat{w}(k)$ where $\hat{w} = S_{wv}(CPC^\top + R_v)^{-1}(y - C\hat{x}^* - C_d\hat{d}^*)$.

2.3.1 Offset-free sufficient conditions

Muske and Badgwell (2002) first established sufficient conditions under which a linear

offset-free MPC (e.g., (2.18)–(2.20)) with a separable disturbance model, i.e.,

$$B_d = [\overline{B}_d \ 0], \quad C_d = [0 \ \overline{C}_d]$$

applied to a linear plant converges to the controlled variable setpoints r_{sp} . This was generalized to LADMs form (1.4) by Pannocchia and Rawlings (2003). Finally, Morari and Maeder (2012) generalized the conditions to nonlinear plants and models. We restate the offset-free conditions for linear models in the following theorem.

Theorem 2.4 (Pannocchia and Rawlings (2003)). *Consider a system controlled by the offset-free MPC (2.18)–(2.20) with a constant setpoint r_{sp} . Assume that*

- (i) *the SSTP and regulator are feasible at all times,*
- (ii) *the disturbance state is of the same dimension as the measurement ($n_d = p$), and*
- (iii) *the LADM (1.4) is detectable.*

If the closed-loop system is stable and the constraints are not active at steady state, then there is zero offset in the controlled variables at steady state, i.e., $\lim_{k \rightarrow \infty} r(k) = r_{\text{sp}}$.

Remark 2.5. The contrapositive of Theorem 2.4 is significant as well. If we assume hypotheses (i)–(iii) of Theorem 2.4 hold, then if there is offset in the controlled variables, we have either hit a constraint or the system is unstable.

Remark 2.6. Despite the fact that Theorem 2.4 does not explicitly mention control of nonlinear plants, the results are widely applicable to both linear and nonlinear plants with asymptotically constant disturbances. This is because Theorem 2.4 does not establish sufficient conditions for controller stability, but simply states sufficient conditions for which a stable controller also has zero offset. In fact, Pannocchia and Rawlings (2003) demonstrate the validity of Theorem 2.4 in the control of a highly nonlinear, non-isothermal reactor model.

2.3.2 Disturbance model equivalence and design

One can infer from Theorem 2.4 that, to achieve offset-free control with offset-free MPC, it is important to have a detectable model. To this end, we have the following result.

Lemma 2.7 (Pannocchia and Rawlings (2003)). *The LADM (1.4) is detectable if and only if the standard LGSS model (1.3) is detectable and*

$$\text{rank} \begin{bmatrix} A - I_n & B_d \\ C & C_d \end{bmatrix} = n + n_d \quad (2.22)$$

The so-called *offset-free rank condition* (2.22) is important in formulating disturbance models for the offset-free MPC algorithm. One can replace the third condition of Theorem 2.4 with the rank condition (2.22). It turns out that, in the same way that a state-space realization is only unique up to a similarity transformation, any *detectable* disturbance model is only unique up to a similarity transformation. In fact, the Kalman filter behavior is equivalent under this similarity transformation, so if disturbances are “misassigned” in the model there is no effect on the closed-loop system.

Lemma 2.8 (Rajamani et al. (2009)). *Consider the alternate LADM*

$$x^+ = Ax + Bu + \tilde{B}_d \tilde{d} + w \quad (2.23a)$$

$$\tilde{d}^+ = \tilde{d} + \tilde{w}_d \quad (2.23b)$$

$$y = Cx + \tilde{C}_d \tilde{d} + v \quad (2.23c)$$

$$\begin{bmatrix} w \\ \tilde{w}_d \\ v \end{bmatrix} \stackrel{\text{iid}}{\sim} \text{N}(0, \tilde{S}_d). \quad (2.23d)$$

If the LGSS model (1.3) is detectable, then the LADMs (1.4) and (2.23) are detectable if and only if both satisfy the offset-free rank condition (2.22). Moreover, there exists a choice of \tilde{S}_d such that (1.4) and (2.23) have equivalent Kalman filter innovations.

The consequence of Lemma 2.8 is that, given a standard LGSS model (1.3), one can “design” the disturbance model to be maximally interpretable, so long as it satisfies the rank condition (2.22). To this end, we propose the follow general guidelines for choosing the disturbance model:

- If \hat{A} does not contain integrators, use an output disturbance model.
- If \hat{A} contains integrators and $n_u = n_y$, use an input disturbance model, $(B_d, C_d) = (B, 0)$.
- Otherwise, use some combination of input and output disturbances, i.e. $(B_d, C_d) = (B\tilde{I}_1, \tilde{I}_2)$ where \tilde{I}_1 and \tilde{I}_2 are diagonal matrices with zeros and ones on the diagonal and collectively p nonzero elements.

Models in these forms retain interpretability while ensuring that the offset-free rank condition (2.22) is satisfied.

2.4 Nonlinear offset-free model predictive control

Morari and Maeder (2012) first extended the offset-free MPC results of Section 2.3 to designs with *nonlinear models* of the following form:

$$x^+ = \hat{f}(x, u, d) \tag{2.24a}$$

$$d^+ = d \tag{2.24b}$$

$$y = \hat{h}(x, u, d). \tag{2.24c}$$

Pannocchia et al. (2015) later summarized and extended these results to consider the SSTP explicitly and consider special cases guaranteeing estimator convergence. These authors consider nonlinear versions of the linear offset-free MPC (2.7), (2.19), and (2.20). Specifically, they

consider nonlinear observers of the form,

$$\hat{x}^+ = \hat{f}(\hat{x}, u, \hat{d}) + \kappa_x(y - \hat{h}(\hat{x}, u, \hat{d})) \quad (2.25a)$$

$$\hat{d}^+ = \hat{d} + \kappa_d(y - \hat{h}(\hat{x}, u, \hat{d})) \quad (2.25b)$$

nonlinear SSTPs of the form,

$$\min_{\bar{x}_s, \bar{u}_s} \ell_s(\bar{u}_s - u_{\text{sp}}, \hat{h}(\bar{x}_s, \bar{u}_s, \hat{d}(k)) - y_{\text{sp}}) \quad (2.26a)$$

$$\text{subject to } \bar{x}_s = \hat{f}(\bar{x}_s, \bar{u}_s, \hat{d}(k)) \quad (2.26b)$$

$$r_{\text{sp}}(k) = g(\bar{u}_s, \hat{h}(\bar{x}_s, \bar{u}_s, \hat{d}(k))) \quad (2.26c)$$

$$(\bar{u}_s, \hat{h}(\bar{x}_s, \bar{u}_s, \hat{d}(k))) \in \mathbb{Z} \quad (2.26d)$$

and nonlinear regulators of the form,

$$\min_{\tilde{x}, \tilde{u}} \frac{1}{2} \sum_{i=0}^{N-1} \ell(\tilde{x}_i - x_s(k), \tilde{u}_i - u_s(k)) + V_f(\tilde{x}_N - x_s(k)) \quad (2.27a)$$

$$\text{subject to } \tilde{x}_0 = \hat{x}(k), \quad (2.27b)$$

$$\tilde{x}_{i+1} = \hat{f}(\tilde{x}_i, \tilde{u}_i, \hat{d}(k)) \quad \forall i \in \mathbb{I}_{0:N-1}, \quad (2.27c)$$

$$(\tilde{u}_i, \hat{h}(\tilde{x}_i, \tilde{u}_i, \hat{d}(k))) \in \mathbb{Z} \quad \forall i \in \mathbb{I}_{0:N-1}. \quad (2.27d)$$

where κ_x, κ_d are the observer feedback laws, g is the (continuous) controlled variable function, \mathbb{Z} is a constraint set, and ℓ_s, ℓ, V_f are positive definite¹ steady-state, stage, and terminal cost functions. As before, solutions to (2.26) are denoted (x_s, u_s) , solutions to (2.27) are denoted $x_i^0(\hat{x}, \hat{d}, x_s, u_s)$ and $u_i^0(\hat{x}, \hat{d}, x_s, u_s)$, and the control law is defined as $u = \kappa(\hat{x}, \hat{d}, x_s, u_s) := u_0^0(\hat{x}, \hat{d}, x_s, u_s)$. For the nonlinear offset-free MPC (2.25)–(2.27), we have the following offset-

¹A function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is positive definite if $V(x) = 0 \Leftrightarrow x = 0$.

free sufficiency theorem, which is a trivial modification of (Pannocchia et al., 2015, Thm. 14).

Theorem 2.9. *Consider the plant (1.2) controlled by the nonlinear offset-free MPC (2.25)–(2.27).*

Assume that

(i) *the plant disturbance and reference are asymptotically constant, i.e., $w(k) \rightarrow w_\infty$ and*

$$r_{\text{sp}}(k) \rightarrow r_\infty,$$

(ii) *the observer feedback laws satisfy $\kappa_d(e) = 0 \Leftrightarrow e = 0$,*

(iii) *the SSTP (2.26) and regulator (2.27) are feasible at all times.*

If the closed-loop system is stable and the constraints are not active at steady state, then there is zero offset in the controlled variables at steady state, i.e., $\lim_{k \rightarrow \infty} r(k) = r_\infty$.

Proof. If the closed-loop system reaches a steady state, the state and disturbance estimates must reach a steady state. Denote the steady-state estimates by $(\hat{x}_\infty, \hat{d}_\infty)$ and let $\hat{y}_\infty := \hat{h}(\hat{x}_\infty, u_\infty, \hat{d}_\infty)$. By (2.25), we have $\hat{d}_\infty = \hat{d}_\infty + \kappa_d(y_\infty - \hat{y}_\infty)$ and therefore $\kappa_d(y_\infty - \hat{y}_\infty) = 0$. But by hypothesis (ii), this is equivalent to $y_\infty = \hat{y}_\infty$. Denote the solution to the SSTP (2.26) with $(r_{\text{sp}}(k), \hat{d}(k)) = (r_{\text{sp},\infty}, \hat{d}_\infty)$ by $(x_{s,\infty}, u_{s,\infty})$ and let $y_{s,\infty} := \hat{h}(x_{s,\infty}, u_{s,\infty}, \hat{d}_\infty)$. Since the closed-loop system is at steady state, we have $u_\infty = u_0^0(\hat{x}_\infty, \hat{d}_\infty, x_{s,\infty}, u_{s,\infty})$. Moreover, $\hat{x}_\infty = x_0^0(\hat{x}_\infty, \hat{d}_\infty, x_{s,\infty}, u_{s,\infty})$ due to the constraint (2.27b). Since ℓ is positive definite, we have $(\hat{x}_\infty, u_\infty) = (x_{s,\infty}, u_{s,\infty})$ and therefore $y_\infty = \hat{y}_\infty = \hat{h}(\hat{x}_\infty, u_\infty, \hat{d}_\infty) = \hat{h}(x_{s,\infty}, u_{s,\infty}, \hat{d}_\infty) =: y_{s,\infty}$. Finally, by the constraint (2.26b) and continuity of g , we have

$$\lim_{k \rightarrow \infty} r(k) = \lim_{k \rightarrow \infty} g(u(k), y(k)) = g(u_\infty, y_\infty) = g(u_{s,\infty}, y_{s,\infty}) = r_\infty. \quad \square$$

Remark 2.10. Pannocchia et al. (2015) propose a slightly different observer than (2.25). They do not consider direct feedthrough in the nonlinear plant and model, i.e., $h(x, u, w) = h(x, w)$

and $\hat{h}(x, u, d) = \hat{h}(x, d)$. As in Remark 2.2, the observer is split into a *prediction* step

$$(\hat{x}^*)^+ = \hat{f}(\hat{x}^*, u, \hat{d}^*) \quad (2.28)$$

$$(\hat{d}^*)^+ = \hat{d}^* \quad (2.29)$$

and a *filtering* step

$$\hat{x} = \hat{x}^* + \kappa_x(y - \hat{h}(\hat{x}^*, \hat{d}^*)) \quad (2.30)$$

$$\hat{d} = \hat{d}^* + \kappa_d(y - \hat{h}(\hat{x}^*, \hat{d}^*)) \quad (2.31)$$

which, as in Remark 2.2, has the advantage of providing feedback from the current output. As the nonlinear observer has no statistical relevance, there are no cross-correlation adjustments to be made. Theorem 2.9 holds for either observer design with trivial modifications to the proof.

For the special case of state feedback, Pannocchia et al. (2015) give a disturbance model and estimator design are given for which the offset-free MPC is provably asymptotically stable and offset-free. Pannocchia et al. (2015) also generalize this observer design to *economic* cost functions² and demonstrate convergence to the optimal steady state. A general, output-feedback offset-free *economic* MPC was first proposed by Vaccari and Pannocchia (2017), who use a gradient correction strategy called *modifier adaptation* to ensure the economic MPC, if it converges, achieves the optimal steady-state performance. For further developments of modifier-adaptation for offset-free economic MPC, we refer the reader to Pannocchia (2018); Faulwasser and Pannocchia (2019); Vaccari et al. (2021).

To the best of our knowledge, there are no stability results for linear, let alone nonlinear, offset-free MPC for the intended setting: persistent disturbances and plant-model mismatch.

²By economic cost functions, we simply mean costs that are not necessarily positive semidefinite, although they are usually a quantification of the net operating cost (or negative profit).

Some authors have proposed tracking MPC designs with offset-free behavior (Limon et al., 2008; Betti et al., 2013; Falugi and Mayne, 2013; Falugi, 2015; Limon et al., 2018; Köhler et al., 2020; Berberich et al., 2022a; Galuppini et al., 2023; Soloperto et al., 2023), but they all assume access to plant dynamic equations (1.2). Part III of this thesis will cover offset-free performance and asymptotic stability of nonlinear offset-free designs, similar to (2.25)–(2.27), with positive definite quadratic costs.

Part I

Identification

Chapter 3

Constrained maximum likelihood identification for offset-free control

As mentioned in Chapter 1, model quality is the main contributor to the performance of industrial MPC implementations (Canney, 2003; Darby and Nikolaou, 2012). For offset-free MPC, integral action is provided through the disturbance estimates, so the control performance depends on the estimator dynamics. Tuning of integrating disturbance models can be a time-consuming and ad-hoc procedure, requiring simplified parameterizations (e.g., diagonal covariance matrices). In prior work, we have suggested identification as the preferred strategy for acquiring LADMs (Kuntz and Rawlings, 2022; Kuntz et al., 2023). In this chapter, we further develop ML identification because of its desirable statistical properties (consistency, asymptotic efficiency) and ability to handle general model structures and constraints (Åström, 1979; Ljung, 1999). We remark that other identification methods (subspace identification, autoregressive modeling, etc.) are not suitable for LADM identification as they cannot impose the model structure we require.

Design constraints can be included in tuning procedures to avoid undesirable filter behaviors (slow response time, fictitious high frequencies) that are passed to the control performance through the integrating disturbance estimates. Control-relevant design constraints and prior knowledge have sometimes been incorporated into identification problems (Piga

et al., 2019; Formentin and Chiuso, 2021; Berberich et al., 2023). However, there are no general approaches to shaping the closed-loop filter behavior in ML identification. To address this gap, we consider ML identification with eigenvalue constraints implemented via the LMI regions commonly used in robust control (Chilali and Gahinet, 1996; Chilali et al., 1999).

LMI region constraints have been used in subspace identification (Miller and De Callafon, 2013). However, subspace identification cannot be used for LADM identification as it is not possible to impose the required disturbance model structure. Open-loop stability constraints have been included in the expectation maximization (EM) algorithm (Umenberger et al., 2018), but this formulation is not obviously generalized to filter stability or general LMI region constraints.

While EM is an algorithm for ML, it does not have strong convergence guarantees. While it can be shown that the EM iterates produce, almost surely, an increasing sequence of likelihood values (Shumway and Stoffer, 1982; Gibson and Ninness, 2005), slow convergence at low noise levels has been reported on a range of problems (Umenberger et al., 2018; Redner and Walker, 1984; Bermond and Cardoso, 1999; Petersen et al., 2005; Petersen and Winther, 2005; Olsson et al., 2007). Interior point, and even gradient methods (Olsson et al., 2007), are therefore preferable to the EM approach.

As originally posed by Chilali and Gahinet (1996); Chilali et al. (1999), LMI regions are strict semidefinite matrix inequalities. While Miller and De Callafon (2013) used relaxed LMI regions with nonstrict inequalities, as we show in Section 3.4, the constraint sets are not closed, and thus problematic as optimization constraints. To address this issue, we formulate tightened LMI region constraints that define a closed constraint set. This formulation introduces nonlinear matrix inequalities and semidefinite matrix arguments, making the ML problem a nonlinear semidefinite program (NSDP).

To efficiently convert the NSDP to a nonlinear program (NLP), we generalize the Burer-Monteiro-Zhang (BMZ) method (Burer et al., 2002a,b), which was originally used to convert

sparse semidefinite matrix arguments into vector arguments with minimal dimension. An additional advantage of the BMZ method over standard Cholesky factor substitution is that structural knowledge of the plant design (e.g., flowsheet or network structure) can be imposed in the model parameterization in an efficient manner. Finally, while this work is primarily motivated by identification of LADMs and offset-free MPC implementations, we remark that any linear Gaussian state-space model can be identified, with eigenvalue constraints, using this approach.

3.1 Maximum likelihood estimation

Maximum likelihood (ML) estimation is a popular statistical method for parametric modeling. In general, the ML estimation problem is to find parameters that maximize the *conditional density* of the following model:

$$\mathbf{y}_{N-1} | \mathbf{u}_{N-1} \sim p_N(\mathbf{y}_{N-1} | \mathbf{u}_{N-1}, \theta) \quad (3.1)$$

where $\mathbf{y}_{N-1} := (y(0), \dots, y(N-1)) \in \mathbb{R}^{pN}$ is the output (or measured) data, $\mathbf{u}_{N-1} := (u(0), \dots, u(N-1)) \in \mathbb{R}^{mN}$ is the input (or actuator) data, $\theta \in \mathbb{R}^{n_\theta}$ are the model parameters, and $p_N : \mathbb{R}^{pN} \times \mathbb{R}^{mN} \times \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}_{\geq 0}$ is the conditional density function for N observations. Maximizing the conditional density is equivalent to minimizing its negative logarithm, so the ML estimate $\hat{\theta}_N$ is typically defined as a solution to

$$\max_{\theta \in \Theta} L_N(\theta) := \ln p(\mathbf{y}_{N-1} | \mathbf{u}_{N-1}, \theta) \quad (3.2)$$

where Θ is a problem-specific constraint set and $L_N : \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ is the (log-)likelihood function. The likelihood function L_N suppresses notation of the data $(\mathbf{u}_{N-1}, \mathbf{y}_{N-1})$. In statistical theory, L_N can be viewed as a random variable (i.e., a function of the data as a random variable). For

algorithms, however, it is more convenient to simply view L_N as the objective function.

In system identification, models are typically causal, meaning the current output $y(k)$ is only a function of past inputs and outputs ($\mathbf{u}_{k-1}, \mathbf{y}_{k-1}$) (and optionally the current input $u(k)$). In this case, it is convenient to successively condition on past data to give the equivalent problem,

$$\max_{\theta \in \Theta} L_N(\theta) = \sum_{k=0}^{N-1} \ln p(y_k | \mathbf{u}_{k-1}, \mathbf{y}_{k-1}, \theta). \quad (3.3)$$

If the parameters themselves are random variables with a known distribution, or we have some prior belief about what they should be, the model can include the prior distribution

$$\theta \sim p_0(\theta) \quad (3.4)$$

and we can define the following *maximum a posteriori* (MAP) problem, a close sibling of the ML problem (3.2):

$$\max_{\theta \in \Theta} L_N(\theta) + \ln p_0(\theta). \quad (3.5)$$

3.2 Problem statement

We consider stochastic LTI models in innovation form:

$$\hat{x}_{k+1} = A(\theta)\hat{x}_k + B(\theta)u_k + K(\theta)e_k \quad (3.6a)$$

$$y_k = C(\theta)\hat{x}_k + D(\theta)u_k + e_k \quad (3.6b)$$

$$e_k \stackrel{\text{iid}}{\sim} \mathcal{N}(0, R_e(\theta)) \quad (3.6c)$$

where $\hat{x} \in \mathbb{R}^n$ are the model states, $u \in \mathbb{R}^{n_u}$ are the inputs, $y \in \mathbb{R}^{n_y}$ are the outputs, $e \in \mathbb{R}^{n_y}$ are the innovation errors, and $\theta \in \Theta$ are the model parameters. The model functions $\mathcal{M}(\cdot) := (A(\cdot), B(\cdot), C(\cdot), D(\cdot), \hat{x}_0(\cdot), K(\cdot), R_e(\cdot))$ are assumed to be known. While the model \mathcal{M} is

kept fairly general throughout, it is advantageous to assume the model is identifiable in Θ . Last, for brevity, we often drop the dependence on the parameters $\theta \in \Theta$ and write the model functions as $\mathcal{M} = (A, B, C, D, \hat{x}_0, K, R_e)$.

While the subsequent developments apply to any model of the form (3.6), our main motivation is to identify the LADM,

$$\hat{s}_{k+1} = A_s(\theta)\hat{s}_k + B_d(\theta)\hat{d}_k + B_s(\theta)u_k + K_s(\theta)e_k \quad (3.7a)$$

$$\hat{d}_{k+1} = \hat{d}_k + K_d(\theta)e_k \quad (3.7b)$$

$$y_k = C_s(\theta)\hat{s}_k + C_d(\theta)\hat{d}_k + D(\theta)u_k + e_k \quad (3.7c)$$

$$e_k \stackrel{\text{iid}}{\sim} \text{N}(0, R_e(\theta)) \quad (3.7d)$$

where $\hat{s} \in \mathbb{R}^{n_s}$ denote plant states and $\hat{d} \in \mathbb{R}^{n_d}$ denote integrating disturbances. The LADM (3.7) is clearly a special case of (3.6) and can be put back into the standard form (3.6) by consolidating the plant and disturbance states $\hat{x}_k := \begin{bmatrix} \hat{s}_k^\top & \hat{d}_k^\top \end{bmatrix}^\top$ and defining

$$\begin{aligned} A &:= \begin{bmatrix} A_s & B_d \\ 0 & I \end{bmatrix}, & B &:= \begin{bmatrix} B_s \\ 0 \end{bmatrix}, \\ C &:= [C_s \quad C_d], & K &:= \begin{bmatrix} K_s \\ K_d \end{bmatrix}. \end{aligned}$$

Typically the LADM (3.7) is parameterized with (A_s, C_s) in observability canonical form (Denham, 1974), (B_d, C_d) fixed,¹ (B_s, K_s, K_d, R_e) fully parameterized, and $(D, \hat{s}_0, \hat{d}_0) = (0, 0, 0)$. Alternatively, we could choose a physics-based or gray-box plant model for the plant dynamics (A_s, B_s, C_s, D) .

¹With (A_s, B_s, C_s, D) fixed, all (B_d, C_d) such that (3.7) is observable are equivalent up to a similarity transform (Rajamani et al., 2009). Thus, (B_d, C_d) are chosen by the practitioner to maximize interpretability of the disturbance estimates.

3.2.1 Constrained maximum likelihood identification

The ML identification problem is defined as follows:

$$\max_{\theta \in \Theta} L_N(\theta) := -\frac{N}{2} \ln \det R_e(\theta) - \frac{1}{2} \sum_{k=0}^{N-1} |e_k(\theta)|_{[R_e(\theta)]^{-1}}^2 \quad (3.8)$$

where the $e_k(\theta)$ are given by the recursion (3.6) (Åström, 1979, p. 557), (Ljung, 1999, p. 219). Often, we may wish to regularize with respect to a previous parameter estimate $\bar{\theta}$, or incorporate an available prior distribution of the parameters $p_0(\theta)$. In either case, we consider the maximum a posteriori (MAP) estimation problem,

$$\max_{\theta \in \Theta} L_N(\theta) - R_0(\theta) \quad (3.9)$$

where $R_0(\theta) \propto -\ln p_0(\theta)$ is the regularization term, typically chosen as a distance from $\bar{\theta}$ (Sjöberg et al., 1993; Johansen, 1997). For example, a Gaussian prior or generalized ℓ_2 penalty can be implemented as

$$R_0(\theta) := \frac{1}{2} |\text{vec}(\theta) - \text{vec}(\bar{\theta})|_{V^{-1}}^2 \quad (3.10)$$

where $\bar{\theta} \in \Theta$ is the prior estimate, vec is a vectorization operator,² and $V \succ 0$ is the prior estimate variance. Such penalties are useful for model updating and re-identification. We typically use the penalty (3.10) with $V = \rho^{-1}I$. Later on, we transform the parameters θ into a more convenient space for optimization and find it more convenient to define the prior directly in the transformed space.

For plants of the form (3.6), the ML estimates are consistent and asymptotically efficient (Åström, 1979). In a standard setting, the plant is of the form (3.6) with $A - KC$ stable, and its coefficients are asymptotically recovered by (3.8). With sufficient data, the identified

²The vectorization operator may depend on the parameterization, as θ may contain both a vector portion and a sparse (semidefinite) matrix portion. The vectorization should only preserve the uniquely defined nonzero elements of the sparse matrix.

filter is stable. However, the LADM (3.7) is an intentional misspecification of the plant. Under certain regularity assumptions, we are consistent with respect to the estimates nearest in *relative entropy rate*, taken between the plant and model measurement distributions

$$\theta^* := \max_{\theta \in \Theta} N^{-1} \mathbb{E}[L_N(\theta)]$$

where the expectation is taken over the true distribution of measurements $(y_k)_{k=0}^{N-1}$ (White, 1984; Douc and Moulines, 2012). Identified LADM filters do not necessarily inherit stability from the plant, so we must design Θ to guarantee offset-free control.

3.2.2 Constraints

The constraint set Θ should capture both estimator design specifications and system knowledge. At a bare minimum, we require nondegeneracy of the innovation errors,

$$R_e(\theta) \succ 0 \tag{3.11}$$

and stability of the estimator,

$$\rho(A(\theta) - K(\theta)C(\theta)) < 1. \tag{3.12}$$

Other useful constraints include spectral abscissa bounds,

$$\alpha(-\tilde{A}(\theta)) < 0, \tag{3.13}$$

and bounds on the argument of the eigenvalues,

$$0 < |\text{Im}(\mu)|/\text{Re}(\mu) < \tan(\omega), \quad \forall \mu \in \lambda(\tilde{A}(\theta)) \tag{3.14}$$

for either the open-loop stability $\tilde{A} = A$ or estimator stability $\tilde{A} = A - KC$ matrices, to eliminate artificial high-frequency dynamics that may affect the control performance.

Chemical processes exhibit sparse interactions between units (mass and energy flows), especially for large-scale plants (Daoutidis et al., 2019; Tang et al., 2023). Sparse parameterizations of (A, B, C, D, K) are easily encoded, but the sparse parameterization of R_e is less obviously accomplished. While the covariance R_e for a centralized Kalman filter is dense even for sparse plants, correlations between distant units are small (Motee and Jadbabaie, 2009). Thus, it suffices to consider only nearest-neighbor correlations, e.g.,

$$R_e = \begin{bmatrix} R_{1,1} & R_{1,2} & & & \\ R_{1,2}^\top & R_{2,2} & \ddots & & \\ & \ddots & \ddots & & \\ & & & R_{N_u-1,N_u} & \\ & & & R_{N_u-1,N_u}^\top & R_{N_u,N_u} \end{bmatrix} \quad (3.15)$$

where $R_{i,j} \in \mathbb{R}^{p_u \times p_u}$ is the covariance between the innovations of the i -th and j -th process unit innovations. In (3.15), the sparse formulation introduces just $O(N_u p_u^2)$ variables compared to $O(N_u^2 p_u^2)$ variables for the dense formulation. Another algorithm goal is to simultaneously and efficiently enforce both (3.11) and (3.15). Finally, we remark that such constraints can be applied to the ML identification of any networked system with a time-invariant topology, as in Zamani et al. (2015).

3.2.3 Other parameterizations

The remainder of this section presents some other formulations of the ML identification problem. While we do not consider these formulations explicitly in our algorithm formulation (Section 3.3) or case studies (Chapter 5), the methods are readily generalized to these formulations.

Time-varying Kalman filter formulations

More generally, we could consider models of the following form:

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k + w_k \quad (3.16a)$$

$$y_k = C(\theta)x_k + D(\theta)u_k + v_k \quad (3.16b)$$

$$x_0 \sim \mathcal{N}(\hat{x}_0(\theta), \hat{P}_0(\theta)) \quad (3.16c)$$

$$\begin{bmatrix} w_k \\ v_k \end{bmatrix} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, S(\theta)) \quad (3.16d)$$

where $w \in \mathbb{R}^n$ and $v \in \mathbb{R}^{n_y}$ are the process and measurement noises and

$$\mathcal{M} := (A, B, C, D, \hat{x}_0, \hat{P}_0, S)$$

are the model functions. The noise covariance matrix $S(\theta)$ may be partitioned as

$$S(\theta) = \begin{bmatrix} Q_w(\theta) & S_{wv}(\theta) \\ [S_{wv}(\theta)]^\top & R_v(\theta) \end{bmatrix} \quad (3.17)$$

where $Q_w(\theta) \in \mathbb{S}_+^n$ is the process noise covariance, $S_{wv}(\theta)$ is the cross-covariance, and $R_v(\theta) \in \mathbb{S}_+^{n_y}$ is the measurement noise covariance. Throughout, we impose the stronger requirement $R_v(\theta) \succ 0$ on the measurement noise covariance.

For the model (3.16), the ML problem is defined as

$$\max_{\theta \in \Theta} L_N(\theta) := -\frac{1}{2} \sum_{k=0}^{N-1} \ln \det \mathcal{R}_k(\theta) - |e_k(\theta)|_{[\mathcal{R}_k(\theta)]^{-1}}^2 \quad (3.18)$$

where the e_k and \mathcal{R}_k are defined by the Kalman filtering equations

$$\hat{x}_{k+1} = A\hat{x}_k + Bu_k + \mathcal{K}_k e_k \quad (3.19a)$$

$$y_k = C\hat{x}_k + Du_k + e_k \quad (3.19b)$$

$$e_k \sim \mathcal{N}(0, \mathcal{R}_k) \quad (\text{indep.}) \quad (3.19c)$$

where

$$\hat{P}_{k+1} := A\hat{P}_k A^\top + Q_w - \mathcal{K}_k \mathcal{R}_k \mathcal{K}_k^\top \quad (3.19d)$$

$$\mathcal{K}_k := (A\hat{P}_k C^\top + S_{wv}) \mathcal{R}_k^{-1} \quad (3.19e)$$

$$\mathcal{R}_k := C\hat{P}_k C^\top + R_v. \quad (3.19f)$$

We remark that $R_v \succ 0$ suffices to guarantee the innovations are uniformly nondegenerate, i.e., $\mathcal{R}_k \succ 0$. However, stability of the filter is more difficult to guarantee as the early iterates $A - \mathcal{K}_k C$ may not be stable, even though the overall filter is stable, or vice versa. Instead, it is necessary to check that a stabilizing solution to the Riccati equation exists, which we elaborate on in the next formulation.

Time-invariant Kalman filter formulations

In most situations, the state error covariance matrix converges exponentially fast to a steady-state solution $\hat{P}_k \rightarrow \hat{P}$, so it suffices to consider the original steady-state filter model (3.6). In terms of the model (3.16), the steady-state filter takes the form $K := (A\hat{P}C^\top + S_{wv})R_e^{-1}$ and $R_e := C\hat{P}C^\top + R_v$, where \hat{P} is the unique, stabilizing solution to the discrete algebraic Riccati equation (DARE),

$$\hat{P} = A\hat{P}A^\top + Q_w - (A\hat{P}C^\top + S_{wv}) \times (C\hat{P}C^\top + R_v)^{-1} (A\hat{P}C^\top + S_{wv})^\top. \quad (3.20)$$

Recall a solution to the DARE (3.20) is stabilizing if the resulting $A_K := A - KC$ is stable.

Convergence of \hat{P}_k to \hat{P} is equivalent to the solution to the DARE (3.20) being unique and stabilizing. We generally assume such a solution exists, but for completeness, we state the following proposition, adapted from (Silverman, 1976, Thm. 18(iii)) (see Appendix 3.A.1 for proof).

Proposition 3.1. *Assume $R_v \succ 0$ and consider the full rank factorization*

$$\begin{bmatrix} Q_w & S_{wv} \\ S_{wv}^\top & R_v \end{bmatrix} = \begin{bmatrix} \tilde{B} \\ \tilde{D} \end{bmatrix} [\tilde{B}^\top \quad \tilde{D}^\top]$$

Then the following statements are equivalent:

1. *The DARE (3.20) has a unique, stabilizing solution $\hat{P} \succeq 0$.*
2. *The error covariance converges exponentially fast $\hat{P}_k \rightarrow \hat{P}$ for any $\hat{P}_0 \succeq 0$.*
3. *(A, C) is detectable and $(A - FC, \tilde{B} - F\tilde{D})$ is stabilizable for all $F \in \mathbb{R}^{n \times n_y}$.*

Remark 3.2. The hypothesis of Proposition 3.1 holds if we constrain A to be stable or (A, C) to be observable.

Remark 3.3. The cross-covariance S_{wv} complicates the filter stability analysis. With $S_{wv} = 0$, it would suffice to assume (A, C) detectable and (A, Q_w) stabilizable. With nonzero S_{wv} , however, a more elaborate stabilizability condition is needed. (Silverman, 1976, Thm. 18) considers the regulation problem with a cross-weighting term and semidefinite input weights. Proposition 3.1 specializes this result to the filter problem with positive definite R_v .

Remark 3.4. While $R_e(\theta)$ and $K(\theta)$ could be defined via $\hat{P}(\theta)$, taken as the function that returns solutions to the DARE (3.20) and therefore enforcing filters stability, it is more convenient to directly parameterize these matrices as in (3.6).

Minimum determinant formulation

Suppose in the model (3.6), that R_e is parameterized fully, and separately from the other terms, i.e.,

$$\mathcal{M}(\tilde{\theta}, R_e) = \left(A(\tilde{\theta}), B(\tilde{\theta}), C(\tilde{\theta}), D(\tilde{\theta}), \hat{x}_0(\tilde{\theta}), K(\tilde{\theta}), R_e \right).$$

Moreover, assume R_e is constrained separately as well, i.e.,

$$\Theta = \tilde{\Theta} \times \mathbb{S}_{++}^{n_y}.$$

Then we can always solve (3.8) stagewise, first in R_e , and then in the remaining variables $\tilde{\theta}$. Solving the inner problem gives the solution

$$\hat{R}_e(\tilde{\theta}) := \frac{1}{N} \sum_{k=0}^{N-1} e_k(\tilde{\theta}) [e_k(\tilde{\theta})]^\top$$

where we use the fact that e_k is only dependent on $\tilde{\theta}$, and we assume $\hat{R}_e(\tilde{\theta}) \succ 0$ for all $\tilde{\theta} \in \tilde{\Theta}$.

The outer problem can be written

$$\min_{\tilde{\theta} \in \tilde{\Theta}} \det \hat{R}_e(\tilde{\theta}). \quad (3.21)$$

The problem (3.21) is of relevance because it avoids posing (3.8) as a NSDP. It has been used both in the early ML identification literature (Åström and Eykhoff, 1971; Ljung, 1978; Åström, 1980) and in recent works (McKelvey et al., 2004; Ribarits et al., 2005; Li et al., 2007). None of these works consider filter stability constraints. To the best of our knowledge, only Umenberger et al. (2018) consider the ML problem (3.18) with stability constraints, but they consider open-loop stability (i.e., $\rho(A) < 1$) and use the EM algorithm.

Remark 3.5. For real-world data, $\det \hat{R}_e(\beta, \tilde{\Sigma}) = 0$ is not attainable because that would imply some direction of y_k were perfectly modeled. Therefore, $\hat{R}_e(\tilde{\theta}) \succ 0$ for all $\tilde{\theta} \in \tilde{\Theta}$ is a reasonable assumption.

3.3 Algorithm outline

3.3.1 Constraint set formulation

More generally, we seek to (i) impose eigenvalue constraints on any model function $\tilde{A}(\theta)$ and (ii) impose a sparsity structure on any semidefinite model function $\tilde{Q}(\theta)$.

Eigenvalue constraints

First, we define a LMI region.

Definition 3.6. We call $\mathcal{D} \subseteq \mathbb{C}$ an *LMI region* if

$$\mathcal{D} = \{ z \in \mathbb{C} \mid f_{\mathcal{D}}(z) := M_0 + M_1 z + M_1^\top \bar{z} \succ 0 \}$$

for some *generating matrices* $(M_0, M_1) \in \mathbb{S}^m \times \mathbb{R}^{m \times m}$. We call $f_{\mathcal{D}} : \mathbb{C} \rightarrow \mathbb{S}^m$ the *characteristic function* of \mathcal{D} .

The following lemma defines the four basic LMI regions: shifted half-planes, circles centered on the real axis, conic sections, and horizontal bands. For a general discussion of LMI regions properties, see Chilali and Gahinet (1996); Kushel (2019).

Lemma 3.7. For each $s, x_0 \in \mathbb{R}$, the subsets

$$\mathcal{D}_1(s) := \{ z \in \mathbb{C} \mid \operatorname{Re}(z) > s \}$$

$$\mathcal{D}_2(s, x_0) := \{ z \in \mathbb{C} \mid |z - x_0| < s \}$$

$$\mathcal{D}_3(s, x_0) := \{ z \in \mathbb{C} \mid |\operatorname{Im}(z)| < s(\operatorname{Re}(z) - x_0) \}$$

$$\mathcal{D}_4(s) := \{ z \in \mathbb{C} \mid |\operatorname{Im}(z)| < s \}$$

are LMI regions with characteristic functions

$$\begin{aligned} f_{\mathcal{D}_1(x_0)}(z) &:= -2x_0 + z + \bar{z} \\ f_{\mathcal{D}_2(s, x_0)}(z) &:= \begin{bmatrix} s & -x_0 \\ -x_0 & s \end{bmatrix} + \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} z + \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \bar{z} \\ f_{\mathcal{D}_3(s, x_0)}(z) &:= -2sx_0 I_2 + \begin{bmatrix} s & 1 \\ -1 & s \end{bmatrix} z + \begin{bmatrix} s & -1 \\ 1 & s \end{bmatrix} \bar{z} \\ f_{\mathcal{D}_4(s)}(z) &:= -2s I_2 + \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} z + \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \bar{z}. \end{aligned}$$

Proof. The first identity follows from the formula $2\operatorname{Re}(z) = z + \bar{z}$. For the second identity, we have $f_{\mathcal{D}_2(s, x_0)}(z) = \begin{bmatrix} s & z-x_0 \\ \bar{z}-x_0 & s \end{bmatrix} \succ 0$ if and only if $s > 0$ and $s^2 > |z - x_0|^2$, or equivalently, $|z - x_0| < s$. For the third identity, we have $f_{\mathcal{D}_3(s, x_0)}(z) = \begin{bmatrix} 2s(\operatorname{Re}(z)-x_0) & 2\iota\operatorname{Im}(z) \\ -2\iota\operatorname{Im}(z) & 2s(\operatorname{Re}(z)-x_0) \end{bmatrix} \succ 0$ if and only if $2s(\operatorname{Re}(z) - x_0) > 0$ and $4s^2(\operatorname{Re}(z) - x_0)^2 > 4|\operatorname{Im}(z)|^2$, or equivalently, $|\operatorname{Im}(z)| < s(\operatorname{Re}(z) - x_0)$. For the fourth identity, we have $f_{\mathcal{D}_4(s)}(z) = \begin{bmatrix} 2s & 2\iota\operatorname{Im}(z) \\ -2\iota\operatorname{Im}(z) & 2s \end{bmatrix} \succ 0$ if and only if $2s > 0$ and $4s^2 > 4|\operatorname{Im}(z)|^2$, or equivalently, $|\operatorname{Im}(z)| < s$. \square

Remark 3.8. For continuous-time systems, $-\mathcal{D}_1(\alpha)$ corresponds to a minimum decay rate of $\alpha > 0$, $\mathcal{D}_3(-\tan(\omega), 0)$ corresponds to a minimum damping ratio $\cos(\omega)$, and $\mathcal{D}_2(r, 0) \cap \mathcal{D}_3(-\tan(\omega), 0)$ implies to a maximum undamped natural frequency $r \sin(\omega)$, where $\alpha, r > 0$ and $\omega \in [0, \pi/2]$ (Chilali and Gahinet, 1996). For discrete-time systems, $\mathcal{D}_2(r, 0)$ corresponds to a minimum decay rate of $-\ln r$, and $\mathcal{D}_2(r, 0) \cap \mathcal{D}_3(\tan(\omega), 0)$ implies a minimum damping ratio $-\cos(\tan^{-1}(\omega/\ln r))$ and maximum natural frequency $(\ln(r)^2 + \omega^2)/\Delta$, where $r > 0$, $\omega \in [0, \pi/2]$, and Δ is the sample time.

Remark 3.9. For any LMI region \mathcal{D} (including those in Lemma 3.7), the set \mathcal{D} is convex, open, and symmetric about the real axis. The intersection of two LMI regions $\mathcal{D} := \mathcal{D}_1 \cap \mathcal{D}_2$ is an LMI region with the characteristic function $f_{\mathcal{D}}(z) = f_{\mathcal{D}_1}(z) \oplus f_{\mathcal{D}_2}(z)$. By this property, we can construct any convex polyhedron that is symmetric about the real axis by intersecting left and right half-planes, horizontal strips, and conic sections. Moreover, since any convex

region can be approximated, to any desired accuracy, by a convex polyhedron, the set of LMI regions is dense in the space of convex subsets of \mathbb{C} that are symmetric about the real axis. An LMI region \mathcal{D} with characteristic function $f_{\mathcal{D}}$ also has characteristic function $M f_{\mathcal{D}}(\cdot) M^{\top}$ for any nonsingular $M \in \mathbb{R}^{m \times m}$. For an in-depth discussion of LMI region geometry and other properties, see Kushel (2019).

In Chilali and Gahinet (1996), it is shown a matrix $\tilde{A} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ has eigenvalues in a LMI region \mathcal{D} if and only if the following system of matrix inequalities is feasible:

$$M_{\mathcal{D}}(\tilde{A}, P) \succ 0, \quad P \succ 0 \quad (3.22)$$

where the *matrix characteristic function* $M_{\mathcal{D}} : \mathbb{R}^{\tilde{n} \times \tilde{n}} \times \mathbb{S}^{\tilde{n}} \rightarrow \mathbb{S}^{\tilde{n}\tilde{m}}$ of \mathcal{D} is defined by

$$M_{\mathcal{D}}(\tilde{A}, P) := M_0 \otimes P + M_1 \otimes (\tilde{A}P) + M_1^{\top} \otimes (\tilde{A}P)^{\top}. \quad (3.23)$$

From this equivalence, we can build tractable eigenvalue constraints. For the constraint (3.12), Lemma 3.7 gives the generating matrices $(M_0, M_1) := ([\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}], [\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}])$ and we have the matrix inequalities

$$\begin{bmatrix} P & (A - KC)P \\ P(A - KC)^{\top} & P \end{bmatrix} \succ 0, \quad P \succ 0$$

which is a well-known LMI for checking stability (Boyd et al., 1994). Similarly, to implement (3.13), we can use the generating matrices $(M_0, M_1) := (0, 1)$, and to implement (3.14), we can use $(M_0, M_1) := \left(\left[\begin{array}{cc} -2 \tan(\omega) & 0 \\ 0 & -2 \tan(\omega) \end{array} \right], \left[\begin{array}{cc} \tan(\omega) & 1 \\ -1 & \tan(\omega) \end{array} \right] \right)$.

The system of matrix inequalities (3.22) contains only strict inequalities, so we “tighten”

them as follows:

$$M_{\mathcal{D}}(\tilde{A}, P) \succeq M, \quad P \succeq 0, \quad \text{tr}(VP) \leq \varepsilon^{-1} \quad (3.24)$$

where $\varepsilon > 0$, $V \in \mathbb{S}_{++}^{\tilde{n}}$, and $M \in \mathbb{S}_{++}^{\tilde{n}\tilde{m}}$. The set of $\tilde{A} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ for which (3.24) is feasible defines a *closed* set for which $\lambda(\tilde{A}) \subseteq \mathcal{D}$. In Section 3.4, we show this fact and other properties of the constraint (3.24).

Sparsity structure

To encode sparsity information, we adapt the notation of Burer et al. (2002a). Define the index sets $\mathcal{L}^n := \{(i, j) \in \mathbb{I}_{1:n}^2 \mid i \geq j\}$ and $\mathcal{D}^n := \{(i, i) \in \mathbb{I}_{1:n}^2\}$ corresponding to the sparsity patterns of $n \times n$ lower triangular and diagonal matrices. With a slight abuse of notation, we define the direct sum of index sets $\mathcal{I} \subseteq \mathcal{L}^n$ and $\mathcal{J} \subseteq \mathcal{L}^m$ by

$$\mathcal{I} \oplus \mathcal{J} := \mathcal{I} \cup \{(i+n, j+n) \mid (i, j) \in \mathcal{J}\} \subseteq \mathcal{L}^{n+m}.$$

For each $\mathcal{I} \subseteq \mathcal{L}^n$, define the sets

$$\begin{aligned} \mathbb{S}^n[\mathcal{I}] &:= \{S \in \mathbb{S}^n \mid S_{ij} = 0 \forall (i, j) \in \mathcal{L}^n \setminus \mathcal{I}\} \\ \mathbb{L}^n[\mathcal{I}] &:= \{L \in \mathbb{L}^n \mid L_{ij} = 0 \forall (i, j) \notin \mathcal{I}\} \\ \mathbb{L}_{++}^n[\mathcal{I}] &:= \{L \in \mathbb{L}_{++}^n \mid L_{ij} = 0 \forall (i, j) \notin \mathcal{I}\}. \end{aligned}$$

Finally, let $\text{vecs}_{\mathcal{I}} : \mathbb{S}^n \rightarrow \mathbb{R}^{|\mathcal{I}|}$ denote the operator that vectorizes the $|\mathcal{I}|$ entries of the argument corresponding to the index set \mathcal{I} .

Constraint definition

To combine the LMI region and sparsity constraints, we partition the parameter into vector and sparse symmetric matrix parts, i.e., $\theta = (\beta, \Sigma)$, and define the constraint set Θ by

$$\Theta = \{ (\beta, \Sigma) \in \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma}[\mathcal{I}_\Sigma] \mid g(\beta, \Sigma) = 0, h(\beta, \Sigma) \leq 0, \Sigma \succeq H(\beta), \mathcal{A}(\beta, \Sigma) \succeq 0 \} \quad (3.25)$$

where $\mathcal{D}^{n_\Sigma} \subseteq \mathcal{I}_\Sigma \subseteq \mathcal{L}^{n_\Sigma}$, $\mathcal{D}^{n_\mathcal{A}} \subseteq \mathcal{I}_\mathcal{A} \subseteq \mathcal{L}^{n_\mathcal{A}}$, $g : \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma} \rightarrow \mathbb{R}^{n_g}$, $h : \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma} \rightarrow \mathbb{R}^{n_h}$, $H : \mathbb{R}^{n_\beta} \rightarrow \mathbb{S}^{n_\Sigma}$, and $\mathcal{A} : \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma} \rightarrow \mathbb{S}^{n_\mathcal{A}}[\mathcal{I}_\mathcal{A}]$. The purpose of the partition $\theta = (\beta, \Sigma)$ is to clearly delineate the sparse semidefinite matrix argument Σ from the remaining parameters β . The index set \mathcal{I}_Σ defines the sparsity pattern of Σ and H , and the index set $\mathcal{I}_\mathcal{A}$ defines the sparsity pattern of \mathcal{A} .

Remark 3.10. Assumption 3.13 rules out direct use strict inequalities, e.g., $R_e(\theta) \succ 0$ or $R_v(\theta) \succ 0$. To satisfy nondegeneracy requirements, we use the closed constraint $R_e(\theta) \succeq \delta I_p$ with a small backoff $\delta > 0$.

Remark 3.11. Typically, the index set \mathcal{I}_Σ encodes block diagonal structures, e.g., for the model (3.16), $\Sigma = \hat{P}_0 \oplus Q_w \oplus R_v \in \mathbb{S}^{2n+p}[\mathcal{I}_\Sigma]$ where $\mathcal{I}_\Sigma := \mathcal{L}^n \oplus \mathcal{L}^n \oplus \mathcal{L}^p$. However, more general structures (e.g., (3.15)) can be stated. For the time-varying formulation (3.18), we may further restrict Q_w and R_v to take block tridiagonal and diagonal structures, e.g.,

$$Q_w = \begin{bmatrix} Q_{1,1} & Q_{1,2} & & & \\ Q_{1,2}^\top & Q_{2,2} & \ddots & & \\ & \ddots & \ddots & & \\ & & & Q_{\tilde{n}-1,\tilde{n}} & \\ & & & Q_{\tilde{n}-1,\tilde{n}}^\top & Q_{\tilde{n},\tilde{n}} \end{bmatrix}, \quad R_v = R_1 \oplus \dots \oplus R_{\tilde{n}}$$

that arise in sequentially interconnected processes such as chemical plants. Adding a $Q_{1,\tilde{n}}$ block can account for an overall recycle loop. Note that if we parameterize the block tridiagonal Q_w via a sparse shaping matrix (i.e., $Q_w = G_w G_w^\top$), then there are more parameters than

if the sparsity of Q_w is known.

Remark 3.12. As alluded to in Section 3.2, the Riccati equation solution has a dense solution, but the entries far from the core sparsity pattern decay rapidly. Thus, we can approximate an eigenvalue constraint, e.g., $P - APA^\top \succ 0$, as a function that maps to the same sparsity pattern as A Motee and Jadbabaie (2009); Haber and Verhaegen (2016); Motee and Sun (2017); Massei and Saluzzi (2024).

3.3.2 Cholesky factorization and elimination

At this juncture, the ML and MAP problems (3.8) and (3.9) with the constraints (3.25) are in standard NSDP form and can be solved with any dedicated NSDP solver, e.g., Fiala et al. (2013); Kočvara and Stingl (2015). However, such solvers are neither as widely available nor as well-understood as NLP solvers such as IPOPT (Wächter and Biegler, 2006).

The Burer-Monteiro-Zhang (BMZ) method is a Cholesky factorization-based substitution and elimination algorithm that can convert certain NSDPs to NLPs (Burer et al., 2002a,b). In Section 3.5, we consider a generalization of this algorithm to (approximately) transform a given NSDP into a NLP while only introducing $|\mathcal{I}_A|$ new variables. This generalization requires the following assumption:

Assumption 3.13. The model functions \mathcal{M} are twice differentiable and the constraint functions \mathcal{C} are differentiable. Moreover, $\text{cl}(\Theta_{++}) = \Theta$ where

$$\Theta_{++} := \{ (\beta, \Sigma) \in \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma}[\mathcal{I}_\Sigma] \mid g(\beta, \Sigma) = 0, h(\beta, \Sigma) \leq 0, \Sigma \succ H(\beta), \mathcal{A}(\beta, \Sigma) \succ 0 \}. \quad (3.26)$$

In Section 3.5, we construct functions

$$\mathcal{T} : \mathbb{R}^{n_\beta} \times \mathbb{L}_{++}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_{++}^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}] \rightarrow \mathbb{R}^{n_\beta} \times \mathbb{S}^{n_\Sigma}[\mathcal{I}_\Sigma]$$

$$\mathcal{A}_{\mathcal{T}} : \mathbb{R}^{n_\beta} \times \mathbb{L}_{++}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_{++}^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}] \rightarrow \mathbb{S}_{++}^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}]$$

and define transformed constraint functions

$$g_{\mathcal{T}}(\phi) := \begin{bmatrix} g(\mathcal{T}(\phi)) \\ \text{vecs}_{\mathcal{I}_{\mathcal{A}}}(\mathcal{A}(\mathcal{T}(\phi)) - \mathcal{A}_{\mathcal{T}}(\phi)) \end{bmatrix} \quad (3.27a)$$

$$h_{\mathcal{T}}(\phi) := h(\mathcal{T}(\phi)) \quad (3.27b)$$

and a transformed constraint set

$$\Phi := \{ \phi \in \mathbb{R}^{n_\beta} \times \mathbb{L}_{++}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_{++}^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}] \mid g_{\mathcal{T}}(\phi) = 0, h_{\mathcal{T}}(\phi) \leq 0 \} \quad (3.28)$$

such that, under Assumption 3.13, \mathcal{T} is an invertible map from Φ to Θ_{++} . Finally, to eliminate the strict inequalities on the diagonal entries of $(L_\Sigma, L_{\mathcal{A}}) \in \mathbb{L}_{++}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_{++}^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}]$, we introduce a fixed lower bound $\varepsilon > 0$ on the diagonal entries,

$$\Phi_\varepsilon := \{ \phi \in \mathbb{R}^{n_\beta} \times \mathbb{L}_\varepsilon^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_\varepsilon^{n_{\mathcal{A}}}[\mathcal{I}_{\mathcal{A}}] \mid g_{\mathcal{T}}(\phi) = 0, h_{\mathcal{T}}(\phi) \leq 0 \} \quad (3.29)$$

where we have defined, for any $\varepsilon > 0$ and $\mathcal{I} \subseteq \mathcal{L}^n$,

$$\mathbb{L}_\varepsilon^n[\mathcal{I}] := \{ L \in \mathbb{L}^n[\mathcal{I}] \mid L_{ii} \geq \varepsilon \forall i \in \mathbb{I}_{1:n} \}.$$

We define the *approximate* transformed problem as

$$\max_{\phi \in \Phi_\varepsilon} L_N(\mathcal{T}(\phi)) + R_0(\mathcal{T}(\phi)). \quad (3.30)$$

Algorithm 1 Identification of an innovation form model (3.6) with eigenvalue constraints and the Cholesky factor-based substitution and elimination scheme.

Require: Model functions $\mathcal{M} = (A, B, C, D, \hat{x}_0, K, R_e)$, regularization term R_0 , initial parameters $\theta_0 = (\beta, \Sigma_0)$ constraint functions $(g, h_0, H_0, \mathcal{A}_0)$ and sparsity patterns $(\mathcal{I}_{\Sigma_0}, \mathcal{I}_{\mathcal{A}_0})$, a series of LMI region constraints $(\mathcal{D}_i, \tilde{A}_i(\cdot))_{i=1}^{n_c}$, and small $\varepsilon, \varepsilon_i > 0$.

- 1: For each $i \in \mathbb{I}_{1:n_c}$, let $M_{\mathcal{D}_i} : \mathbb{R}^{n_i \times n_i} \times \mathbb{S}^{n_i} \rightarrow \mathbb{S}^{n_i m_i}$ denote the matrix characteristic function for \mathcal{D}_i .
- 2: Extend the parameters $\Sigma := \Sigma_0 \oplus (\bigoplus_{i=1}^{n_c} P_i)$ and $\theta := (\beta, \Sigma)$ with $P_i \in \mathbb{S}^{n_i}$.
- 3: Extend the constraint functions

$$h(\theta) := \begin{bmatrix} h_0(\theta_0) \\ \text{tr}(V_1 P_1) - \varepsilon_1^{-1} \\ \vdots \\ \text{tr}(V_{n_c} P_{n_c}) - \varepsilon_{n_c}^{-1} \end{bmatrix}$$

$$H(\beta) := H_0(\beta) \oplus (\bigoplus_{i=1}^{n_c} 0_{n_i \times n_i})$$

$$\mathcal{A}(\theta) := \mathcal{A}_0(\theta_0) \oplus \left(\bigoplus_{i=1}^{n_c} M_{\mathcal{D}_i}(\tilde{A}_i(\theta_0), P_i) - \varepsilon_i I \right).$$

- 4: Extend the index sets

$$\mathcal{I}_{\Sigma} := \mathcal{I}_{\Sigma_0} \oplus (\bigoplus_{i=1}^{n_c} \mathcal{L}^{n_i})$$

$$\mathcal{I}_{\mathcal{A}} := \mathcal{I}_{\mathcal{A}_0} \oplus (\bigoplus_{i=1}^{n_c} \mathcal{L}^{n_i m_i}).$$

- 5: Form the functions $\mathcal{T}, \mathcal{T}^{-1}$, and $\tilde{\mathcal{A}}$ as in Section 3.5.
 - 6: Form the transformed constraint functions (3.27).
 - 7: Solve (3.29) and (3.30), and let $\hat{\phi}$ denote the solution.
 - 8: Let $\hat{\theta} := \mathcal{T}(\hat{\phi})$.
-

If $\hat{\phi}$ solves the problem (3.30), then $\hat{\theta} := \mathcal{T}(\hat{\phi})$ *approximately* solves the MAP problem (3.9).

We recover the ML problem (3.8) and its approximate solutions with $R_0 \equiv 0$.

3.3.3 Algorithm summary

Algorithm 1 provides an example of our approach towards solving the identification problem (3.9) with eigenvalue constraints and the Cholesky factor-based substitution and elimination scheme.

3.4 Eigenvalue constraints

In this section, we elaborate on the LMI region constraints previewed in Section 3.3. Throughout, assume the LMI region \mathcal{D} is nonempty, not equal to \mathbb{C} , and its characteristic function $f_{\mathcal{D}}$ and generating matrices (M_0, M_1) are fixed. Our goal in this section is to define, using only matrix inequalities, a *closed* set of matrices $A \in \mathbb{R}^{n \times n}$ such that $\lambda(A) \subseteq \mathcal{D}$. For this section, the matrix $A \in \mathbb{R}^{n \times n}$ need not have any relation to the model function in (3.6), and can in fact be any square matrix of any dimension (e.g., the filter stability matrix $A - KC$, the plant stability matrix A_s from (3.7), or any submatrix thereof). Throughout this section, we assume the matrix characteristic function $M_{\mathcal{D}}$ is fixed.

3.4.1 LMI region constraints

Originally, Chilali and Gahinet (1996) proved the following theorem relating the eigenvalues of $A \in \mathbb{R}^{n \times n}$ to feasibility of a system of matrix inequalities.

Theorem 3.14 ((Chilali and Gahinet, 1996, Thm. 2.2)). *For any $A \in \mathbb{R}^{n \times n}$, we have $\lambda(A) \subseteq \mathcal{D}$ if and only if*

$$M_{\mathcal{D}}(A, P) \succ 0, \quad P \succ 0. \quad (3.31)$$

holds for some $P \in \mathbb{S}^n$.

Ultimately, we seek matrix inequalities that define a *closed* set of constraints. Due to the strictness of the inequalities (3.31), it is unlikely that (Chilali and Gahinet, 1996, Thm. 2.2) achieves this goal.

3.4.2 Relaxed constraints

In Miller and De Callafon (2013), the following relaxation of (3.31) was considered,

$$M_{\mathcal{D}}(A, P) \succeq 0, \quad P \succ 0. \quad (3.32)$$

Since $M_{\mathcal{D}}(A, P)$ is linear in P , feasibility of (3.32) is equivalent to feasibility of

$$M_{\mathcal{D}}(A, P) \succeq 0, \quad P \succeq P_0 \quad (3.33)$$

for some fixed $P_0 \in \mathbb{S}_{++}^n$.³

An attempt was made in (Miller and De Callafon, 2013, Thm. 1) to characterize the eigenvalues of matrices $A \in \mathbb{R}^{n \times n}$ for which (3.32) is feasible, but this theorem does not correctly treat eigenvalues on the LMI region's boundary $\partial\mathcal{D}$. We restate (Miller and De Callafon, 2013, Thm. 1) below as a conjecture and disprove it with a simple counterexample.

Conjecture 3.15 ((Miller and De Callafon, 2013, Thm. 1)). *The matrix $A \in \mathbb{R}^{n \times n}$ satisfies $\lambda(A) \subset \text{cl}(\mathcal{D})$ if and only if (3.32) holds for some $P \in \mathbb{S}^n$.*

Counterexample. Let \mathcal{D} be the left half-plane, consider the Jordan block $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$, and suppose $P = \begin{bmatrix} p_{11} & p_{12} \\ p_{12} & p_{22} \end{bmatrix} \in \mathbb{S}^2$ such that (3.32) holds. Then $\lambda(A) \subset \text{cl}(\mathcal{D})$ and

$$0 \preceq M_{\mathcal{D}}(A, P) = - \begin{bmatrix} 2p_{12} & p_{22} \\ p_{22} & 0 \end{bmatrix}$$

which implies $p_{12} = p_{22} = 0$, a contradiction of (3.32). *

The correction to Conjecture 3.15 requires a more careful treatment of eigenvalues lying on the the LMI region's boundary $\partial\mathcal{D}$. Specifically, we show in the following proposition that

³For any $P_0 \succ 0$ and P satisfying (3.32), define the scaling factor $\gamma := \|P_0\|_2 \|P^{-1}\|_2$ and a rescaled solution $P^* := \gamma P$. Then $P^* \succeq P_0$ and $M_{\mathcal{D}}(A, P^*) = \gamma M_{\mathcal{D}}(A, P) \succeq 0$.

feasibility of (3.32) for a given $A \in \mathbb{R}^{n \times n}$ is equivalent to the eigenvalues of A being in $\text{cl}(\mathcal{D})$, with all non-simple eigenvalues lying in \mathcal{D} (see Appendix 3.A.2 for proof).

Proposition 3.16. *The matrix $A \in \mathbb{R}^{n \times n}$ satisfies $\lambda(A) \subseteq \text{cl}(\mathcal{D})$ and $\lambda \in \mathcal{D}$ for all non-simple eigenvalues $\lambda \in \lambda(A)$ if and only if (3.32) holds for some $P \in \mathbb{S}^n$.*

3.4.3 Tightened constraints

Instead of the “relaxed” constraints (3.32), we consider “tightened” constraints of the form

$$M_{\mathcal{D}}(A, P) \succeq M, \quad P \succeq 0, \quad \text{tr}(VP) \leq \varepsilon^{-1} \quad (3.34)$$

where $M \in \mathbb{S}_+^{nm}$ and $V \in \mathbb{S}_{++}^n$ are fixed and chosen in a way that (3.34) implies (3.31). While we allow M to be semidefinite,⁴ in the following proposition, we show $M \succ 0$ always suffices.

Proposition 3.17. *Suppose $M \in \mathbb{S}_{++}^{nm}$ and $V \in \mathbb{S}_{++}^n$. Then (3.34) implies (3.31) for all $A \in \mathbb{R}^{n \times n}$ and $\varepsilon > 0$.*

Proof. With $M \succ 0$ and (3.34), we automatically have $M_{\mathcal{D}}(A, P) \succ 0$. It remains to show (3.34) implies $P \succ 0$. For contradiction suppose (3.34) and $M \succ 0$, but $P \not\succeq 0$. Then there exists a nonzero $v \in \mathbb{R}^n$ such that $Pv = 0$, and

$$(I_m \otimes v)^\top M_{\mathcal{D}}(A, P)(I_m \otimes v) = M_0 \otimes (v^\top Pv) + M_1 \otimes (v^\top APv) + M_1^\top \otimes (v^\top PA^\top v) = 0$$

a contradiction of the assumption $M_{\mathcal{D}}(A, P) \succeq M \succ 0$. □

Remark 3.18. The tightened constraint (3.34) was inspired by a similar set of constraints was introduced by Diehl et al. (2009) to “smooth” the spectral radius. Specifically, feasibility

⁴For some LMI regions, $M \succeq 0$ is advantageous. For example, we can always take $M := \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \otimes Q$ with $Q \succ 0$ for circular LMI regions. Then we can reduce the constraint dimension by taking the Schur complement.

of the nonlinear system

$$s^2P - APA^\top = W, \quad P \succeq 0, \quad \text{tr}(VP) \leq \varepsilon^{-1} \quad (3.35)$$

implies $\rho(A) < s$ where $W, V \in \mathbb{S}_{++}^n$ and $s, \varepsilon > 0$ are fixed (Diehl et al., 2009, Thms. 5.4, 5.6). Similarly, the spectral abscissa was “smoothed” in (Vanbiervliet et al., 2009, Thms. 2.5, 2.6), and it is straightforward to generalize (Diehl et al., 2009, Thms. 5.4, 5.6) to show feasibility of

$$(A - sI)P + P(A - sI)^\top = -W, \quad P \succeq 0, \quad \text{tr}(VP) \leq \varepsilon^{-1} \quad (3.36)$$

implies $\alpha(A) < s$ where $W, V \in \mathbb{S}_{++}^n$, $s \in \mathbb{R}$, and $\varepsilon > 0$ are fixed. The authors do not discuss LMI regions and the results are not obviously generalizable to them.

3.4.4 Constraint topology

Consider the constraint sets,

$$\mathbb{A}_{\mathcal{D}}^n := \{ A \in \mathbb{R}^{n \times n} \mid \exists P \in \mathbb{S}^n : (3.31) \text{ holds} \}$$

$$\tilde{\mathbb{A}}_{\mathcal{D}}^n := \{ A \in \mathbb{R}^{n \times n} \mid \exists P \in \mathbb{S}^n : (3.32) \text{ holds} \}$$

$$\mathbb{A}_{\mathcal{D}}^n(\varepsilon) := \{ A \in \mathbb{R}^{n \times n} \mid \exists P \in \mathbb{S}^n : (3.34) \text{ holds} \}.$$

The following proposition characterizes the topology of $\mathbb{A}_{\mathcal{D}}^n$ and $\tilde{\mathbb{A}}_{\mathcal{D}}^n$ (see Appendix 3.A.3 for proof).

Proposition 3.19. (a) $\mathbb{A}_{\mathcal{D}}^n$ is open.

(b) $\tilde{\mathbb{A}}_{\mathcal{D}}^n$ is not open if (i) $n \geq 2$ or (ii) $\partial\mathcal{D} \cap \mathbb{R}$ is nonempty.

(c) $\tilde{\mathbb{A}}_{\mathcal{D}}^n$ is not closed if (i) $n \geq 4$ or (ii) $\partial\mathcal{D} \cap \mathbb{R}$ is nonempty and $n \geq 2$.

(d) $\text{cl}(\mathbb{A}_{\mathcal{D}}^n) = \{ A \in \mathbb{R}^{n \times n} \mid \lambda(A) \subset \text{cl}(\mathcal{D}) \}$.

Proposition 3.19 reveals a weakness of the relaxed constraints (3.32) and (3.33). Since $\tilde{\mathbb{A}}_{\mathcal{D}}^n$ is not closed, any feasible path towards a matrix $A \in \text{cl}(\mathbb{A}_{\mathcal{D}}^n) \setminus \tilde{\mathbb{A}}_{\mathcal{D}}^n$ has no feasible limiting P . In fact, P will grow unbounded along the path of iterates.

To analyze the topology of $\mathbb{A}_{\mathcal{D}}^n(\varepsilon)$, we take a barrier function approach. Consider the parameterized linear SDP,

$$\phi_{\mathcal{D}}(A) := \inf_{P \in \mathbb{S}_+^n} \text{tr}(VP) \text{ subject to } M_{\mathcal{D}}(A, P) \succeq M. \quad (3.37)$$

The optimal value function $\phi_{\mathcal{D}} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a barrier function for the constraint $A \in \mathbb{A}_{\mathcal{D}}^n$. Theorem 3.20 establishes properties of $\phi_{\mathcal{D}}$ and its ε^{-1} -sublevel sets (see Appendix 3.A.4 for proof).

Theorem 3.20. *Let $V \in \mathbb{S}_{++}^n$ and $M \in \mathbb{S}_+^n$ such that $M_{\mathcal{D}}(A, P) \succeq M$ implies $M_{\mathcal{D}}(A, P) \succ 0$. Then*

- (a) $\phi_{\mathcal{D}}$ is continuous on $\mathbb{A}_{\mathcal{D}}$;
- (b) for each $\varepsilon > 0$, $\mathbb{A}_{\mathcal{D}}^n(\varepsilon)$ is equivalent to the ε^{-1} -sublevel set of $\phi_{\mathcal{D}}$, i.e.,

$$\mathbb{A}_{\mathcal{D}}^n(\varepsilon) = \{ A \in \mathbb{R}^{n \times n} \mid \phi_{\mathcal{D}}(A) \leq \varepsilon^{-1} \} \quad (3.38)$$

and both are closed; and

- (c) $\mathbb{A}_{\mathcal{D}}^n(\varepsilon) \nearrow \mathbb{A}_{\mathcal{D}}^n$ as $\varepsilon \searrow 0$.

Remark 3.21. To reconstruct (3.35) via Theorem 3.20, we set $M = sW \oplus 0_{n \times n}$ for any $W, V \succ 0$ and $s > 0$ and apply the Schur complement lemma to $M_{\mathcal{D}_2}(A, P)/s - M/s$, where \mathcal{D}_2 is the circle defined in Lemma 3.7 with $x_0 = 0$, and $M_{\mathcal{D}_2}$ is defined by the generating matrices used in Lemma 3.7. Then the ε^{-1} -sublevel set of $\phi_{\mathcal{D}_2}$ equals the set of $A \in \mathbb{R}^{n \times n}$ for which (3.35) is feasible.

Similarly, we can reconstruct the set of $A \in \mathbb{R}^{n \times n}$ for which (3.36) is feasible as ε^{-1} -

sublevel sets of $\phi_{\mathcal{D}_1}$, where \mathcal{D}_1 is the shifted half-plane defined in Lemma 3.7, and $M = W$ for any $W, V \succ 0$.

3.5 Cholesky substitution and elimination

In this section, we seek to approximate certain NSDPs by NLPs. Specifically, we consider the NSDP

$$\min_{(\beta, \Sigma) \in \Theta} f(\beta, \Sigma) \quad (3.39)$$

where Θ is defined as in (3.25). This covers both ML (3.8) and MAP (3.9) problems with constraints (3.25). We combine Cholesky factor-based substitution with an elimination scheme to convert the NSDP to a NLP while adding just $|\mathcal{I}_A|$ variables to the optimization problem.

For this section, we define the following notation. For each $\mathcal{I} \subseteq \mathcal{L}^n$, let $\pi_{\mathcal{I}}^L : \mathbb{R}^{n \times n} \rightarrow \mathbb{L}^n[\mathcal{I}]$ and $\pi_{\mathcal{I}} : \mathbb{R}^{n \times n} \rightarrow \mathbb{S}^n[\mathcal{I}]$ denote the orthogonal projections (in the Frobenius norm) from $\mathbb{R}^{n \times n}$ onto the subspaces $\mathbb{L}^n[\mathcal{I}]$ and $\mathbb{S}^n[\mathcal{I}]$, respectively. Let $\text{chol} : \mathbb{S}_{++}^n \rightarrow \mathbb{L}_{++}^n$ denote the invertible function that maps a positive definite matrix to its Cholesky factor.

3.5.1 Burer-Monteiro-Zhang method

We first consider the simplified constraint set

$$\mathcal{P} := \{ (x, Q) \in \mathbb{R}^m \times \mathbb{S}^n[\mathcal{I}] \mid Q \succeq H(x) \} \quad (3.40)$$

where $\mathcal{D}^n \subseteq \mathcal{I} \subseteq \mathcal{L}^n$ and $H : \mathbb{R}^m \rightarrow \mathbb{S}^n$. As in Burer et al. (2002a), we parameterize the matrix argument Q in a way that automatically enforces the constraint $Q \succeq H(x)$ while introducing just n scalar inequality constraints.

Recall $Q \succeq H$ if and only if $Q = H + LL^\top$ for the *unique* matrix $L = \text{chol}(Q - H) \in \mathbb{L}_{++}^n$.

Algorithm 2 Cholesky factorization algorithm for solving systems of the form (3.42) based on (Burer et al., 2002a, Lem. 1).

Require: $\mathcal{D}^n \subseteq \mathcal{I} \subseteq \mathcal{L}^n$, $L^{\mathcal{I}} \in \mathbb{L}_{++}^n[\mathcal{I}]$, and $H \in \mathbb{S}^n$

- 1: $(\mathcal{J}, L^{\mathcal{J}}) \leftarrow (\mathcal{L}^n \setminus \mathcal{I}, 0_{n \times n})$
 - 2: **for** each $(i, j) \in \mathcal{J}$ in ascending lexicographic order **do**
 - 3: $L_{ij}^{\mathcal{J}} \leftarrow -\frac{1}{L_{jj}^{\mathcal{I}}}(H_{ij} + \sum_{k=1}^{j-1}(L_{ik}^{\mathcal{I}} + L_{ik}^{\mathcal{J}})(L_{jk}^{\mathcal{I}} + L_{jk}^{\mathcal{J}}))$
 - 4: **end for**
 - 5: **return** $L^{\mathcal{J}}$
-

With $\mathcal{J} := \mathcal{L}^n \setminus \mathcal{I}$, we can split L into a sum of $L^{\mathcal{I}} \in \mathbb{L}_{++}^n[\mathcal{I}]$ and $L^{\mathcal{J}} \in \mathbb{L}^n[\mathcal{J}]$, giving

$$Q = H + (L^{\mathcal{I}} + L^{\mathcal{J}})(L^{\mathcal{I}} + L^{\mathcal{J}})^{\top}. \quad (3.41)$$

But $Q \in \mathbb{S}^n[\mathcal{I}]$, so we can apply the vectorization operator $\text{vecs}_{\mathcal{J}}$ on both sides to give

$$\text{vecs}_{\mathcal{J}}(H + (L^{\mathcal{I}} + L^{\mathcal{J}})(L^{\mathcal{I}} + L^{\mathcal{J}})^{\top}) = 0. \quad (3.42)$$

Equation (3.42) defines $|\mathcal{J}|$ equations to solve for the $|\mathcal{J}|$ variables of $L^{\mathcal{J}}$, where each $L_{ij}^{\mathcal{J}}$ is fully specified by H_{ij} and the $L_{i'j'}$ with $(i', j') < (i, j)$.⁵ In Algorithm 2, we compute the $L_{ij}^{\mathcal{J}}$ in ascending lexicographic order via Cholesky factorization.

Notice that each $L^{\mathcal{J}}$ is fully defined by H and $L^{\mathcal{I}}$ via Algorithm 2, so we have proven the following lemma.

Lemma 3.22 ((Burer et al., 2002a, Lem. 1)). *For each $(H, L^{\mathcal{I}}) \in \mathbb{S}^n \times \mathbb{L}^n[\mathcal{I}]$ such that $L_{ii}^{\mathcal{I}} \neq 0$ for each $i \in \mathbb{I}_{1:n}$, there is a unique $L^{\mathcal{J}} \in \mathbb{L}^n[\mathcal{J}]$ satisfying (3.42).*

With a slight abuse of notation, we let $L^{\mathcal{J}} = L^{\mathcal{J}}(H, L^{\mathcal{I}})$ denote the function defined by Algorithm 2, which maps each $(H, L^{\mathcal{I}}) \in \mathbb{S}^n \times \mathbb{L}_{++}^n[\mathcal{I}]$ to the matrix $L^{\mathcal{J}} \in \mathbb{L}^n[\mathcal{J}]$ *uniquely*

⁵The lexicographic order $<$ on \mathbb{I}^2 is defined by $(i, j) < (i', j')$ if $i < i'$ or $(i = i') \wedge (j < j')$.

satisfying (3.42). Moreover, we let

$$Q(H, L^{\mathcal{I}}) := H + (L^{\mathcal{I}} + L^{\mathcal{J}}(H, L^{\mathcal{I}}))(L^{\mathcal{I}} + L^{\mathcal{J}}(H, L^{\mathcal{I}}))^{\top}$$

as in (3.41). Clearly $Q(H, L^{\mathcal{I}}) \succ H$ is satisfied by definition. Finally, we define the transformation

$$T(x, L^{\mathcal{I}}) := (x, Q(H(x), L^{\mathcal{I}})) \quad (3.43)$$

which has the inverse

$$T^{-1}(x, Q) := (x, \pi_{\mathcal{I}}^L[\text{chol}(Q - H(x))]) \quad (3.44)$$

and we have the following lemma.

Lemma 3.23 ((Burer et al., 2002a, Lem. 2)). *The function T defined by (3.43) is a bijection between $\mathbb{R}^m \times \mathbb{L}_{++}^n[\mathcal{I}]$ and $\text{int}(\mathcal{P})$.*

Differentiability of T and T^{-1} follow from differentiability of H and Algorithm 2. In fact, these functions are as smooth as H . More importantly, the bijection T allows us to transform the minimum of a continuous function over \mathcal{P} to an infimum over $\mathbb{R}^m \times \mathbb{L}_{++}^n[\mathcal{I}]$, given by the following theorem.

Theorem 3.24 ((Burer et al., 2002a, Thm. 1)). *If $f : \mathcal{P} \rightarrow \mathbb{R}$ is continuous and attains a minimum in \mathcal{P} , then*

$$\min_{(x, Q) \in \mathcal{P}} f(x, Q) = \inf_{(x, L^{\mathcal{I}}) \in \mathbb{R}^m \times \mathbb{L}_{++}^n[\mathcal{I}]} f(T(x, L^{\mathcal{I}})). \quad (3.45)$$

We reiterate the proof of Theorem 3.24 for illustrative purposes.

Proof. Continuity of f implies its minimum over \mathcal{P} equals its infimum over $\text{int}(\mathcal{P})$, i.e.,

$$\min_{(x,Q) \in \mathcal{P}} f(x, Q) = \inf_{(x,Q) \in \text{int}(\mathcal{P})} f(x, Q)$$

Since T is a bijection, we can transform the optimization variables as follows:

$$\inf_{(x,Q) \in \text{int}(\mathcal{P})} f(x, Q) = \inf_{(x,L^{\mathcal{I}}) \in T^{-1}(\text{int}(\mathcal{P}))} f(T(x, L^{\mathcal{I}})).$$

Finally, since $\mathbb{R}^m \times \mathbb{L}_{++}^n[\mathcal{I}] = T^{-1}(\text{int}(\mathcal{P}))$, we have (3.45). \square

3.5.2 Generalized Burer-Monteiro-Zhang method

We return to constraints of the form (3.25). Recall Assumption 3.13 requires the matrix inequalities are strictly feasible in the constraint set. We use a similar procedure to Section 3.5.1, but Algorithm 2 must be applied to *each* strict inequality $\Sigma \succ H$ and $\mathcal{A}(\beta, \Sigma) \succ 0$.

For the sparse symmetric matrix Σ and matrix inequality $\Sigma \succ H(\beta)$, the procedure is the same as in Section 3.5.1. Let $L^{\mathcal{J}\Sigma} = L^{\mathcal{J}\Sigma}(H, L^{\mathcal{I}\Sigma})$ denote the function defined by Algorithm 2 with $L^{\mathcal{I}} = L^{\mathcal{I}\Sigma}$, $\mathcal{I} = \mathcal{I}_{\Sigma}$, and $n = n_{\Sigma}$. Then

$$\Sigma(\beta, L^{\mathcal{I}\Sigma}) := H + (L^{\mathcal{I}\Sigma} + L^{\mathcal{J}\Sigma}(H, L^{\mathcal{I}\Sigma}))(L^{\mathcal{I}\Sigma} + L^{\mathcal{J}\Sigma}(H, L^{\mathcal{I}\Sigma}))^{\top}$$

guarantees $\Sigma(H, L^{\mathcal{I}\Sigma}) \succ H$ and $\Sigma(H, L^{\mathcal{I}\Sigma}) \in \mathbb{S}^{n_{\Sigma}}[\mathcal{I}_{\Sigma}]$ for all $(H, L^{\mathcal{I}\Sigma}) \in \mathbb{S}^{n_{\Sigma}} \times \mathbb{L}_{++}^{n_{\Sigma}}[\mathcal{I}_{\Sigma}]$. In other words, Σ is fully defined and the constraint $\Sigma \succ H$ automatically satisfied by $(H, L^{\mathcal{I}\Sigma}) \in \mathbb{S}^{n_{\Sigma}} \times \mathbb{L}_{++}^{n_{\Sigma}}[\mathcal{I}_{\Sigma}]$.

For the general matrix inequality $\mathcal{A}(\beta, \Sigma) \succeq 0$, the procedure is slightly different. Let $L^{\mathcal{J}\mathcal{A}} = L^{\mathcal{J}\mathcal{A}}(L^{\mathcal{I}\mathcal{A}})$ denote function defined by Algorithm 2 with $L^{\mathcal{I}} = L^{\mathcal{I}\mathcal{A}}$, $\mathcal{I} = \mathcal{I}_{\mathcal{A}}$, $n = n_{\mathcal{A}}$,

and $H = 0$. Define the functions

$$\mathcal{A}(L^{\mathcal{I}_A}) := (L^{\mathcal{I}_A} + L^{\mathcal{J}_A}(L^{\mathcal{I}_A}))(L^{\mathcal{I}_A} + L^{\mathcal{J}_A}(L^{\mathcal{I}_A}))^\top$$

which guarantees $\mathcal{A}(L^{\mathcal{I}_A}) \in \mathbb{S}_{++}^{n_A}[\mathcal{I}_A]$ for all $L^{\mathcal{I}_A} \in \mathbb{L}_{++}^{n_A}[\mathcal{I}_A]$. However, the constraint is not fully eliminated; we are left with $|\mathcal{I}_A|$ equality constraints in the transform space,

$$\text{vecs}_{\mathcal{I}_A}(\mathcal{A}(\beta, \Sigma(H(\beta), L^{\mathcal{I}_\Sigma})) - \mathcal{A}(L^{\mathcal{I}_A})) = 0$$

with the other $|\mathcal{L}^{n_A} \setminus \mathcal{I}_A|$ constraints automatically guaranteed by Algorithm 2.

To define the new constraints, we require the variable transformations

$$\mathcal{T}(\beta, L^{\mathcal{I}_\Sigma}, L^{\mathcal{I}_A}) := (\beta, \Sigma(H(\beta), L^{\mathcal{I}_\Sigma})) \quad (3.46a)$$

$$\mathcal{A}_{\mathcal{T}}(\beta, L^{\mathcal{I}_\Sigma}, L^{\mathcal{I}_A}) := \mathcal{A}(L^{\mathcal{I}_A}) \quad (3.46b)$$

which are well-defined for all $(\beta, L^{\mathcal{I}_\Sigma}, L^{\mathcal{I}_A}) \in \mathbb{R}^{n_\beta} \times \mathbb{L}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}^{n_A}[\mathcal{I}_A]$. With the functions (3.46), we define the transformed constraint functions $(g_{\mathcal{T}}, h_{\mathcal{T}})$ and the transformed constraint set $\Phi \subseteq \mathbb{R}^{n_\beta} \times \mathbb{L}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}^{n_A}[\mathcal{I}_A]$ according to (3.27) and (3.28). The inverse transform is

$$\mathcal{T}^{-1}(\beta, \Sigma) := (\beta, \pi_{\mathcal{I}_\Sigma}^L[\text{chol}(\Sigma - H(\beta))], \pi_{\mathcal{I}_A}^L[\text{chol}(\mathcal{A}(\beta, \Sigma))]) \quad (3.47)$$

for all $(\beta, \Sigma) \in \Theta_{++}$, and we have the following lemma.

Lemma 3.25. *The function \mathcal{T} defined by (3.46) is a bijection between Φ and Θ_{++} .*

Proof. First, we have $\mathcal{T}(\Phi) \subseteq \Theta_{++}$ since the transformed constraints guarantee the constraints $g(\beta, \Sigma) = 0$, $h(\beta, \Sigma) \leq 0$, $\Sigma \succ H(\beta)$, and $\mathcal{A}(\beta, \Sigma) \succ 0$ for any $(\beta, \Sigma) := \mathcal{T}(\phi)$ and $\phi \in \Phi$. Next, it is clear by construction that $\mathcal{T}^{-1} \circ \mathcal{T}$ is the identity map on Φ . Therefore \mathcal{T}

is injective. Similarly, we have $\mathcal{T}^{-1}(\Theta_{++}) \subseteq \Phi$ by construction, and $\mathcal{T} \circ \mathcal{T}^{-1}$ is the identity map on Θ_{++} , so $\mathcal{T} : \Phi$ is surjective. \square

Under Assumption 3.13, the functions \mathcal{T} , \mathcal{T}^{-1} , and $\mathcal{A}_{\mathcal{T}}$ are as smooth as H , and moreover, the bijection \mathcal{T} transforms a minimum over Θ into an infimum over Φ .

Proposition 3.26. *If Assumption 3.13 holds and $f : \Theta \rightarrow \mathbb{R}$ is continuous and attains a minimum in Θ , then*

$$\min_{\theta \in \Theta} f(\theta) = \inf_{\phi \in \Phi} f(\mathcal{T}(\phi)).$$

Proof. The proof follows that of Theorem 3.24, noting that Assumption 3.13 gives $\text{cl}(\Theta_{++}) = \Theta$ and therefore the minimum of f over Θ equals the infimum of f over Θ_{++} . \square

3.5.3 Approximate solutions

As mentioned in Section 3.3, we consider a lower bound $\varepsilon > 0$ on the diagonal elements of $(L^{\mathcal{I}_{\Sigma}}, L^{\mathcal{I}_{\mathcal{A}}})$. We define the tightened constraint set Φ_{ε} by (3.29). In the following theorem we show, under Assumption 3.13 and continuity of f , the infimum of $f \circ \mathcal{T}$ over Φ_{ε} converges to the minimum of f over Θ (see Appendix 3.A.5 for proof).

Theorem 3.27. *Suppose f is continuous and attains a minimum in Θ . Define*

$$\begin{aligned} \mu_0 &:= \min_{\theta \in \Theta} f(\theta) \\ \mu_{\varepsilon} &:= \inf_{\phi \in \Phi_{\varepsilon}} f(\mathcal{T}(\phi)). \end{aligned} \tag{3.48}$$

If Assumption 3.13 holds, then $\mu_{\varepsilon} \searrow \mu$ as $\varepsilon \searrow 0$.

In fact, with a few additional requirements on the objective f , convergence of approximate problem solutions to the solution of the original problem is guaranteed by the following theorem (see Appendix 3.A.5 for proof).

Theorem 3.28. *Suppose f is continuous and Assumption 3.13 holds. Consider the set-valued function $\hat{\theta} : \mathbb{R}_{\geq 0} \rightarrow \mathcal{P}(\Theta)$, defined as $\hat{\theta}_\varepsilon := \operatorname{argmin}_{\theta \in \mathcal{T}(\Phi_\varepsilon)} f(\theta)$ for all $\varepsilon > 0$, and $\hat{\theta}_0 := \operatorname{argmin}_{\theta \in \Theta} f(\theta)$. If there exists $\alpha \in \mathbb{R}$ and compact $C \subseteq \Theta$ such that*

$$\Theta_{f \leq \alpha} := \{ \theta \in \Theta \mid f(\theta) \leq \alpha \}$$

is contained in C and $\Theta_{f \leq \alpha} \cap \Theta_{++}$ is nonempty, then there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon_0 \in [0, \bar{\varepsilon})$,

- (a) *f achieves a minimum in Θ and $\hat{\theta}_0$ is nonempty;*
- (b) *if $\varepsilon_0 > 0$, then f achieves a minimum in $\mathcal{T}(\Theta_{\varepsilon_0})$ and $\hat{\theta}_{\varepsilon_0}$ is nonempty;*
- (c) *μ_ε is continuous and $\hat{\theta}_\varepsilon$ is outer semicontinuous at $\varepsilon = \varepsilon_0$; and*
- (d) *if $\hat{\theta}_0$ is a singleton, then $\limsup_{\varepsilon \searrow 0} \hat{\theta}_\varepsilon = \hat{\theta}_0$.*

Remark 3.29. Originally, Burer et al. (2002a) used a log-barrier approach to handle the strict inequalities implied by $L \in \mathbb{L}_{++}^n[\mathcal{I}]$ and achieve global convergence for a class of linear SDPs. For problems of the form (3.45), the log-barrier term eliminates all remaining constraints. However, for problems of the form (3.39) many constraints remain in addition to the strict inequalities on the diagonal elements of $(L^{\mathcal{I}_\Sigma}, L^{\mathcal{I}_\mathcal{A}}) \in \mathbb{L}_{++}^{n_\Sigma}[\mathcal{I}_\Sigma] \times \mathbb{L}_{++}^{n_\mathcal{A}}[\mathcal{I}_\mathcal{A}]$.

3.6 Summary and discussion

We stated a ML identification problem with model structure and constraints suitable for implementing offset-free controllers (Section 3.2). An algorithm for solving the ML identification problem with standard software was outlined (Section 3.3). Tightened LMI region constraints were introduced and shown to define closed sets of system matrices (Section 3.4, Theorem 3.20). A substitution and elimination scheme for approximating NSDPs as NLPs was presented (Section 3.5, Theorems 3.27 and 3.28). See Chapter 5 for real-world case studies of

these methods. We conclude with a discussion of computational issues, possible future applications and research directions, and an appendix of proofs that were deferred from the main chapter text.

Computational concerns The main limitation of eigenvalued-constrained ML is computational cost. While constrained ML retains linear scaling in sample size N , each LMI region constraint on an arbitrary system matrix $\tilde{A} \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$ requires an additional $O(\tilde{n}^2(m^2 + 1))$ variables and $O(\tilde{n}^2 m^2)$ equality constraints. These requirements can be significantly reduced for spectral abscissa bounds $\mathcal{D}_1(s)$ and stability constraints $\mathcal{D}_2(s, 0)$. As mentioned in Remark 3.18, these constraints are quite similar to the “smooth” spectral radii and abscissa constraints of Diehl et al. (2009); Vanbiervliet et al. (2009), which only add $O(\tilde{n}^2)$ variables and $O(\tilde{n}^2)$ equality constraints. For eigenvalues constrained to the LMI regions $\mathcal{D}_1(s)$ or $\mathcal{D}_2(s, x_0)$, implementing these constraints as a special case can reduce the computational cost significantly.

For a standard, black-box LADM (3.7) with $n_d = n_y$, a canonical form for (A_s, B_s, C_s) , and $(D, \hat{s}_0, \hat{d}_0) = (0, 0, 0)$, there are $O(n_s(n_u + n_y) + p^2)$ variables before constraints are added, and $O(n_s^2)$ variables after. Thus, fitting black-box models of large-scale systems is computationally prohibitive. However, as discussed in Section 3.2, large-scale chemical plants and networked systems may be represented by significantly fewer variables: $O(N_u n_u (p_u + m_u) + N_u p_u^2)$ without constraints, or $O(N_u n_u^2)$ with constraints, where N_u is the number of units or nodes, and n_u, m_u, p_u are the number of states, inputs, and outputs per unit or node.

Data-driven control The approach discussed so far is an indirect data-driven control design of offset-free MPC. A potential alternative is the direct data-driven control approach, where the control law is designed according to data (Berberich et al., 2021; Dorfler et al., 2022; Berberich et al., 2022a; Yuan and Cortés, 2022; Bianchin et al., 2023). The drawback of this ap-

proach is its reliance on Willem’s Fundamental Lemma (Willems et al., 2005), which assumes the data is generated from a plant of the model class and does not allow structured models. We also remark the models considered in this chapter have far more general noise models than those considered in direct data-driven control works.

Recent work on direct data-driven control has incorporated likelihood functions with measurement noise models into the control design (Yin et al., 2023). To the best of our knowledge, no current work has considered process noise, Kalman filter forms, or structuring the model with uncontrollable integrators for offset-free MPC. There is a future possibility of direct data-driven offset-free MPC design with both optimal control and estimation performance.

Linear identification of nonlinear systems The main difficulty of linear identification of nonlinear systems is plant-model mismatch. With ML identification, properties of the estimates are dependent on the plant’s stochastic behavior (Jr., 1982; White, 1984). For stationary, input-free models, the solution to the mismatched problem can be interpreted as (asymptotically) minimizing the Kullback-Leibler divergence between the power spectral densities of the model and plant (Anderson et al., 1978). However, there are still gaps in the treatment of inputs, state-space models, and arbitrary nonlinear plants. Moreover, there are no guarantees placed on this “closest” model, as it no longer aligns with the plant and therefore does not inherit any physically relevant properties from it. The constraints considered herein address concerns that the “closest” model may be unphysical. However, exactly what the closest model may entail is highly speculative and an area of future research.

Closed-loop performance monitoring Since ML identified models are more distributionally accurate, they are more suitable to the performance monitoring technique of Zagrobelny et al. (2013). Integrated identification and offset-free controller validation may be possible by

combining this method with ours. Another promising application is model re-identification, as mentioned in Chapter 1. From the asymptotic distribution, decision functions can be constructed to map the on-line MPC performance to a re-identification signal. These decision functions can be constructed so as to not alarm unless sufficiently exciting data is available. Decision-theoretic re-identification therefore has lower cost and risk compared to classic adaptive control or online reinforcement learning methods that require a persistently exciting identification signal. This approach could bring statistical data efficiency to the fields of adaptive control and online reinforcement learning.

Appendices

3.A Additional proofs

3.A.1 Proof of Proposition 3.1

Silverman (1976) presents complete characterization of the DARE solutions for regulation problems with cross terms. However, this admits additional nullspace terms into the gain matrix which the Kalman filtering problem does not allow. We avoid nullspace terms through the assumption $R_v \succ 0$ and therefore streamline the proof of Proposition 3.1.

For the following definitions and lemmas, we denote by $\mathcal{W} := (A, B, C, D)$ the system matrices corresponding to the noise-free system (2.1).

Definition 3.30. The system \mathcal{W} is *left invertible on* $\mathbb{I}_{0:k-1}$ if

$$0 = \begin{bmatrix} D & & & & \\ CB & D & & & \\ \vdots & \ddots & \ddots & & \\ CA^{k-2}B & \dots & CB & D & \end{bmatrix} \begin{bmatrix} u_0 \\ \vdots \\ u_{k-1} \end{bmatrix}$$

implies $u_0 = 0$. The system \mathcal{W} is *left invertible* if there is some $j \in \mathbb{I}_{>0}$ such that \mathcal{W} is left invertible on $\mathbb{I}_{0:k-1}$ for all $k \geq j$.

Definition 3.31. The system \mathcal{W} is *strongly detectable* if $y_k \rightarrow 0$ implies $x_k \rightarrow 0$.

The following lemmas are taken directly from (Silverman, 1976, Thms. 8, 18(iii)), but the proofs are omitted for the sake of brevity.

Lemma 3.32 ((Silverman, 1976, Thm. 8)). *If \mathcal{W} is left invertible, then \mathcal{W} is strongly detectable if and only if $(A - BF, C - DF)$ is detectable for all F of appropriate dimension.*

Lemma 3.33 ((Silverman, 1976, Thm. 18(iii))). *If \mathcal{W} is left invertible, then the DARE*

$$P = A^\top P A - (A^\top P B + C^\top D)(B^\top P B + D^\top D)^{-1}(B^\top P A + D^\top C)$$

has a unique, stabilizing solution⁶ if and only if \mathcal{W} is stabilizable and semistrongly detectable.

For the remainder of this section, we consider the full rank factorization

$$\begin{bmatrix} Q_w & S_{wv} \\ S_{wv}^\top & R_v \end{bmatrix} = \begin{bmatrix} \tilde{B} \\ \tilde{D} \end{bmatrix} \begin{bmatrix} \tilde{B}^\top & \tilde{D}^\top \end{bmatrix}$$

and the dual system $\tilde{\mathcal{W}} := (A^\top, C^\top, \tilde{B}^\top, \tilde{D}^\top)$ to analyze the properties of the original system (3.16). The following lemma relates the properties $R_v \succ 0$ and left invertability of $\tilde{\mathcal{W}}$.

Lemma 3.34. *If $R_v \succ 0$ then $\tilde{\mathcal{W}}$ is left invertible.*

Proof. Left invertability on $\mathbb{I}_{0:k-1}$ is equivalent to

$$0 = \begin{bmatrix} \tilde{D}^\top & & & & & & \\ & \tilde{B}^\top C^\top & & \tilde{D}^\top & & & \\ & \vdots & & \ddots & & \ddots & \\ \tilde{B}^\top (A^\top)^{k-2} C^\top & \dots & \tilde{B}^\top C^\top & & \tilde{D}^\top & & \end{bmatrix} \begin{bmatrix} u_0 \\ \vdots \\ u_{k-1} \end{bmatrix} \quad (3.49)$$

implying $u_0 = 0$. But $R_v = \tilde{D}\tilde{D}^\top \succ 0$, so \tilde{D}^\top has a zero nullspace. For each $k \in \mathbb{I}_{>0}$, the coefficient matrix of (3.49) has a zero nullspace. Thus, $u_0 = 0$ and $\tilde{\mathcal{W}}$ is left invertible. \square

⁶Contrary to in Section 3.2, here we mean the solution P is stabilizing when $A - BK(P)$ is stable, where $K(P) := (B^\top P B + D^\top D)^{-1} B^\top P$.

Finally, we can prove Proposition 3.1.

Proof of Proposition 3.1. By Lemma 3.34, we have that $\tilde{\mathcal{W}}$ is left invertible. Therefore, by Lemma 3.33, the DARE (3.20) has a unique, stabilizing solution if and only if $\tilde{\mathcal{W}}$ is stabilizable and strongly detectable. But by Lemma 3.32 and duality, the latter statement is true if and only if (A, C) is detectable and $(A - FC, \tilde{B} - F\tilde{D})$ is stabilizable for all $F \in \mathbb{R}^{n \times n_y}$. \square

3.A.2 Proof of Proposition 3.16

Throughout this appendix, we define the set of $n \times n$ Hermitian, Hermitian positive definite, and Hermitian positive semidefinite matrices as \mathbb{H}^n , \mathbb{H}_{++}^n , and \mathbb{H}_+^n . Notice that $f_{\mathcal{D}}$ maps to Hermitian matrices so we can write it as $f : \mathbb{C} \rightarrow \mathbb{H}^m$. We define the extension of $M_{\mathcal{D}}$ to complex arguments $M_{\mathcal{D}} : \mathbb{C}^{n \times n} \times \mathbb{H}_+^n \rightarrow \mathbb{H}^{nm}$ as

$$M_{\mathcal{D}}(A, P) := M_0 \otimes P + M_1 \otimes (AP) + M_1^{\top} \otimes (AP)^{\text{H}}.$$

To show Proposition 3.16, we need a preliminary result about Hermitian positive semidefinite matrices, generalized from Lemma A.1 of Chilali and Gahinet (1996).

Lemma 3.35. *For any $M \in \mathbb{H}^n$, if $M \succeq 0$ ($M \succ 0$) then $\text{Re}(M) \succeq 0$ ($\text{Re}(M) \succ 0$).*

Proof. With $M = \text{Re}(M) + \iota \text{Im}(M)$, it is clear M Hermitian implies $\text{Re}(M)$ is symmetric and $\text{Im}(M)$ is skew-symmetric. Thus $v^{\top} M v = v^{\top} \text{Re}(M) v$ for all $v \in \mathbb{R}^n$, and positive (semi)definiteness of M implies positive (semi)definiteness of $\text{Re}(M)$. \square

In proving Proposition 3.16, we take the approach of Chilali and Gahinet (1996) but are careful to distinguish eigenvalues on the interior \mathcal{D} from those on the boundary $\partial\mathcal{D}$.

Proof of Proposition 3.16. (\Leftarrow) Suppose that $M_{\mathcal{D}}(A, P) \succeq 0$ for some $P \succ 0$ and let $\lambda \in \lambda(A)$.

Then there exists a nonzero $v \in \mathbb{C}^n$ for which $v^H A = \lambda v^H$. Consider the identity

$$\begin{aligned}
(I_m \otimes v)^H M_{\mathcal{D}}(A, P)(I_m \otimes v) &= M_0 \otimes v^H P v + M_1 \otimes (v^H A P v) + M_1^T \otimes (v^H P A^T v) \\
&= M_0 \otimes v^H P v + M_1 \otimes (\bar{\lambda} v^H P v) + M_1^T \otimes (\lambda v^H P v) \\
&= v^H P v (M_0 + M_1 \lambda + M_1^T \bar{\lambda}) \\
&= v^H P v f_{\mathcal{D}}(\lambda).
\end{aligned}$$

The assumption $P \succ 0$ implies $v^H P v > 0$, and $M_{\mathcal{D}}(A, P) \succeq 0$ further implies $f_{\mathcal{D}}(\lambda) \succeq 0$. Therefore $\lambda \in \text{cl}(\mathcal{D})$.

Next suppose $\lambda \in \lambda(A)$ is non-simple and $\lambda \in \partial \mathcal{D}$. Then there exists nonzero $v_1, v_2 \in \mathbb{C}^n$ (linearly independent) such that $v^H f_{\mathcal{D}}(\lambda) v = 0$, $v_1^H A = \lambda v_1^H$, and $v_2^H A = \lambda v_2^H + v_1$. Because \mathcal{D} is open, $\lambda \in \partial \mathcal{D} = \text{cl}(\mathcal{D}) \setminus \mathcal{D}$ must satisfy both $f_{\mathcal{D}}(\lambda) \succeq 0$ and $f_{\mathcal{D}}(\lambda) \not\prec 0$. Therefore $f_{\mathcal{D}}(\lambda)$ is singular, and there exists a nonzero vector $v \in \mathbb{C}^m$ such that $v^H f_{\mathcal{D}}(\lambda) v = 0$. With the 2×2 matrices

$$\begin{aligned}
\tilde{P} &= \begin{bmatrix} p_{11} & p_{12} \\ p_{12} & p_{22} \end{bmatrix} := \begin{bmatrix} v_1^H \\ v_2^H \end{bmatrix} P \begin{bmatrix} v_1 & v_2 \end{bmatrix} \succ 0 \\
\tilde{J} &:= \lambda I_2 + \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}
\end{aligned}$$

we have $\begin{bmatrix} v_1 & v_2 \end{bmatrix}^H A = \tilde{J} \begin{bmatrix} v_1 & v_2 \end{bmatrix}^H$ and therefore

$$\begin{aligned}
(I_m \otimes \begin{bmatrix} v_1 & v_2 \end{bmatrix})^H M_{\mathcal{D}}(A, P)(I_m \otimes \begin{bmatrix} v_1 & v_2 \end{bmatrix}) &= M_0 \otimes \tilde{P} + M_1 \otimes \tilde{J} \tilde{P} + M_1^T \otimes (\tilde{J} \tilde{P})^T \\
&= M_{\mathcal{D}}(\tilde{J}, \tilde{P}) \succeq 0.
\end{aligned}$$

Next, we have

$$\begin{aligned}
 \tilde{M} &:= K_{2,m} M_{\mathcal{D}}(\tilde{J}, \tilde{P}) K_{2,m}^{\top} \\
 &= \tilde{P} \otimes M_0 + \tilde{J} \tilde{P} \otimes M_1 + (\tilde{J} \tilde{P})^{\top} \otimes M_1^{\top} \\
 &= \tilde{P} \otimes f_{\mathcal{D}}(\lambda) + \begin{bmatrix} p_{12}(M_1 + M_1^{\top}) & p_{22}M_1 \\ p_{22}M_1^{\top} & 0 \end{bmatrix} \succeq 0.
 \end{aligned}$$

Finally,

$$(I_2 \otimes v)^{\text{H}} \tilde{M} (I_2 \otimes v) = \begin{bmatrix} p_{12}v^{\text{H}}(M_1 + M_1^{\top})v & p_{22}v^{\text{H}}M_1v \\ p_{22}v^{\text{H}}M_1^{\top}v & 0 \end{bmatrix} \succeq 0.$$

But $\tilde{P} \succ 0$ implies $p_{22} > 0$, so the above matrix inequality implies $v^{\text{H}}M_1v = 0$. Moreover, with $v^{\text{H}}f_{\mathcal{D}}(\lambda)v = 0$, we also have $v^{\text{H}}M_0v = 0$ and therefore $f(z) \equiv 0$ and \mathcal{D} is empty, a contradiction. Therefore each $\lambda \in \lambda(A)$ non-simple implies $\lambda \in \mathcal{D}$.

(\Rightarrow) Suppose $\lambda(A) \subset \text{cl}(\mathcal{D})$ and $\lambda \in \lambda(A)$ non-simple implies $\lambda \in \mathcal{D}$.

If $A = \lambda$ is a (possibly complex) scalar, then it lies in $\text{cl}(\mathcal{D})$ by assumption, and therefore $M_{\mathcal{D}}(\lambda, p) = pf_{\mathcal{D}}(\lambda) \succeq 0$ for all $p > 0$.

If $A = \lambda I_n + N$ is a (possibly complex) Jordan block, where $N \in \mathbb{R}^{n \times n}$ is a shift matrix and $n > 1$, then $\lambda \in \mathcal{D}$ and $f_{\mathcal{D}}(\lambda) \succ 0$. Let $T_k := \text{diag}(k^{n-1}, \dots, k, 1)$ for each $k \in \mathbb{I}_{>0}$. Then $T_k^{-1}AT_k = \lambda I_n + k^{-1}N \rightarrow \lambda I_n$ as $k \rightarrow \infty$. Moreover, because $M_{\mathcal{D}}$ is continuous, we have

$$M_{\mathcal{D}}(T_k^{-1}AT_k, I_n) \rightarrow M_{\mathcal{D}}(\lambda I_n, I_n) = f_{\mathcal{D}}(\lambda) \otimes I_n \succ 0.$$

Therefore there exists some $k_0 \in \mathbb{I}_{>0}$ such that $M_{\mathcal{D}}(T_k^{-1}AT_k, I_n) \succ 0$ for all $k \geq k_0$. With

$P := T_k T_k^\top$, we have

$$\begin{aligned} M_{\mathcal{D}}(A, P) &= M_0 \otimes T_k T_k^\top + M_1 \otimes (A T_k T_k^\top) + M_1^\top \otimes (A T_k T_k^\top)^\top \\ &= (I_m \otimes T_k)(M_0 \otimes I_n + M_1 \otimes T_k^{-1} A T_k + M_1^\top \otimes (T_k^{-1} A T_k)^\top)(I_m \otimes T_k)^\top \\ &= (I_m \otimes T_k) M_{\mathcal{D}}(T_k^{-1} A T_k, I_n) (I_m \otimes T_k)^\top \succ 0. \end{aligned}$$

Finally, for any $A \in \mathbb{R}^{n \times n}$, let $A = V(\bigoplus_{i=1}^p J_i)V^{-1}$ denote the Jordan decomposition of A , where $J_i = \lambda_i I_{n_i} + N_i$, $\lambda_i \in \lambda(A)$, N_i are shift matrices, and $n = \sum_{i=1}^p n_i$. We have already shown that for each $i \in \mathbb{I}_{1:p}$, there exists $P_i \succ 0$ such that $M_{\mathcal{D}}(J_i, P_i) \succeq 0$. Then with $\tilde{P} := V(\bigoplus_{i=1}^p P_i)V^{-1}$, we have

$$\begin{aligned} &(I_m \otimes V^{-1}) M_{\mathcal{D}}(A, \tilde{P}) (I_m \otimes V^{-1})^H \\ &= M_0 \otimes \left(\bigoplus_{i=1}^p P_i \right) + M_1 \otimes \left(\bigoplus_{i=1}^p J_i P_i \right) + M_1 \otimes \left(\bigoplus_{i=1}^p J_i P_i \right)^\top \\ &= K_{n,m} \left(\bigoplus_{i=1}^p K_{m,n_i} M_{\mathcal{D}}(J_i, P_i) K_{m,n_i}^\top \right) K_{n,m}^\top \succeq 0 \end{aligned}$$

and therefore $M_{\mathcal{D}}(A, \tilde{P}) \succeq 0$. Last, Lemma 3.35 gives $M_{\mathcal{D}}(A, P) \succeq 0$ with $P := \text{Re}(\tilde{P})$ since

$$M_{\mathcal{D}}(A, P) = M_{\mathcal{D}}(A, \text{Re}(\tilde{P})) = \text{Re}(M_{\mathcal{D}}(A, \tilde{P})). \quad \square$$

3.A.3 Proof of Proposition 3.19

To show Proposition 3.19(a), we first require the following eigenvalue sensitivity result due to (Golub and Van Loan, 2013, Thm. 7.2.3).

Theorem 3.36 ((Golub and Van Loan, 2013, Thm. 7.2.3)). *For any $A \in \mathbb{C}^{n \times n}$, denote its Schur decomposition by $A = Q(D + N)Q^H$, where $Q \in \mathbb{C}^{n \times n}$ is unitary, $D \in \mathbb{C}^{n \times n}$ is diagonal, and*

$N \in \mathbb{C}^{n \times n}$ is strictly upper triangular.⁷ Let p be the smallest positive integer for which $M^p = 0$ where $M_{ij} := |N_{ij}|$. Then, for any $E \in \mathbb{R}^{n \times n}$ and $\mu \in \lambda(A + E)$,

$$\min_{\lambda \in \lambda(A)} |\mu - \lambda| \leq \max \{ c \|E\|, (c \|E\|)^{1/p} \}$$

where $c := \sum_{k=0}^{p-1} \|N\|^k$.

Proof of Proposition 3.19. Throughout this proof, we show a set S is not open (or not closed) by demonstrating that S^c (or S) does not contain all its limit points.

(a)—For any $A \in \mathbb{A}_{\mathcal{D}}^n$, continuity of $f_{\mathcal{D}}$ gives the existence of a function $\delta(\lambda) > 0$ such that $f_{\mathcal{D}}(z) \succ 0$ for all $|z - \lambda| < \delta(\lambda)$ and $\lambda \in \lambda(A)$. Let $\delta := \min_{\lambda \in \lambda(A)} \delta(\lambda)$. By Theorem 3.36 and norm equivalence, there exist $c > 0$ and $p \in \mathbb{I}_{1:n}$ such that

$$\max_{\mu \in \lambda(A+E)} \min_{\lambda \in \lambda(A)} |\lambda - \mu| \leq \max \{ c \|E\|_{\mathbb{F}}, (c \|E\|_{\mathbb{F}})^{1/p} \}$$

for all $E \in \mathbb{R}^{n \times n}$. Therefore there exists a $\varepsilon > 0$ such that

$$\max_{\mu \in \lambda(A+E)} \min_{\lambda \in \lambda(A)} |\lambda - \mu| < \delta$$

for all $E \in \mathcal{B} := \{ E' \in \mathbb{R}^{n \times n} \mid \|E'\|_{\mathbb{F}} < \varepsilon \}$. Finally, $A + \mathcal{B}$ is a neighborhood of A contained in $\mathbb{A}_{\mathcal{D}}^n$, and, since $A \in \mathbb{A}_{\mathcal{D}}^n$ was chosen arbitrarily, $\mathbb{A}_{\mathcal{D}}^n$ is open.

(b)(i)—Because \mathcal{D} is open, nonempty, and not equal to \mathcal{D} , $\partial\mathcal{D}$ is nonempty. Let $\lambda \in \partial\mathcal{D}$ and $\lambda_k \in \mathcal{D}^c$ be a sequence for which $\lambda_k \rightarrow \lambda$. By symmetry, we also have $\bar{\lambda} \in \mathcal{D}$ and $\bar{\lambda}_k \in \mathcal{D}^c$.

For $n = 2$, we have $A := \begin{bmatrix} \operatorname{Re}(\lambda) & -\operatorname{Im}(\lambda) \\ \operatorname{Im}(\lambda) & \operatorname{Re}(\lambda) \end{bmatrix} \in \mathbb{R}^{2 \times 2}$ has eigenvalues $\lambda, \bar{\lambda} \in \mathcal{D}$, and $A_k := \begin{bmatrix} \operatorname{Re}(\lambda_k) & -\operatorname{Im}(\lambda_k) \\ \operatorname{Im}(\lambda_k) & \operatorname{Re}(\lambda_k) \end{bmatrix} \in \mathbb{R}^{2 \times 2}$ has eigenvalues $\lambda_k, \bar{\lambda}_k \in \mathcal{D}^c$ for each $k \in \mathbb{I}_{>0}$. The corresponding

⁷A matrix U is strictly upper triangular if $U_{ij} = 0$ for all $i \geq j$.

eigenvectors are $\begin{bmatrix} \pm\iota \\ 1 \end{bmatrix} \in \mathbb{C}^2$. Therefore $A \in \tilde{\mathbb{A}}_{\mathcal{D}}^2$ but $A_k \in (\tilde{\mathbb{A}}_{\mathcal{D}}^2)^c$ for each $k \in \mathbb{I}_{>0}$, and the limit $A_k \rightarrow A$ gives us that $(\tilde{\mathbb{A}}_{\mathcal{D}}^2)^c$ does not contain all its limit points.

For $n > 2$, let $A_0 \in \tilde{\mathbb{A}}_{\mathcal{D}}^{n-2}$, and we can extend the prior argument with the sequence $B_k := A_k \oplus A_0 \in (\tilde{\mathbb{A}}_{\mathcal{D}}^n)^c$, $k \in \mathbb{I}_{>0}$ that converges to $B := A \oplus A_0 \in \tilde{\mathbb{A}}_{\mathcal{D}}^n$.

(b)(ii)—By part (b)(i), it suffices to consider the case $n = 1$. By closure and convexity of \mathcal{D} , $\mathcal{D} \cap \mathbb{R}$ is either a closed line segment, a closed ray, or \mathbb{R} itself. In other words, $\mathcal{D} \cap \mathbb{R}$ is open if and only if it has no endpoints. Moreover, since $\partial\mathcal{D} \cap \mathbb{R}$ is the set of the endpoints of $\mathcal{D} \cap \mathbb{R}$, $\mathcal{D} \cap \mathbb{R}$ is open if and only if $\partial\mathcal{D} \cap \mathbb{R}$ is empty. Finally, since $\tilde{\mathbb{A}}_{\mathcal{D}}^1 = \mathcal{D} \cap \mathbb{R}$, $\tilde{\mathbb{A}}_{\mathcal{D}}^1$ is open if and only if $\partial\mathcal{D} \cap \mathbb{R}$ is empty.

(c)(i)—Let $\lambda \in \partial\mathcal{D}$. Suppose $n = 4$. Then $\bar{\lambda} \in \partial\mathcal{D}$ by symmetry. Because \mathcal{D} is open, there exists a sequence $\lambda_k \in \mathcal{D}$ such that $\lambda_k \rightarrow \lambda$, and by symmetry, we also have $\bar{\lambda}_k \in \mathcal{D}$ and $\bar{\lambda}_k \rightarrow \bar{\lambda}$. Consider again the 2×2 matrices A and A_k from part (b)(i), which have eigenvalues $\lambda, \bar{\lambda} \in \mathcal{D}$ and $\lambda_k, \bar{\lambda}_k \in \mathcal{D}^c$, respectively. Then the block matrices $B := \begin{bmatrix} A & I_2 \\ 0 & A \end{bmatrix} \in \mathbb{R}^{4 \times 4}$ and $B_k := \begin{bmatrix} A_k & I_2 \\ 0 & A_k \end{bmatrix} \in \mathbb{R}^{4 \times 4}$ have the same eigenvalues, but this time the eigenvectors are $\begin{bmatrix} \pm\iota \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ \pm\iota \\ 1 \end{bmatrix} \in \mathbb{C}^4$ and the eigenvalues are non-simple. Since λ is a non-simple eigenvalue on the boundary of \mathcal{D} , we have $B \notin \tilde{\mathbb{A}}_{\mathcal{D}}^4$. However, λ_k are all in the interior of \mathcal{D} , so $B_k \in \tilde{\mathbb{A}}_{\mathcal{D}}^4$. Since $B_k \rightarrow B$, the set $\tilde{\mathbb{A}}_{\mathcal{D}}^4$ does not contain all its limit points.

On the other hand, let $\lambda \in \partial\mathcal{D}$ and suppose $n > 4$. Similarly to part (b)(i), with any $\tilde{A}_0 \in \tilde{\mathbb{A}}_{\mathcal{D}}^{n-4}$, we can extend the argument for the $n = 4$ case with the sequence $\tilde{A}_k := B_k \oplus \tilde{A}_0 \in \tilde{\mathbb{A}}_{\mathcal{D}}^n$, $k \in \mathbb{I}_{>0}$ that converges to $\tilde{A} := B \oplus \tilde{A}_0 \in (\tilde{\mathbb{A}}_{\mathcal{D}}^n)^c$.

(c)(ii)—Let $\lambda \in \partial\mathcal{D} \cap \mathbb{R}$ and $n \geq 2$. Because \mathcal{D} is convex, open, and nonempty, there exists $\varepsilon > 0$ such that exactly one of the real intervals $(\lambda, \lambda + \varepsilon)$ or $(\lambda - \varepsilon, \lambda)$ is contained in \mathcal{D} , whereas the other is contained in $\text{int}(\mathcal{D}^c)$. Without loss of generality, assume $(\lambda - \varepsilon, \lambda) \subseteq \mathcal{D}$.⁸ Then $A_k := (\lambda - \varepsilon/k)I_n + N_n \in \tilde{\mathbb{A}}_{\mathcal{D}}^n$ for each $k \in \mathbb{I}_{>0}$, but $A_k \rightarrow \lambda I_n + N_n \in (\tilde{\mathbb{A}}_{\mathcal{D}}^n)^c$ and therefore $\tilde{\mathbb{A}}_{\mathcal{D}}^n$ does not contain all its limit points.

⁸Otherwise, take the reflection about the imaginary axis $-\mathcal{D}$ and $-\tilde{\mathbb{A}}_{\mathcal{D}}^n$.

(d)—Since $\overline{\mathbb{A}}_{\mathcal{D}}^n := \{ A \in \mathbb{R}^{n \times n} \mid \lambda(A) \subset \text{cl}(\mathcal{D}) \}$ contains $\mathbb{A}_{\mathcal{D}}^n$, it suffices to show any $A \in \overline{\mathbb{A}}_{\mathcal{D}}^n$ is a limit point of $\mathbb{A}_{\mathcal{D}}^n$. Denote the Jordan form by $A = V (\bigoplus_{i=1}^p \mu_i I_{n_i} + N_{n_i}) V^{-1}$, where $V \in \mathbb{R}^{n \times n}$ is invertible, $\mu_i \in \lambda(A)$, $n = \sum_{i=1}^p n_i$, and $N_i \in \mathbb{R}^{n_i \times n_i}$ is a shift matrix. Because $\mu_i \in \text{cl}(\mathcal{D})$, there exists a sequence $\mu_{i,k} \in \mathcal{D}$ such that $\mu_{i,k} \rightarrow \mu_i$. Then $A_k := V (\bigoplus_{i=1}^p \mu_{i,k} I_{n_i} + N_i) V^{-1} \in \mathbb{A}_{\mathcal{D}}^n$ and $A_k \rightarrow A$. \square

3.A.4 Proof of Theorem 3.20

To prove Theorem 3.20(a,b), we use sensitivity results on the value functions of parameterized nonlinear SDPs,

$$V(y) := \inf_{x \in \mathbb{X}(y)} F(x, y) \quad (3.50)$$

where the set-valued function $\mathbb{X} : \mathbb{R}^m \rightarrow \mathcal{P}(\mathbb{R}^n)$ is defined by

$$\mathbb{X}(y) := \{ x \in \mathbb{R}^n \mid G(x, y) \succeq 0 \}.$$

Consider also the graph of the set-valued function \mathbb{X} ,

$$\mathbb{Z} := \{ (x, y) \in \mathbb{R}^{n+m} \mid G(x, y) \succeq 0 \}.$$

Notice that \mathbb{Z} is closed if G is continuous. We say Slater's condition holds at $y \in \mathbb{R}^m$ if there exists $x \in \mathbb{R}^n$ such that $x \in \text{int}(\mathbb{X}(y))$, or equivalently, $G(x, y) \succ 0$. In the following proposition, we specialize (Bonnans and Shapiro, 2000, Prop. 4.4) to nonlinear SDPs.

Proposition 3.37. *Let $y_0 \in \mathbb{R}^m$ and suppose*

- (i) *F and G are continuous on \mathbb{R}^{n+m} ;*
- (ii) *there exist $\alpha \in \mathbb{R}$ and compact $C \subset \mathbb{R}^n$ such that, for each y in a neighborhood of y_0 , the*

level set

$$\text{lev}_{\leq \alpha} F(\cdot, y) := \{ x \in \mathbb{X}(y) \mid F(x, y) \leq \alpha \}$$

is nonempty and contained in C ; and

(iii) Slater's condition holds at y_0 .

Then $F(\cdot, y)$ attains a minimum on $\mathbb{X}(y)$ for all $y \in N_y$, and $V(y)$ is continuous at $y = y_0$.

Proof. See Proposition 4.4 and the discussions in pp. 264, 483–484, 491–492 of Bonnans and Shapiro (2000). □

Finally, we prove Theorem 3.20 with Proposition 3.37.

Proof of Theorem 3.20. Let $\text{vec} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n^2}$ and $\text{vecs} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{(1/2)(n+1)n}$ denote the vectorization and symmetric vectorization operators, respectively.

(a)—Let $x := \text{vecs}(P)$, $y := \text{vec}(A)$, $F(x, y) := \text{tr}(VP)$, and $G(x, y) := P \oplus (M_{\mathcal{D}}(A, P) - M)$. We aim to use Proposition 3.37 to show the continuity of $\phi_{\mathcal{D}}$ on $\mathbb{A}_{\mathcal{D}}^n$. Let $A_0 \in \mathbb{A}_{\mathcal{D}}^n$. Condition (i) of Proposition 3.37 holds by assumption. Slater's condition (iii) holds because for any $P \succ 0$ such that $M_{\mathcal{D}}(A_0, P) \succ 0$, we can define $P_0 := \gamma P \succ 0$ for some $\gamma > \gamma_0 := \|M\| \times \|[M_{\mathcal{D}}(A_0, P)]^{-1}\|$ to give

$$M_{\mathcal{D}}(A_0, P_0) = \gamma M_{\mathcal{D}}(A_0, P) \succ \gamma_0 M_{\mathcal{D}}(A_0, P) \succeq M.$$

Moreover, by continuity of $M_{\mathcal{D}}$, there exists a neighborhood N_A of A_0 such that $M_{\mathcal{D}}(A, P_0) \succ M$ for all $A \in N_A$. Letting $\alpha := \text{tr}(VP_0) > 0$, we have that the set

$$\{ P \in \mathbb{S}_+^n \mid \text{tr}(VP) \leq \alpha \}$$

is compact and contains the nonempty level set

$$\{ P \in \mathbb{P}(A) \mid \text{tr}(VP) \leq \alpha \}$$

for all $A \in N_A$. Taking the image of each of the above sets under the `vecs` operation gives condition (ii) of Proposition 3.37. All the conditions of Proposition 3.37 are thus satisfied for each $A_0 \in \mathbb{A}_{\mathcal{D}}^n$, and we have $\phi_{\mathcal{D}}$ is continuous on $\mathbb{A}_{\mathcal{D}}^n$.

(b)—Continuity of $\phi_{\mathcal{D}}$ on $\mathbb{A}_{\mathcal{D}}^n$ implies closure of the sublevel sets of $\phi_{\mathcal{D}}$, and (3.38) follows by definition of $\mathbb{A}_{\mathcal{D}}^n(\varepsilon)$.

(c)—First, $M_{\mathcal{D}}(A, P) \succ 0$ implies $P \succ 0$ by Proposition 3.17. Moreover, for any $P \succ 0$ such that $M_{\mathcal{D}}(A, P) \succ 0$, we have $M_{\mathcal{D}}(A, P) \succeq \gamma M_{\mathcal{D}}(A, P) \succeq M$ with $P := \gamma P$ and $\gamma := \|M\| \times \|[M_{\mathcal{D}}(A, P)]^{-1}\|$, so feasibility of (3.22) is equivalent to feasibility of

$$M_{\mathcal{D}}(A, P) \succ M, \quad P \succeq 0$$

and therefore $\bigcup_{\varepsilon > 0} \mathbb{A}_{\mathcal{D}}^n(\varepsilon) = \mathbb{A}_{\mathcal{D}}^n$. But $\mathbb{A}_{\mathcal{D}}^n(\varepsilon)$ is monotonically decreasing,⁹ so $\mathbb{A}_{\mathcal{D}}^n(\varepsilon) \nearrow \bigcup_{\varepsilon > 0} \mathbb{A}_{\mathcal{D}}^n(\varepsilon) = \mathbb{A}_{\mathcal{D}}^n$ as $\varepsilon \searrow 0$. \square

3.A.5 Proof of Theorems 3.27 and 3.28

Starting with Theorem 3.27:

Proof of Theorem 3.27. Since μ_{ε} is nondecreasing and bounded from below by μ , it suffices to show that for each $\delta > 0$, there exists a $\bar{\varepsilon} > 0$ such that $\mu_{\bar{\varepsilon}} - \mu < \delta$.

Let $\theta^* \in \Theta$ denote a point for which $\mu = f(\theta^*)$. If $\theta^* \in \Theta_{++}$, we could simply choose $\bar{\varepsilon} > 0$ large enough to put θ^* in $\mathcal{T}(\Phi_{\bar{\varepsilon}})$ and achieve $\mu_{\bar{\varepsilon}} - \mu = 0 < \delta$.

⁹By “monotonically decreasing” we mean $\varepsilon \leq \varepsilon' \Rightarrow \mathbb{A}_{\mathcal{D}}^n(\varepsilon) \supseteq \mathbb{A}_{\mathcal{D}}^n(\varepsilon')$.

Instead, we assume $\theta^* \notin \Theta_{++}$. By Assumption 3.13, there exists a sequence $\theta_k \in \Theta_{++}$, $k \in \mathbb{I}_{>0}$ such that $\theta_k \rightarrow \theta$ as $k \rightarrow \infty$. Defining $\nu_k := f(\theta_k)$, we have $\nu_k \rightarrow \mu$ by continuity of f . Therefore, there exists some $k_0 \in \mathbb{I}_{>0}$ such that $\nu_k - \mu < \delta$ for all $k \geq k_0$. For each $\theta_k \in \Theta_{++}$, there exists a unique $\phi_k = (\beta_k, L_k^{\mathcal{I}_\Sigma}, L_k^{\mathcal{I}_A}) \in \Phi$ such that $\theta_k = \mathcal{T}(\phi_k)$ (by Lemma 3.25). Let $\bar{\varepsilon}$ be the minimum over all the diagonal elements of $L_{k_0}^{\mathcal{I}_\Sigma}$ and $L_{k_0}^{\mathcal{I}_A}$. Then $(\beta_{k_0}, L_{k_0}^{\mathcal{I}_\Sigma}, L_{k_0}^{\mathcal{I}_A}) \in \Phi_{\bar{\varepsilon}}$ by construction, $\nu_{k_0} \geq \mu_{\bar{\varepsilon}}$ by optimality, and $\mu_{\bar{\varepsilon}} - \mu \leq \nu_{k_0} - \mu < \delta$. \square

As in Appendix 3.A.4, we use sensitivity results of Bonnans and Shapiro (2000) on optimization problems to prove Theorem 3.28. This time, however, we consider the continuity of the value function for parameterized NLPs on Banach spaces. Let \mathcal{X} , \mathcal{Y} , and \mathcal{K} be Banach spaces and consider the parameterized NLP,

$$V(y) := \inf_{x \in \mathbb{X}(y)} F(x, y) \quad (3.51)$$

where the set-valued function $\mathbb{X} : \mathcal{Y} \rightarrow \mathcal{P}(\mathcal{X})$ is defined by

$$\mathbb{X}(y) := \{ x \in \mathcal{X} \mid G(x, y) \in K \}$$

for some $G : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{K}$ and $K \subseteq \mathcal{K}$ is closed. Let $X^0(y)$ denote the (possibly empty) set of solutions to (3.51). Define the graph of the set-valued function $\mathbb{X}(\cdot)$ by

$$\mathbb{Z} := \{ (x, y) \in \mathcal{X} \times \mathcal{Y} \mid G(x, y) \in K \}.$$

Notice that \mathbb{Z} is closed if G is continuous and K is closed.

Proposition 3.38 ((Bonnans and Shapiro, 2000, Prop. 4.4)). *Let $y_0 \in \mathcal{Y}$ and assume:*

- (i) *F and G are continuous on $\mathcal{X} \times \mathcal{Y}$ and K is closed;*
- (ii) *there exist $\alpha \in \mathbb{R}$ and a compact set $C \subseteq \mathcal{X}$ such that, for every y in a neighborhood of*

y_0 , the level set

$$\{x \in \mathbb{X}(y) \mid f(x, y) \leq \alpha\}$$

is nonempty and contained in C ; and

- (iii) for any neighborhood N_x of the solution set $X^0(y_0)$, there exists a neighborhood N_y of y_0 such that $N_x \cap \mathbb{X}(y)$ is nonempty for all $y \in N_y$;

then $V(y)$ is continuous and $X^0(y)$ is outer semicontinuous at $y = y_0$.

Proof of Theorem 3.28. First, we must specify $\bar{\varepsilon}$. For each $\theta \in \Theta_{++}$, let

$$\varepsilon(\theta) := \max \{ \varepsilon > 0 \mid \theta \in \mathcal{T}(\Phi_\varepsilon) \}$$

where the maximum is achieved since there is a finite number of diagonal elements of the Cholesky factors that must be lower bounded. Now we specify $\bar{\varepsilon}$ as the supremum of $\varepsilon(\theta)$ over all $\theta \in \Theta_{f \leq \alpha} \cap \Theta_{++}$,

$$\bar{\varepsilon} := \sup \{ \varepsilon(\theta) \mid \theta \in \Theta_{f \leq \alpha} \cap \Theta_{++} \}$$

so that, for any $\varepsilon \in (0, \bar{\varepsilon})$, $\Theta_{f \leq \alpha} \cap \mathcal{T}(\Phi_\varepsilon)$ is nonempty and is contained in the compact set C .

(a)—Following the proof of (Bonnans and Shapiro, 2000, Prop. 4.4), we have (i) F is continuous and (ii) the level set $\Theta_{f \leq \alpha}$ is nonempty and contained in the compact set C , which implies $\Theta_{f \leq \alpha}$ is a compact level set and therefore the minimum of f over $\Theta_{f \leq \alpha}$ is achieved and equals the minimum over Θ . Moreover, $\hat{\theta}_0$ must be nonempty.

(b)—Similarly to part (a), we have, for each $\varepsilon \in (0, \bar{\varepsilon})$, that the level set $\Theta_{f \leq \alpha} \cap \mathcal{T}(\Phi_\varepsilon)$ is nonempty and contained in the compact set C , so f achieves its minimum over $\mathcal{T}(\Phi_\varepsilon)$ and $\hat{\theta}_\varepsilon$ is nonempty.

(c)—Consider the graph of the constraint function,

$$\mathbb{Z} := \{ (\theta, \varepsilon) \in \Theta \times \mathbb{R}_{\geq 0} \mid \theta \in \mathcal{T}(\Phi_\varepsilon) \text{ if } \varepsilon > 0 \}.$$

Consider a sequence $(\theta_k, \varepsilon_k) \in \mathbb{Z}, k \in \mathbb{I}_{>0}$ that is convergent $(\theta_k, \varepsilon_k) \rightarrow (\theta, \varepsilon)$. Then $\varepsilon \geq 0$, otherwise the sequence would not converge. Moreover, $\theta \in \Theta$ since $\theta_k \in \mathcal{T}(\Phi_{\varepsilon_k}) \subseteq \Theta$ for all $k \in \mathbb{I}_{>0}$ and Θ contains all its limit points. If $\varepsilon = 0$, then $(\theta, \varepsilon) \in \mathbb{Z}$ trivially. On the other hand, if $\varepsilon > 0$, then $\varepsilon(\theta_k)$ converges to $\varepsilon(\theta)$ because \mathcal{T} is continuous and the max can be taken over a finite number of elements of $\mathcal{T}^{-1}(\theta_k)$. Moreover, $\varepsilon(\theta_k)$ and upper bounds ε_k because $\theta_k \in \mathcal{T}(\Phi_{\varepsilon_k})$, so $\varepsilon(\theta) \geq \varepsilon$. Finally, we have $\theta \in \mathcal{T}(\Phi_\varepsilon)$, $(\theta, \varepsilon) \in \mathbb{Z}$, and \mathbb{Z} is closed.

Let $\varepsilon_0 \geq 0$ and N_θ be a neighborhood of $\hat{\theta}_{\varepsilon_0}$. With

$$\delta := \sup \{ \varepsilon(\theta) \mid \theta \in N_\theta \} > 0$$

we have $N_\theta \cap \Theta$ and $N_\theta \cap \mathcal{T}(\Phi_\varepsilon)$ are nonempty for all $\varepsilon \in (0, \varepsilon_0 + \delta)$.

Finally, the requirements of Proposition 3.38 are satisfied for all $\varepsilon_0 \in [0, \bar{\varepsilon})$, so μ_ε is continuous and $\hat{\theta}_\varepsilon$ is outer semicontinuous at $\varepsilon = \varepsilon_0$.

(d)—The last statement follows by the definition of outer semicontinuity and the fact that the lim sup is nonempty. □

Chapter 4

Maximum likelihood estimation of disturbance models

Despite the strong foundations and attractive statistical properties of maximum likelihood (ML) identification, there are many reasons it may not be suitable to some applications. The first and most obvious shortcoming is the availability of an initial guess. The methods of Chapter 3 are optimization-based and therefore require a sensible initial guess to feed to the optimizer. Second, computation may be difficult. While we do not find ML identification computationally prohibitive in the case studies of Chapter 5, as discussed in Section 3.6, the optimization may scale poorly to large-scale systems. Third, for some applications, it is necessary to have a computationally lightweight framework with strong convergence guarantees, minimal library requirements (e.g., only linear algebra packages), and/or closed-form solutions.

To avoid nonlinear optimization, practitioners and theorists have long used regression-based identification methods such as ARX models or subspace identification. While these techniques are unsuitable for directly identifying the linear augmented disturbance model (LADM) (1.4), they may be augmented or modified to allow LADM identification. For practitioners, an attractive path from tuning to identification of disturbance models may be through such an augmented method.

In this chapter, we present methods for augmenting standard identification methods with disturbance modeling capabilities. The key prerequisite for such an augmentation is the availability of a state sequence estimate produced during the standard identification algorithm. From there, successive ML estimation problems can be solved to estimate a disturbance sequence and the noise covariance matrix. Throughout, we pose the subproblems with ML estimation, sometimes even as approximations of the ML identification problem (3.8), and focus on formulations that lead to closed-form solutions of the subproblems.

Additional notation The following additional notation is used throughout this chapter. For any signal $(a(k))_{k \in \mathbb{I}_{\geq 0}}$, we denote the length- n past and future horizons as

$$A_{-n}(k) := \begin{bmatrix} a(k-1) \\ \vdots \\ a(k-n) \end{bmatrix}, \quad A_n(k) := \begin{bmatrix} a(k) \\ \vdots \\ a(k+n-1) \end{bmatrix}$$

4.1 Literature review

Traditionally, MPC implementations have relied on linear finite impulse response (FIR) plant models (Qin and Badgwell, 2003; Darby and Nikolaou, 2012) with which an algorithm such as dynamic matrix control (DMC) (Cutler and Ramaker, 1980) or Identification and Command (IDCOM) (Richalet et al., 1978) is implemented. A few products, such as the Shell Multivariable Optimizing Controller (SMOC) (Marquis and Broustail, 1988; Yousfi and Tournier, 1991) and Adersa's predictive functional control (PFC) algorithm, rely solely on a linear state-space plant model. Darby and Nikolaou (2012) note that recent MPC products have shifted away from FIR models and towards linear state-space models. This shift is motivated by a number of shortcomings of the FIR approach, most notably: (1) the inability to handle unstable and integrating systems without modification, (2) the overparameterization of the underlying linear system (especially for slow processes), (3) the difficulty of formulating estimators,

and (4) the fact that FIR models are a special case of the linear state-space model (Lee et al., 1994; Lundström et al., 1995).

Other plant model formulations include autoregressive models (e.g., ARMA and CARIMA models) (Clarke et al., 1987a,b; Clarke, 1991; Sun et al., 2011) and transfer function models (Ljung, 1999). Both model types require complicated estimator formulations and their identification algorithms are typically formulated for single-input single-output (SISO) systems. As such, multi-input multi-output (MIMO) models are typically constructed from individually fit SISO models. Transfer function models must be realized as state-space models in order to formulate controller constraints. As with FIR models, every autoregressive and transfer function model can be realized as a state-space model (Ho and Kalman, 1966; Akaike, 1974).

To identify the plant model, practitioners typically fit a linear model to step response data, although it is also possible to linearize a physics-based plant model (Caveness and Downs, 2005; Rawlings et al., 2020). Neither approach provides the noise covariance estimates required to design an estimator for MPC implementation. Subspace methods—such as canonical variate analysis (CVA) (Larimore, 1983), N4SID (Van Overschee and De Moor, 1994), or MOESP (Verhaegen, 1994)—can be used to identify estimate the process and measurement noise covariances, but they are unsuitable for structuring disturbance model with uncontrollable integrating modes (Muske and Badgwell, 2002; Pannocchia and Rawlings, 2003). Disturbance models may be tuned under strong assumptions on the process and measurement noises (Lee et al., 1994; Lee and Yu, 1994), but the required assumptions are not general, producing suboptimal estimator performance. Autocovariance least squares (ALS) can identify the complete disturbance model, but it does not identify the plant model (Odelson et al., 2006). Additionally, there is a trade-off between the computational complexity of ALS and the variance of the ALS estimates because the optimal least squares weighting matrix is a function of the covariances to be estimated (Rajamani and Rawlings, 2009; Zagrobelny and Rawlings,

2015; Arnold and Rawlings, 2022). Kuntz and Rawlings (2022) presented the first identification algorithm that provides estimates of both the state-space model coefficients and the disturbance noise covariance required to implement an offset-free MPC.

Most of the MPC deployment cost is incurred during plant identification due to the commonality of open-loop identification experiments, where product quality is difficult if not impossible to maintain, and the process must be perturbed from the optimal operating point in order to acquire quality data (Canney, 2003; Zhu, 2006). As a result, closed-loop identification experiments are an opportunity for significant safety and profitability improvements in chemical process control. Closed-loop identification experiments can then be conducted online, at and around the optimal operating point, negating the cost of opening the loop to perform the experiment. New MPCs can be implemented on processes controlled with other methods (PID, DMC, etc.) and existing MPCs be significantly improved with re-identified models. Closed-loop experiments can be conducted via setpoint perturbations that are more predictable and reliable than open-loop input perturbations. Moreover, the control loop is never broken, so the MPC is always enforcing constraints throughout the experiment.

Canney (2003) points out that MPC performance decays over time after deployment, and proposes MPC upkeep be a continuous process of algorithm improvement, where the model, MPC tuning, and organizational details are adjusted as necessary. A closed-loop disturbance model identification method can be applied to continuous offset-free MPC monitoring and upkeep. Previous attempts at continuous MPC monitoring and upkeep simply attempt to detect (and sometimes diagnose the source of) plant-model mismatch (Harrison and Qin, 2009; Pannocchia and De Luca, 2012; Kheradmandi and Mhaskar, 2018). However these algorithms rely on heuristic cutoffs for the alarm thresholds because they are based on LTI system order estimation. With the full set of parameter estimates, there is a future possibility of advanced offset-free MPC monitoring schemes with rigorous performance guarantees.

Closed-loop experimentation requires an existing controller, meaning open-loop experi-

ments for MPC design or PID tuning are still necessary. To this end, we suggest suboptimal but safe experiments be done using traditional step-response designs, or loops be initially closed with PID methods. While the algorithm proposed herein and in Kuntz and Rawlings (2022) will still handle open-loop step responses. At a later date, a closed-loop identification experiment may be run to refine and re-identify the model. The only advantage of open-loop methods (Section 4.2) are their relative simplicity compared to closed-loop methods (Section 4.3).

4.2 Simple methods for disturbance model identification

We modify the models (1.3) and (1.4) slightly to provide a better match with industrial practice. First, we consider linear Gaussian state-space (LGSS) models without feedthrough terms or cross-covariances,

$$x^+ = Ax + Bu + w \quad (4.1a)$$

$$y = Cx + v \quad (4.1b)$$

$$\begin{bmatrix} w \\ v \end{bmatrix} \stackrel{\text{iid}}{\sim} \text{N} \left(0, \begin{bmatrix} Q_w & \\ & R_v \end{bmatrix} \right) \quad (4.1c)$$

rather than the general model (1.3). Similarly, we neglect the feedthrough term in the linear augmented disturbance model (LADM),

$$x^+ = Ax + Bu + B_d d + w \quad (4.2a)$$

$$d^+ = d + w_d \quad (4.2b)$$

$$y = Cx + C_d d + v \quad (4.2c)$$

$$\begin{bmatrix} w \\ w_d \\ v \end{bmatrix} \stackrel{\text{iid}}{\sim} \text{N}(0, S_d). \quad (4.2d)$$

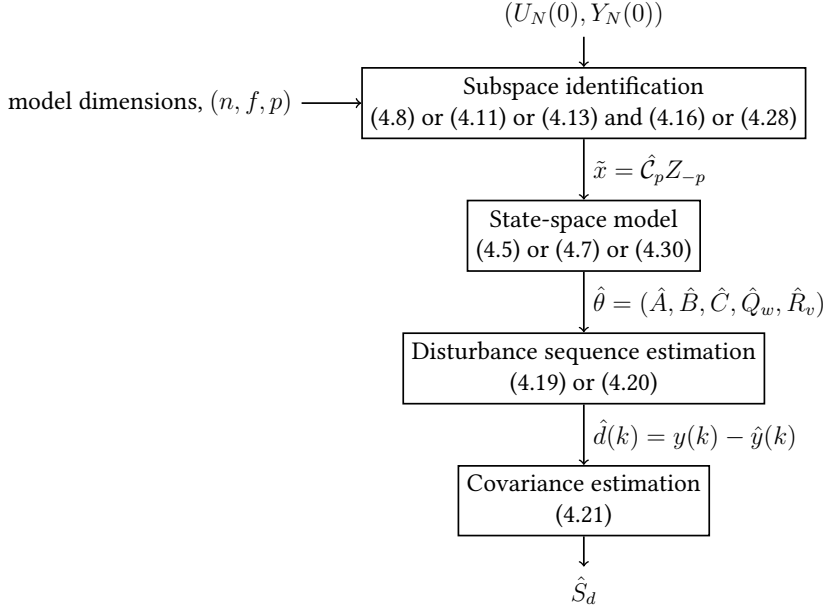


Figure 4.1: Outline of the disturbance model identification method.

The goal of the identification algorithm is to estimate the parameters of the LADM (4.2) from only input-output data (u, y) . An outline of the algorithm is given in Figure 4.1. We split the algorithm outline into two main parts: (i) standard state-space identification methods and (ii) disturbance model augmentation. In this section, we focus on simple algorithms with closed-form solutions. More a more general algorithm that works with closed-loop data and provides a degree of statistical efficiency is given in Section 4.3.

4.2.1 Standard identification

First, we consider estimation of the parameters $\theta = (A, B, C, Q_w, R_v)$ of the model (4.1). Throughout this process, a focus is given to the estimation of a state sequence $\tilde{X}_{N-p+1}(p)$ (where p is to be defined). If one were available, we could pose the ML problem corresponding to the *joint density* in $(X_{N_s}(p+1), Y_{N_s}(p))$,

$$\max_{\theta} L_N^{\text{SS}}(X_{N_s+1}(p), \theta) := \ln p(X_{N_s}(p+1), Y_{N_s}(p) | x(p), U_{N_s}(p), \theta) \quad (4.3)$$

using the estimated state sequence in place of $X_{N_s+1}(p)$, where $N_s := N - p$.

For the general LGSS model (1.3), we have the log-likelihood function

$$L_N^{\text{SS}}(\tilde{X}_{N_s+1}(p), \theta) \propto -\frac{N_s}{2} \ln \det \Sigma - \frac{1}{2} \sum_{k=p}^{N-1} |s(k) - \Theta t(k)|_{\Sigma^{-1}}^2. \quad (4.4)$$

where $s := \begin{bmatrix} \tilde{x}^+ \\ y \end{bmatrix}$, $t := \begin{bmatrix} \tilde{x} \\ u \end{bmatrix}$, $\Theta := \begin{bmatrix} A & B \\ C & D \end{bmatrix}$, and $\Sigma := \begin{bmatrix} Q_w & S_{wv} \\ S_{wv}^\top & R_v \end{bmatrix}$. Unique solutions to (4.3) are given by

$$\begin{bmatrix} \hat{A} & \hat{B} \\ \hat{C} & \hat{D} \end{bmatrix} = \left(\sum_{k=p}^{N-1} s(k)[t(k)]^\top \right) \left(\sum_{k=p}^{N-1} t(k)[t(k)]^\top \right)^{-1} \quad (4.5a)$$

$$\begin{bmatrix} \hat{Q}_w & \hat{S}_{wv} \\ \hat{S}_{wv}^\top & \hat{R}_v \end{bmatrix} = \frac{1}{N_s} \sum_{k=p}^{N-1} (s(k) - \hat{\Theta}t(k))(s(k) - \hat{\Theta}t(k))^\top \quad (4.5b)$$

so long as the inverse exists (Anderson, 2003, Thm. 8.2.1). For the special case of (4.1), we have the modified log-likelihood function,

$$L_N^{\text{SS}}(\tilde{X}_{N_s+1}(p), \theta) \propto -\frac{N_s}{2} (\ln \det Q_w + \ln \det R_v) - \frac{1}{2} \sum_{k=p}^{N-1} \left[|\tilde{x}(k+1) - A\tilde{x}(k) - Bu(k)|_{Q_w^{-1}}^2 + |y(k) - C\tilde{x}(k)|_{R_v^{-1}}^2 \right]. \quad (4.6)$$

For the log-likelihood (4.4), the problem (4.3) can be separated into (A, B, Q_w) and (C, R_v) subproblems, with unique solutions given by

$$[\hat{A} \quad \hat{B}] = \left(\sum_{k=p}^{N-1} \tilde{x}(k+1)[t(k)]^\top \right) \left(\sum_{k=p}^{N-1} t(k)[t(k)]^\top \right)^{-1} \quad (4.7a)$$

$$\hat{C} = \left(\sum_{k=p}^{N-1} y(k)[\tilde{x}(k)]^\top \right) \left(\sum_{k=p}^{N-1} \tilde{x}(k)[\tilde{x}(k)]^\top \right)^{-1} \quad (4.7b)$$

$$\hat{Q}_w = \frac{1}{N_s} \sum_{k=p}^{N-1} (\tilde{x}(k+1) - \hat{A}\tilde{x}(k) - \hat{B}u(k))(\tilde{x}(k+1) - \hat{A}\tilde{x}(k) - \hat{B}u(k))^\top \quad (4.7c)$$

$$\hat{R}_v = \frac{1}{N_s} \sum_{k=p}^{N-1} (y(k) - \hat{C}\tilde{x}(k))(y(k) - \hat{C}\tilde{x}(k))^\top \quad (4.7d)$$

so long as the inverses exist (Anderson, 2003, Thm. 8.2.1).

The ML problem (4.3) and estimates (4.5) and (4.7) accomplish the second stage in Figure 4.1. The remainder of this subsection is aimed at accomplishing the first stage through simple methods.

States as past inputs and outputs

The most straightforward option is to package past histories of input-output data into the state, i.e.,

$$\tilde{x} := \begin{bmatrix} Y_{-p_y} \\ U_{-p_u} \end{bmatrix} \quad (4.8)$$

where $p_y, p_u \in \mathbb{I}_{\geq 0}$ and $p := \min\{p_y, p_u\}$. While this is a straightforward option, it may drastically inflate the state dimension. Moreover, this approach is equivalent to posing the following autoregressive model:

$$y(k) = \sum_{i=1}^{p_y} H_i y(k-i) + \sum_{j=1}^{p_u} G_j u(k-j) + v(k), \quad v(k) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, R_v). \quad (4.9)$$

We can rewrite (4.9) in the form of (1.3) with

$$A := \begin{bmatrix} H_1 & \dots & H_{p_y-1} & H_{p_y} & G_1 & \dots & G_{p_u-1} & G_{p_u} \\ I_{n_y} & & & & & & & \\ & \ddots & & & & & & \\ & & I_{n_y} & 0_{n_y \times n_y} & 0 & & & \\ & & & & I_{n_u} & & & \\ & & & & & \ddots & & \\ & & & & & & I_{n_u} & 0 \end{bmatrix}, \quad B := \begin{bmatrix} 0_{n_y \times n_u} \\ 0_{n_y \times n_u} \\ \vdots \\ 0_{n_y \times n_u} \\ I_{n_u} \\ 0_{n_u \times n_u} \\ \vdots \\ 0_{n_u \times n_u} \end{bmatrix},$$

$$C := [H_1 \dots H_{p_y-1} H_{p_y} G_1 \dots G_{p_u-1} G_{p_u}], \quad D := 0,$$

$$Q_w := \begin{bmatrix} R_v & \\ & 0 \end{bmatrix}, \quad S_{wv} := \begin{bmatrix} R_v \\ 0 \end{bmatrix}.$$

The past-history state (4.8) creates rank-deficient covariance matrices,

$$\begin{bmatrix} Q_w & S_{wv} \\ S_{wv}^\top & R_v \end{bmatrix} = \begin{bmatrix} R_v & 0 & R_v \\ 0 & 0 & 0 \\ R_v & 0 & R_v \end{bmatrix}$$

and therefore the inverses in (4.5) will never exist. The state-space ML problem can still be solved by directly posing the ML estimation of (4.9), or by using the arguments in Kuntz and Rawlings (2023b). However, it is worth reducing the state dimension to avoid unnecessary computation time and ill-conditioning in the controller and estimator.

State reduction with singular value decomposition

In lieu of the states, we can use the Kalman filter state estimates, which are represented by past input-output data. There exists a steady-state Kalman gain K and innovation error covariance R_e such that $A_K := A - KC$ is stable and

$$\hat{x}^+ = A_K \hat{x} + B_K z \tag{4.10a}$$

$$e := y - C\hat{x} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, R_e) \tag{4.10b}$$

where $B_K := \begin{bmatrix} B & K \end{bmatrix}$, $\hat{x} \in \mathbb{R}^n$ are the state estimates, and $z := \begin{bmatrix} u^\top & y^\top \end{bmatrix}^\top$ is the combined input-output data (Kwakernaak and Sivan, 1972; Hespanha, 2018). Given any $n \leq p \ll N$ chosen large enough so that $A_K^p \approx 0$, we can recursively solve the Kalman predictor (4.10) to write the state as follows,

$$\hat{x}(k) = A_K^p \hat{x}(k-p) + C_p Z_{-p}(k) \approx C_p Z_{-p}(k)$$

where $\mathcal{C}_p := \begin{bmatrix} B_K & A_K B_K & \dots & A_K^{p-1} B_K \end{bmatrix}$. Therefore, estimating \mathcal{C}_p also provides state estimates $\tilde{x} := \mathcal{C}_p Z_{-p}$ for estimation of the parameters θ in the model (4.1).

It is fairly straightforward to construct a matrix \mathcal{C}_p from the data using the singular value decomposition (SVD). Take an economic SVD of the data matrix,

$$H = [Z_{-p}(p) \quad \dots \quad Z_{-p}(N)] \approx U_1 S_1 V_1^\top. \quad (4.11a)$$

where $U_1 \in \mathbb{R}^{n_y p \times n}$ and $V_1 \in \mathbb{R}^{N_s \times n}$ have orthogonal columns and $S_1 \in \mathbb{R}^{n \times n}$ is diagonal with positive diagonal elements. We can define our states as

$$\tilde{x}(k) := U_1^\top Z_{-p}(k) \quad (4.11b)$$

and therefore $Z_{-p}(k) = U_1 \tilde{x}(k)$. But this means $\tilde{x}(k) \approx T \hat{x}(k)$ where $T := \mathcal{C}_p U_1 \in \mathbb{R}^{n \times n}$ is invertible because it is the product of full row and column rank matrices. Thus, our state $\tilde{x}(k)$ is approximately equal to the state estimates, up to a similarity transformation.

Ho-Kalman algorithm

An alternative to the direct ARX formulation (4.9) is to use the so-called Ho-Kalman algorithm to find an estimate of $\hat{\mathcal{C}}_p$. Choose $n < f, p \ll N$ and suppose we have access to the first $\bar{p} := f + p - 1$ Markov parameters $G_i := C A_K^{i-1} B_K, i \in \mathbb{I}_{1:\bar{p}}$. Then we can write

$$\mathcal{H}_{f,p} := \begin{bmatrix} G_1 & \dots & G_p \\ \vdots & & \vdots \\ G_f & \dots & G_{f+p-1} \end{bmatrix} = \mathcal{O}_f \mathcal{C}_p, \quad \mathcal{O}_f := \begin{bmatrix} C \\ C A_K \\ \vdots \\ C A_K^{f-1} \end{bmatrix}. \quad (4.12)$$

Notice that, if the model (4.1) is minimal, the block Hankel matrix $\mathcal{H}_{f,p}$ has rank n . Taking an (exact) economic SVD of $\mathcal{H}_{f,p}$,

$$\mathcal{H}_{f,p} = U_1 S_1 V_1^\top \quad (4.13a)$$

where $U_1 \in \mathbb{R}^{n_y p \times n}$ and $V_1 \in \mathbb{R}^{(n_u+n_y)f \times n}$ have orthogonal columns and $S_1 \in \mathbb{R}^{n \times n}$ is diagonal with positive diagonal elements, we can realize \mathcal{C}_p and \mathcal{O}_f as

$$\mathcal{O}_f := U_1 S_1^{1/2}, \quad \mathcal{C}_p := S_1^{1/2} V_1^\top. \quad (4.13b)$$

To estimate the Markov parameters $G_i, i \in \mathbb{I}_{1:\bar{p}}$, we use the Kalman predictor form (4.10) to write the following ARX model:

$$y(k) = CA_K^{\bar{p}} \hat{x}(k - \bar{p}) + CC_{\bar{p}} Z_{-\bar{p}}(k) + e(k) \approx CC_{\bar{p}} Z_{-\bar{p}}(k) + e(k) \quad (4.14)$$

where the coefficient matrix contains the first \bar{p} Markov parameters,

$$CC_{\bar{p}} = [G_1 \quad \dots \quad G_{\bar{p}}].$$

The ML estimation problem for the ARX model parameters is

$$\max_{CC_{\bar{p}}, R_e} L_N^{\text{ARX}}(CC_{\bar{p}}, R_e) \quad (4.15)$$

where the log-likelihood function is given by

$$\begin{aligned} L_N^{\text{ARX}}(CC_{\bar{p}}, R_e) &:= \ln p(Y_{N_s}(\bar{p}) | U_{N_s}(\bar{p}), CC_{\bar{p}}, R_e) \\ &\propto -\frac{N_s}{2} \ln \det R_v - \frac{1}{2} \sum_{k=\bar{p}}^{N_s-1} |y(k) - CC_p Z_{-p}(k)|_{R_e^{-1}}. \end{aligned}$$

The unique solution to (4.15) is given by

$$\widehat{\mathcal{C}\mathcal{C}}_{\bar{p}} = \left(\sum_{k=p}^{N-1} y(k)[Z_{-\bar{p}}(k)]^\top \right) \left(\sum_{k=p}^{N-1} Z_{\bar{p}}(k)[Z_{-\bar{p}}(k)]^\top \right)^{-1}, \quad (4.16a)$$

$$\widehat{R}_e = \sum_{k=p}^{N-1} (y(k) - \widehat{\mathcal{C}\mathcal{C}}_{\bar{p}} Z_{-\bar{p}}(k))(y(k) - \widehat{\mathcal{C}\mathcal{C}}_{\bar{p}} Z_{-\bar{p}}(k))^\top \quad (4.16b)$$

so long as the inverse exists (Anderson, 2003, Thm. 8.2.1). Then we can substitute

$$\widehat{\mathcal{C}\mathcal{C}}_{\bar{p}} = [\hat{G}_1 \quad \dots \quad \hat{G}_{f+p-1}]$$

into (4.12) and compute $\hat{\mathcal{C}}_p$ via the SVD (4.13). This time, due to noise in the estimates, the SVD (4.13) will not be exact. However, the estimates are still, in some sense, “robust” to noise, as shown by Oymak and Ozay (2022).

Selection of the model dimensions (n, f, p) can either be tuned by hand or with information criteria methods. While dimension selection is outside of the scope of this chapter, Bauer (2001); Chiuso (2010); Larimore (2005) each describe selection of the parameters n , f , and p , respectively. In the case studies, we tuned (n, f, p) by hand and validated the chosen state order n with the singular value criterion described by Bauer (2001).

4.2.2 Disturbance model identification

We now augment the standard LGSS model (4.1) with a disturbance model, given estimates of the state sequence $\tilde{X}_{N_s+1}(p)$ and parameters $\hat{\theta} = (\hat{A}, \hat{B}, \hat{C}, \hat{Q}_w, \hat{R}_v)$ of the LGSS model (4.1). This is done by first estimating a disturbance sequence that captures the most long-term modeling error, and then re-estimating the noise covariances based on that disturbance sequence. The shaping matrices (B_d, C_d) of the noise model are inconsequential to the algorithm, except that they must obey the offset-free rank condition (2.22), and that output

disturbance models turn out to be computationally advantageous.

Estimating the disturbance sequence

Given a model of the form (4.1), a disturbance model (B_d, C_d) , and a state sequence $(\tilde{x}(k))$, we treat the disturbance sequence $(d(k))$ as accounting for the long-range model errors. We can write the long-range output at time k as

$$\begin{aligned} y(k) &= \hat{C}\hat{A}^{k-p}\tilde{x}(p) + \sum_{j=p}^{k-1} \hat{C}\hat{A}^{k-j-1}\hat{B}u(j) \\ &\quad + \sum_{j=p}^{k-1} \hat{C}\hat{A}^{k-j-1}(B_d d(j) + w(j)) + C_d d(k) + v(k) \end{aligned}$$

and the *predicted* long-range output at time k is

$$\hat{y}(k) := \hat{C}\hat{A}^{k-p}\tilde{x}(p) + \sum_{j=p}^{k-1} \hat{C}\hat{A}^{k-j-1}\hat{B}u(j). \quad (4.17)$$

Next, we define the long-range prediction error as $z(k) := y(k) - \hat{y}(k)$, which gives

$$z(k) = \sum_{j=p}^{k-1} \hat{C}\hat{A}^{k-j-1}(B_d d(j) + w(j)) + C_d d(k) + v(k).$$

Rewriting this as a linear model,

$$Z_{N_s}(p) = \mathcal{A}D_{N_s}(p) + \mathcal{B}W_{N_s}(p) + V_{N_s}(p) \quad (4.18a)$$

$$\mathcal{B}W_{N_s}(p) + V_{N_s}(p) \sim \mathcal{N}(0, \mathcal{V}) \quad (4.18b)$$

where

$$\mathcal{A} := \begin{bmatrix} C_d & & & & \\ \hat{C}B_d & C_d & & & \\ \vdots & \ddots & \ddots & & \\ \hat{C}\hat{A}^{N-2}B_d & \dots & \hat{C}B_d & C_d & \end{bmatrix}, \quad \mathcal{B} := \begin{bmatrix} 0 & & & & \\ B_1 & 0 & & & \\ \vdots & \ddots & \ddots & & \\ B_{N-1} & \dots & B_1 & 0 & \end{bmatrix},$$

$$B_j := \hat{C}\hat{A}^{j-1} \quad \forall j \geq 1, \quad \mathcal{V} := \mathcal{B}(I \otimes Q_w)\mathcal{B}^\top + I \otimes R_v.$$

The model (4.18) has the ML estimate (Rao, 1971; Magnus and Neudecker, 2019, p. 313),

$$\hat{D}_{N_s}(p) = (\mathcal{A}^\top \mathcal{V}_0^\dagger \mathcal{A})^\dagger \mathcal{A}^\top \mathcal{V}_0^\dagger Z_{N_s}(p) \quad (4.19)$$

where $\mathcal{V}_0 := \mathcal{V} + \mathcal{A}\mathcal{A}^\top$. This is an $O(N^3)$ computation with $O(N^2)$ memory requirements. Notice that when $B_d = 0$ and $C_d = I$, we have $\mathcal{A} = I$, $\mathcal{V}_0 = \mathcal{V} + I$ invertible, and

$$(\mathcal{A}^\top \mathcal{V}_0^\dagger \mathcal{A})^\dagger \mathcal{A}^\top \mathcal{V}_0^\dagger = \mathcal{V}_0 \mathcal{V}_0^{-1} = I.$$

Thus the disturbance estimates (4.19) are equivalently written

$$\hat{d}(k) = z(k) \quad (4.20)$$

which is an $O(N)$ computation without additional memory requirements. It is clear that whenever the system is free of integrators, the simplified solution (4.20) is computationally advantageous. A similarity transformation can be used to find the desired disturbance model after the output disturbance model is found (Rajamani et al., 2009).

Estimating the noise covariances

Given the estimated states and disturbances, one can stack the equations of the model (4.2) to write a simple covariance estimation problem,

$$\tilde{e}(k) := \begin{bmatrix} \tilde{x}(k+1) \\ \hat{d}(k+1) \\ y(k) \end{bmatrix} - \begin{bmatrix} \hat{A} & B_d & \hat{B} \\ 0 & I & 0 \\ \hat{C} & C_d & 0 \end{bmatrix} \begin{bmatrix} \tilde{x}(k) \\ \hat{d}(k) \\ u(k) \end{bmatrix} \stackrel{\text{iid}}{\sim} \text{N}(0, S_d) \quad (4.21)$$

The ML estimate of S_d is therefore $\hat{S}_d = \sum_{k=p}^{N-1} \tilde{e}(k)[\tilde{e}(k)]^\top$ (Anderson, 2003, Thm. 8.2.1), which completes the algorithm.

4.3 Closed-loop subspace identification

In this section, we describe a regularized version of the CVA algorithm of Larimore (1983, 1997, 2005). As in Section 4.2.1, the algorithm's goal is to estimate the parameters $\theta = (A, B, C, Q_w, R_v)$ of the model (4.1) from input-output data $(U_N(0), Y_N(0))$. The algorithm, outlined in Figure 4.2, can be viewed as a nested modeling procedure using maximum likelihood (ML) at each step to compute parameter estimates. We refer the reader to Gong and Samaniego (1981) for a theoretical justification of nested ML estimation. As in Section 4.2.1, the algorithm takes two basic steps. First, we determine a state sequence $\tilde{X}_{N_s+1}(p)$ via approximations of the ML problem corresponding to the *marginal density* in $Y_N(0)$,

$$\max_{\theta} L_N(\theta) := \ln p(Y_N(0)|U_N(0), \theta) \quad (4.22)$$

where $\theta = (A, B, C, Q_w, R_v)$ are the system parameters and $p \geq n$ is an integer to be defined. Then, we solve the ML problem (4.3) corresponding to the *joint density* in $(X_{N_s}(p+1), Y_{N_s}(p))$.

To begin, we choose $n \leq f, p \ll N$. These denote the past and future horizons of the

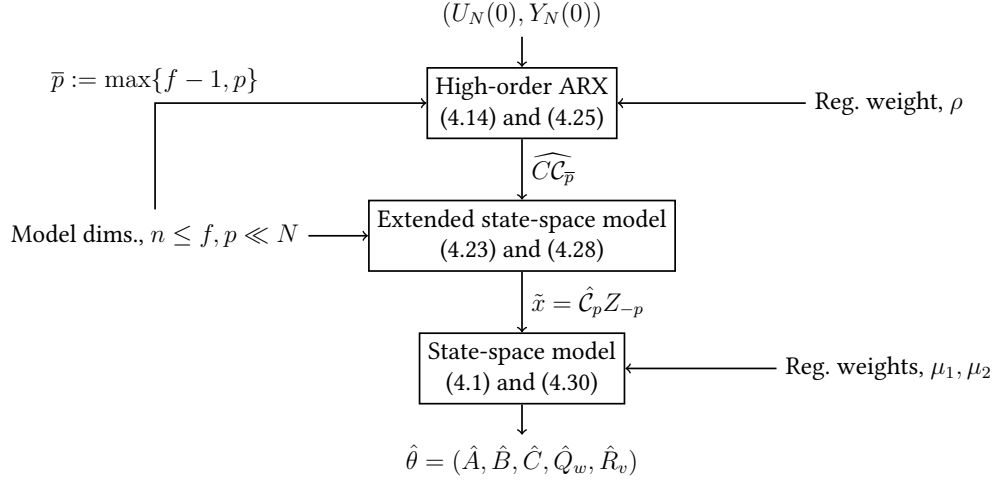


Figure 4.2: Outline of the closed-loop regularized subspace method, based on the work of Larimore (1983, 1997, 2005).

following *extended state-space model*:

$$\begin{aligned}
 Y_f(k) &= \mathcal{O}_f A_K^p \hat{x}(k-p) + \mathcal{H}_{f,p} Z_{-p}(k) + \mathcal{G}_f Z_f(k) + E_f(k) \\
 &\approx \mathcal{H}_{f,p} Z_{-p}(k) + \mathcal{G}_f Z_f(k) + E_f(k).
 \end{aligned} \tag{4.23}$$

where $\mathcal{O}_f, \mathcal{H}_{f,p}, \mathcal{C}_p$ are defined as before, and

$$\mathcal{G}_f := \begin{bmatrix} 0 & & & & & \\ G_1 & 0 & & & & \\ \vdots & \ddots & \ddots & & & \\ G_{f-1} & \dots & G_1 & 0 & & \end{bmatrix}.$$

As in Section 4.2.1, our goal is to estimate the matrix \mathcal{C}_p . Again, if the model (4.1) is minimal, the coefficient matrix $\mathcal{H}_{f,p}$ has rank n . Assuming we have access to some Markov parameter estimates \hat{G}_i , the extended state-space model (4.23) takes the form of a classic rank-reduced regression problem for which closed-form solutions are well-known (Larimore, 1983; Anderson, 1999).¹

¹Since the regressors Z_f are correlated with the errors E_f , joint estimation the parameters $(\mathcal{H}_{f,p}, \mathcal{G}_f)$ produces inconsistent estimates. As a result, we have to eliminate the $\mathcal{G}_f Z_f$ term of the model (4.23) via a “pre-

Before deriving the estimates to the above problems, it is worth discussing other closed-loop subspace algorithms. Other classic subspace algorithms—such as N4SID (Van Overschee and De Moor, 1994) and MOESP (Verhaegen, 1994)—can be used to supply parameters to the disturbance model identification method of Section 4.2.2, so long as a Markov parameter “pre-estimation” step is included. Van Overschee and De Moor (1995) showed that the classic subspace algorithms (CVA, N4SID, and MOESP) are equivalent up to formulation of the estimation objective for estimation of the model (4.23).² Only the method of Larimore (1983, 1997, 2005) uses ML estimation at each step of the algorithm, making it the logical choice for integration with the ML-based disturbance model identification. To use the method of Section 4.2.2, one should take care to use methods that construct state sequences, rather than those that construct the parameters (A, B, C) directly from the matrices $(\mathcal{O}_f, \mathcal{C}_p)$. In fact, any closed-loop state-space identification method that estimates state sequences can be directly integrated with the method of Section 4.2.2.

4.3.1 Estimating the Markov parameters

We must first obtain Markov parameter estimates \hat{G}_i . Jansson (2003) first proposed “pre-estimation” of the Markov parameters G_i from the ARX model (4.14). The likelihood function corresponding to the ARX model (4.14) is an approximation of the likelihood function $L_N(\theta)$,

$$L_N(\theta) \approx L_N^{\text{ARX}}(CC_{\bar{p}}, R_e) := \sum_{k=\bar{p}}^{N-1} \ln p(y(k) | Z_{-\bar{p}}(k), CC_{\bar{p}}, R_e) \\ \propto -\frac{N-\bar{p}}{2} \ln \det R_e - \frac{1}{2} \sum_{k=\bar{p}}^{N-1} |y(k) - CC_{\bar{p}}Z_{-\bar{p}}(k)|_{R_e^{-1}}^2$$

estimation” step. In open-loop subspace methods, a slightly different extended state-space model is used, allowing for consistent estimation of both of the parameters $(\mathcal{H}_{f,p}, \mathcal{G}_f)$, but only under open-loop conditions.

²N4SID and MOESP methods are weighted least squares problems, whereas CVA is an approximate ML problem.

as in (4.15). Closed-form solutions are given by (4.16). The ARX model is an overparameterization of the model (4.1), so it is beneficial to regularize the coefficients, trading a biased estimate for reduced variance,³

$$\max_{CC_{\bar{p}}, R_e > 0} L_N^{\text{ARX}}(CC_{\bar{p}}, R_e) - \frac{\rho}{2} \text{tr}(R_e^{-1} CC_{\bar{p}}(CC_{\bar{p}})^\top) \quad (4.24)$$

which results in the regularized estimates,

$$\widehat{CC}_{\bar{p}} = \left(\sum_{k=\bar{p}}^{N-1} y(k)[Z_{-\bar{p}}(k)]^\top \right) \left(\sum_{k=\bar{p}}^{N-1} Z_{-\bar{p}}(k)[Z_{-\bar{p}}(k)]^\top + \rho I \right)^{-1}. \quad (4.25)$$

The estimates (4.25) are unbiased when $\rho = 0$ and consistent for all $\rho \geq 0$.⁴ Moreover, the estimate errors $\mathcal{E}_{\text{ARX}} := \widehat{CC}_{\bar{p}} - CC_{\bar{p}} = [\hat{G}_1 - G_1, \dots, \hat{G}_{\bar{p}} - G_{\bar{p}}]$ are independent of the innovation sequence $e(k)$ and regression vectors $Z_{-\bar{p}}(k)$.

4.3.2 Estimating the state sequence

It turns out that the likelihood of the extended state-space model (4.23), even though the errors $E_f(k)$ are serially correlated, is an approximation of the likelihood in the ML problem (4.22) (Larimore, 1997). Assume that $N \gg f, p$. Then, for each $s \in \mathbb{I}_{p:p+f-1}$, we can use successive conditioning to write the likelihood as

$$\begin{aligned} L_N(\theta) &\approx \ln p(Y_{(M_s-1)f+s}(s) | Y_{-s}(s), U_{(M_s-1)f+s}(s), \theta) \\ &\approx \sum_{m=0}^{M_s-1} \ell_{f,p}(mf + s) \end{aligned}$$

³The regularizer here is close to using the prior $(CC_{\bar{p}})_i \stackrel{\text{iid}}{\sim} N(0, \rho^{-1} R_e)$ where $(CC_{\bar{p}})_i$ denotes the i -th column of $CC_{\bar{p}}$, but to be equivalent, we would also need to add $(-\rho(n_u + n_y)\bar{p}/2) \ln \det R_e$ to the likelihood.

⁴This neglects numerical errors introduced by the approximation $A_K^p \approx 0$.

where $M_s := \lfloor (N - s)/f \rfloor$ and

$$\ell_{f,p}(k) := \ln p(Y_f(k)|Z_{-p}(k), U_f(k), \mathcal{H}_{f,p}, \mathcal{G}_f, \mathcal{R}_f).$$

Terms at times $k \in \mathbb{I}_{0:N-1} \setminus \mathbb{I}_{s:M_s f + s - 1}$ can be dropped because of the assumption that $N \gg f, p$. Taking the average over s gives

$$\begin{aligned} L_N(\theta) &\approx \frac{1}{f} \sum_{s=p}^{p+f-1} \sum_{m=0}^{M_s-1} \ell_{f,p}(mf + s) \\ &= \frac{1}{f} \sum_{k=p}^{N-f} \ell_{f,p}(k) =: \frac{1}{f} L_N^{\text{ESS}}(\mathcal{H}_{f,p}, \mathcal{G}_f, \mathcal{R}_f). \end{aligned}$$

For closed-loop data, the signals Z_f and E_f are correlated, which may introduce bias into the estimates if all the parameters $(\mathcal{H}_{f,p}, \mathcal{G}_f, \mathcal{R}_f)$ are estimated simultaneously (Qin, 2006). Noting that the future data coefficients \mathcal{G}_f is simply a linear function of the ARX coefficients, i.e. $\mathcal{G}_f = \mathcal{L}(CC_{\bar{p}})$, the future data term in the model (4.23) can be eliminated as follows,

$$\tilde{Y}_f(k) := Y_f(k) - \hat{\mathcal{G}}_f Z_f(k) \approx \mathcal{H}_{f,p} Z_{-p}(k) + \mathcal{E}_{\text{ESS}}(k) \quad (4.26)$$

where $\hat{\mathcal{G}}_f := \mathcal{L}(\widehat{CC_{\bar{p}}})$, and $\mathcal{E}_{\text{ESS}} := \mathcal{L}(\mathcal{E}_{\text{ARX}})Z_f + E_f$ is zero-mean since \mathcal{E}_{ARX} and Z_f are independent. Importantly, the signals Z_{-p} and \mathcal{E}_{ESS} are uncorrelated, so the parameters $(\mathcal{H}_{f,p}, \mathcal{R}_f)$ can be estimated without bias. The corresponding likelihood function is

$$L_N^{\text{ESS}}(\mathcal{H}_{f,p}, \hat{\mathcal{G}}_f, \mathcal{R}_f) \propto -\frac{N - f - p + 1}{2} \ln \det \mathcal{R}_f - \frac{1}{2} \sum_{k=p}^{N-f} |\tilde{Y}_f(k) - \mathcal{H}_{f,p} Z_{-p}(k)|_{\mathcal{R}_f^{-1}}^2$$

and we have the following ML problem:

$$\max_{\text{rank } \mathcal{H}_{f,p}=n, \mathcal{R}_f > 0} L_N^{\text{ESS}}(\mathcal{H}_{f,p}, \hat{\mathcal{G}}_f, \mathcal{R}_f). \quad (4.27)$$

The ML problem (4.27) corresponds to a rank-reduced regression.⁵ To solve it, consider the following definitions,

$$\begin{aligned}
 S_{\tilde{Y}_f, Z_{-p}} &:= \frac{1}{N-f-p+1} \sum_{k=p}^{N-f+1} \tilde{Y}_f(k) [Z_{-p}(k)]^\top \\
 S_{Z_{-p}, Z_{-p}} &:= \frac{1}{N-f-p+1} \sum_{k=p}^{N-f+1} Z_{-p}(k) [Z_{-p}(k)]^\top \\
 S_{\tilde{Y}_f, \tilde{Y}_f} &:= \frac{1}{N-f-p+1} \sum_{k=p}^{N-f+1} \tilde{Y}_f(k) [\tilde{Y}_f(k)]^\top.
 \end{aligned}$$

According to Larimore (1983); Anderson (1999), the ML problem (4.27) has a closed-form solution,

$$\hat{\mathcal{H}}_{f,p} = S_{\tilde{Y}_f, Z_{-p}} J_n^\top J_n$$

where J_n denotes the first n rows of

$$J = U^\top S_{Z_{-p}, Z_{-p}}^{-1/2}$$

and U are the left singular vectors of the following full SVD,

$$S_{Z_{-p}, Z_{-p}}^{-1/2} S_{\tilde{Y}_f, Z_{-p}}^\top S_{\tilde{Y}_f, \tilde{Y}_f}^{-1/2} = U S V^\top.$$

Given these estimates, we have the rank factorization $\hat{\mathcal{H}}_{f,p} = \hat{\mathcal{O}}_f \hat{\mathcal{C}}_p$ where

$$\hat{\mathcal{O}}_f = S_{\tilde{Y}_f, Z_{-p}} J_n^\top$$

and $\hat{\mathcal{C}}_p = J_n$. Moreover, the estimate $\hat{\mathcal{C}}_p$ is a consistent and asymptotically normal estimator

⁵The rank constraint is a consequence of $(\mathcal{O}_f, \mathcal{C}_p)$ showing up in the regression model as the product $\mathcal{H}_{f,p} := \mathcal{O}_f \mathcal{C}_p$, where we assume $(\mathcal{O}_f, \mathcal{C}_p)$ are both rank- n so the states come from a minimal realization.

of \mathcal{C}_p (up to similarity transformation) Anderson (1999). Therefore, we have consistent and asymptotically normal estimates of the states,

$$\tilde{x} := J_n Z_{-p}. \quad (4.28)$$

4.3.3 Estimating the state-space parameters

In practice, we have found that the state estimates (4.28) may contain spurious, unwanted dynamics, so we may regularize this objective in a similar manner to the ARX problem (4.24),

$$\max_{\theta} L_N^{\text{SS}}(\tilde{X}_{N_s+1}(p), \theta) - \frac{\mu_1}{2} \text{tr}(Q_w^{-1}(AA^\top + BB^\top)) - \frac{\mu_2}{2} \text{tr}(R_v^{-1}CC^\top) \quad (4.29)$$

where $\mu_1, \mu_2 > 0$. According to (Anderson, 2003, Thm. 8.2.1), the regularized estimates are

$$[\hat{A} \quad \hat{B}] = \left(\sum_{k=p}^{N-1} \tilde{x}(k+1)[t(k)]^\top \right) \left(\sum_{k=p}^{N-1} t(k)[t(k)]^\top + \mu_1 I \right)^{-1} \quad (4.30a)$$

$$\hat{C} = \left(\sum_{k=p}^{N-1} y(k)[\tilde{x}(k)]^\top \right) \left(\sum_{k=p}^{N-1} \tilde{x}(k)[\tilde{x}(k)]^\top + \mu_2 I \right)^{-1} \quad (4.30b)$$

$$\begin{aligned} \hat{Q}_w &= \frac{1}{N_s} \sum_{k=p}^{N-1} \tilde{x}(k+1)[\tilde{x}(k+1)]^\top - \frac{1}{N_s} \left(\sum_{k=p}^{N-1} \tilde{x}(k+1)[t(k)]^\top \right) \\ &\quad \times \left(\sum_{k=p}^{N-1} t(k)[t(k)]^\top + \mu_1 I \right)^{-1} \left(\sum_{k=p}^{N-1} t(k)[\tilde{x}(k+1)]^\top \right) \end{aligned} \quad (4.30c)$$

$$\begin{aligned} \hat{R}_v &= \frac{1}{N_s} \sum_{k=p}^{N-1} y(k)[y(k)]^\top - \frac{1}{N_s} \left(\sum_{k=p}^{N-1} y(k)[\tilde{x}(k)]^\top \right) \\ &\quad \times \left(\sum_{k=p}^{N-1} \tilde{x}(k)[\tilde{x}(k)]^\top + \mu_2 I \right)^{-1} \left(\sum_{k=p}^{N-1} \tilde{x}(k)[y(k)]^\top \right). \end{aligned} \quad (4.30d)$$

Since \tilde{x} are consistent estimates and independent of the errors (w, v) , the estimates (4.30) are consistent. This completes the closed-loop subspace identification of the model (4.1) from an

input-output sequence.

Part II

Application

Chapter 5

Case studies in combined identification and offset-free control

In this chapter, we apply the methods of Part I to the identification and offset-free control of two application systems: a benchmark temperature control laboratory, and an industrial-scale chemical reactor at Eastman Chemical Company's plant in Kingsport, TN. We remark that these studies primarily use real-world data and experiments rather than simulated experiments. This is important to validate the performance of our turnkey model predictive control (MPC) design, as it ensures we have offset-free performance with regard to physically relevant plant disturbances. Of course, we could simulate the plant and disturbance as a linear augmented disturbance model (LADM) Equation (1.4), but this is a fictitious disturbance model, and the performance may have no relevance to real-world applications.

5.1 Systems of interest

5.1.1 TCLab: a benchmark temperature controller

The TCLab (Figure 5.1), an Arduino-based temperature control laboratory, serves as a low-cost¹ benchmark for linear MIMO control Park et al. (2020). It is a prototypical system

¹The TCLab is available for under \$40 from <https://apmonitor.com/heat.htm> and <https://www.amazon.com/gp/product/B07GMFWMRX>.





Figure 5.1: Benchmark temperature Control Laboratory (TCLab) Park et al. (2020).

of the form (1.4) with internal temperatures as plant states $x = \begin{bmatrix} T_1 & T_2 \end{bmatrix}^\top$, heater voltages as inputs $u = \begin{bmatrix} V_1 & V_2 \end{bmatrix}^\top$, measured temperatures $y = \begin{bmatrix} T_{m,1} & T_{m,2} \end{bmatrix}^\top$ as outputs, and environmental temperatures $d = \begin{bmatrix} T_{a,1} & T_{a,2} \end{bmatrix}^\top$ as disturbances. The control objective is to track setpoints in the measured temperatures, subject to environmental disturbances, such as ambient temperature fluctuations, changes in air circulation, and curious pets.

5.1.2 Eastman industrial-scale chemical reactor

Experiments were also conducted on a reactor at Eastman Chemical's plant in Kingsport, Tennessee. The chosen process is similar to that used in Caveness and Downs (2005). The process produces dimethyl terephthalate (DMT) by reacting terephthalic acid (TPA) with methanol (MeOH). Water is a byproduct of the reaction. The primary equilibrium reaction can be represented as



TPA is a solid and enters the reactor in a slurry with methanol, and additional methanol enters as a vapor. The reactor has two phases. The reaction takes place in a liquid phase, and the DMT product, water, excess methanol, and side products leave the reactor as a vapor

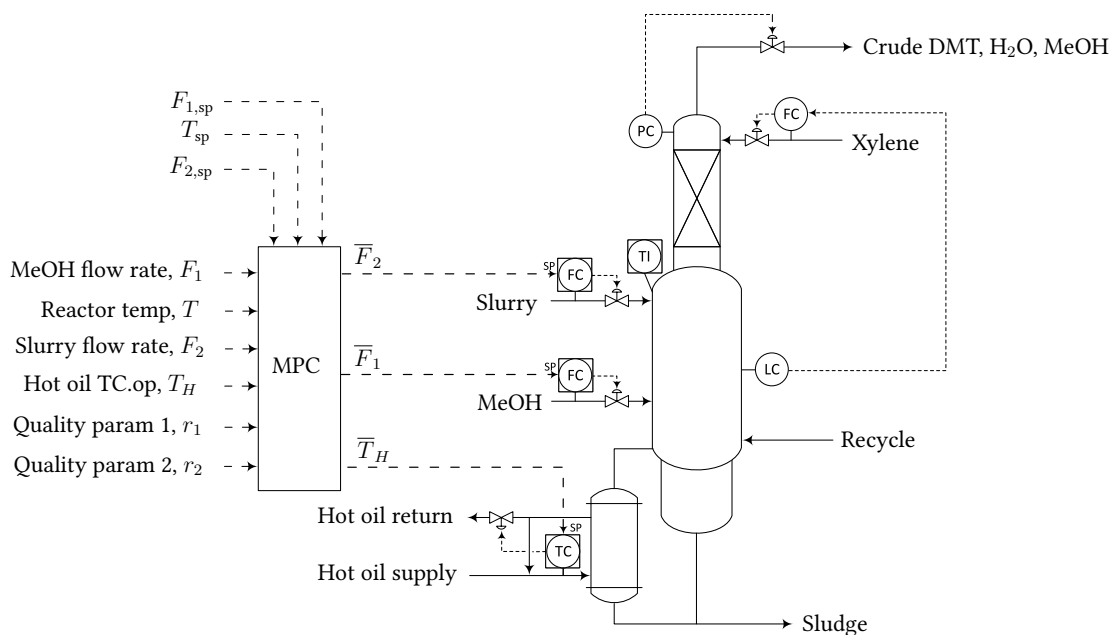


Figure 5.2: Schematic of the DMT reactor and MPC control strategy.

and move forward to a DMT purification section. Xylene is added as reflux to minimize the carryover of an impurity that results from the half reaction of TPA and methanol. Xylene does not participate in the reaction. A schematic of the reactor is shown in Figure 5.2.

The reactor operates under pressure, which is controlled by manipulating a valve in the vapor line. Heat is supplied to the reboiler by circulating hot oil through the shell side of the exchanger. A temperature controller manipulates the flow of hot fluid supplying the circulation loop to control the temperature of the heating fluid entering the reboiler. Liquid level is controlled by manipulating the xylene reflux. Any change in the material balance that affects the composition of methanol in the reactor has a large influence on reactor temperature. Infinite-horizon MPC (2.18), (2.19), and (2.21) is used to control the reactor temperature, T , and the production rate (ultimately set by the slurry feed, F_2) and to maintain the methanol feed, F_1 , at a desired rate. The MPC also handles constraints on two quality-control variables, r_1 and r_2 , and on the hot oil controller valve position (used to infer a temperature pinch/constraint on hot oil temperature, T_H). The manipulated variables are the PID loop

setpoints for the inlet flowrate and utility temperature controllers, denoted $(\bar{F}_1, \bar{F}_2, \bar{T}_H)$.

The control objectives are to achieve offset-free setpoint tracking and disturbance rejection and to avoid violating box constraints on the measured and manipulated variables. For several decades the reactor has run on an MPC designed with a step response model (to be referred to as the “previous MPC model”) and hand-tuned estimator, as described in Caveness and Downs (2005). The inlet flowrate “measurements” are actually “wrap-around” variables, that is, each flowrate “measurement” is generated by passing the corresponding PID setpoint (the MPC’s actuator) through a first-order filter.² We refer to these fictitious flowrate “measurements” as the “wrap-around” variables, and the actual flowrate data, collected from the PID layer, as the raw sensor data. We refer to the complete dataset (3 inputs, 6 outputs) formed with the “wrap-around” variables as the “MPC variables” and the complete dataset formed with the sensor data as the “raw sensor data.” The MPC runs at a sample time of 5 seconds.

5.2 Maximum likelihood identification

In this section, we present two real-world case studies in which Algorithm 1 is used to identify the LADM (3.7) and implement offset-free MPC. In the first case study, we consider the TCLab (Figure 5.1). We identify the TCLab from open-loop data and use the resulting model to design an offset-free MPC. We compare closed-loop control and estimation performance of these models to that of offset-free MPCs designed with the identification methods in Chapter 4. In the second case study, data from an industrial-scale chemical reactor is used to design Kalman filters for the linear augmented disturbance model, and the closed-loop esti-

²Given the clarity of hindsight, we would not design the MPC with these fictitious variables. However, our objective in this paper is not to scrutinize the MPC organizational design (that is, the variable choices) but to identify and validate a flexible replacement model via closed-loop experiments. It is worth pointing out that practitioners and academics alike agree that a significant opportunity in MPC performance gains is in improving the organizational structure of implementations (Darby and Nikolaou, 2012).

mation performance is compared to that of models identified using the methods in Chapter 4.

Throughout these experiments, we use an ℓ_2 regularization term in the transformed space,

$$-\ln p_0(\beta, L^{\mathcal{I}_\Sigma}) \propto R_0(\beta, L^{\mathcal{I}_\Sigma}) := \frac{\rho}{2} \left(|\beta - \bar{\beta}|^2 + \|L_\Sigma(\beta, L^{\mathcal{I}_\Sigma}) - L_\Sigma(\bar{\beta}, \bar{L}^{\mathcal{I}_\Sigma})\|_{\mathbb{F}}^2 \right). \quad (5.1)$$

where $\rho \geq 0$ is the regularization weight and $(\bar{\beta}, \bar{L}^{\mathcal{I}_\Sigma}, \bar{L}^{\mathcal{I}_A})$ denote the initial guess for the optimizer.^{3,4}The variable L_A is not regularized. With $\rho = 0$, the MAP problem (3.9) with the regularizer (5.1) simplifies to the standard ML identification problem (3.8).

The initial guess for the ML and MAP problems is based on a nested ML estimation approach described in Chapter 4. The initial guess methods effectively augment standard identification methods (e.g., principal component analysis (PCA), Ho-Kalman (HK), canonical correlation analysis (CCA) algorithms), so we refer to the initial guess models as “augmented” versions of the standard method being used (e.g., augmented PCA, augmented HK, augmented CCA). Each optimization problem is formulated in CasADi via Algorithm 1 and solved with IPOPT. Information about each model fit and configuration is presented in Table 5.1. Wall times for a single-thread of an Intel Core i9-10850K processor are reported.

5.2.1 Benchmark temperature controller

Unless otherwise specified, the TCLab is modeled as a two-state, two-disturbance system of the form (3.7), with internal temperatures as plant states $s = \begin{bmatrix} T_1 & T_2 \end{bmatrix}^\top$, heater voltages as inputs $u = \begin{bmatrix} V_1 & V_2 \end{bmatrix}^\top$, and measured temperatures $y = \begin{bmatrix} T_{m,1} & T_{m,2} \end{bmatrix}^\top$ as outputs. Throughout, we choose $n_d = p$ to satisfy the offset-free necessary conditions in Muske and Badgwell (2002); Pannocchia and Rawlings (2003), and we consider output disturbance models

³With $L_\Sigma(\bar{\beta}, \bar{L}^{\mathcal{I}_\Sigma}) = 0$, the last term of (5.1) becomes proportional to $\text{tr}(L_\Sigma L_\Sigma^\top) = \text{tr}(\Sigma)$ where $L_\Sigma = L_\Sigma(\beta, L^{\mathcal{I}_\Sigma})$ and $\Sigma = \Sigma(\beta, L^{\mathcal{I}_\Sigma})$.

⁴With $L^{\mathcal{J}_\Sigma}(\beta, L^{\mathcal{I}_\Sigma}) \equiv 0$ (e.g., Σ is block diagonal and $H(\beta) \equiv 0$) the last term of (5.1) is proportional to $\|L^{\mathcal{I}_\Sigma} - \bar{L}^{\mathcal{I}_\Sigma}\|_{\mathbb{F}}^2 = |\text{vec}_{\mathcal{I}_\Sigma}(L^{\mathcal{I}_\Sigma} - \bar{L}^{\mathcal{I}_\Sigma})|^2$.




Figure 5.3: TCLab identification data and noise-free responses $\hat{y}_k = \sum_{j=1}^k \hat{C} \hat{A}^{j-1} \hat{B} u_{k-j}$ of a few selected models.

Table 5.1: TCLab model fitting results. * The augmented PCA/ARX identification methods are not iterative. ** The maximum number of iterations was set at 500.

Model	Results			Configuration					
	Time (s)	Iterations	$L_N(\hat{\theta})$	Method	ρ	\mathcal{D}	ε	ε_i	
Augmented PCA	0.02	N/A*	3823.4	Sec. 4.2	N/A	N/A	N/A	N/A	
Augmented ARX	0.03	N/A*	3807.3	see text	N/A	N/A	N/A	N/A	
Unregularized ML	121.25	500**	-9430.9	Algo. 1	0	\mathcal{C}	10^{-6}	N/A	
Regularized ML 1	123.39	500**	-9431.7	Algo. 1	0.002	\mathcal{C}	10^{-6}	N/A	
Regularized ML 2	9.17	21	-9416.6	Algo. 1	0.005	\mathcal{C}	10^{-6}	N/A	
Constrained ML 1	72.21	97	-9347.2	Algo. 1	0	$\mathcal{D}_1(0.3) \cap \mathcal{D}_2(0.998, 0)$	10^{-6}	0.03	
Constrained ML 2	49.70	62	-9358.2	Algo. 1	0	$\mathcal{D}_1(0.3) \cap \mathcal{D}_2(0.999, 0)$	10^{-6}	0.03	
Reg. & Cons. ML	36.67	40	-9338.4	Algo. 1	0.001	$\mathcal{D}_1(0.3) \cap \mathcal{D}_2(0.998, 0)$	10^{-6}	0.03	

$(B_d, C_d) = (0_{2 \times 2}, I_2)$. We use (A_s, B_s) fully parameterized and $C = I_2$ to guarantee model identifiability and make the states interpretable as internal temperatures. For the remaining model terms, we have (K_x, K_d, R_e) fully parameterized and $(D, \hat{s}_0, \hat{d}_0) = (0, 0, 0)$.

Eight TCLab models are presented.

1. **Augmented PCA:** the 6-state TCLab model used in Kuntz and Rawlings (2022), where principle component analysis on a 400×5100 data Hankel matrix is used to determine the states in the disturbance-free model.
2. **Augmented ARX:** a VARX(1, 1) model, equivalent to a stochastic LTI model with process noise but zero measurement noise.
- 3–5) **Unregularized ML, Regularized ML 1 and 2:** classic ML and MAP models.




Figure 5.4: TCLab models open-loop and closed-loop (filter) eigenvalues.

6–8) **Constrained ML 1 and 2, Reg. & Cons. ML:** eigenvalue-constrained ML and MAP models. LMI region constraints enforce filter stability and impose a lower bound on the real part of the filter eigenvalues.

Each ML model uses Augmented ARX as the initial guess as it has the smallest number of states. The augmented PCA model is, in effect, an unsupervised learner of the state estimates, and therefore does not produce a parsimonious state description.

In Figure 5.3, the identification data is presented along with the noise-free responses $\hat{y}_k = \sum_{j=1}^k \hat{C} \hat{A}^{j-1} \hat{B} u_{k-j}$ of a few selected models. Computation times, numbers of IPOPT iterations, and *unregularized* log-likelihood $L_N(\hat{\theta})$ values are reported in Table 5.1. The open-loop A and closed-loop $A_K := A - KC$ eigenvalues of each model are plotted in Figure 5.4.

Except for the augmented PCA model, all of the open-loop eigenvalues cluster around the same region of the complex plane (Figure 5.4). The closed-loop filter eigenvalues are also placed similarly, although the classic ML models (Unregularized ML, Regularized ML 1 and 2) suffer from slow or even unstable filter eigenvalues, despite achieving lower $L_N(\hat{\theta})$ values than their eigenvalue-constrained counterparts. The models with unstable eigenvalues fail to converge (Table 5.1) as the unstable filter modes make the problem extremely sensitive to changes in the parameter values. While sufficiently high ρ is sufficient to achieve filter stabil-




Figure 5.5: TCLab setpoint tracking tests.

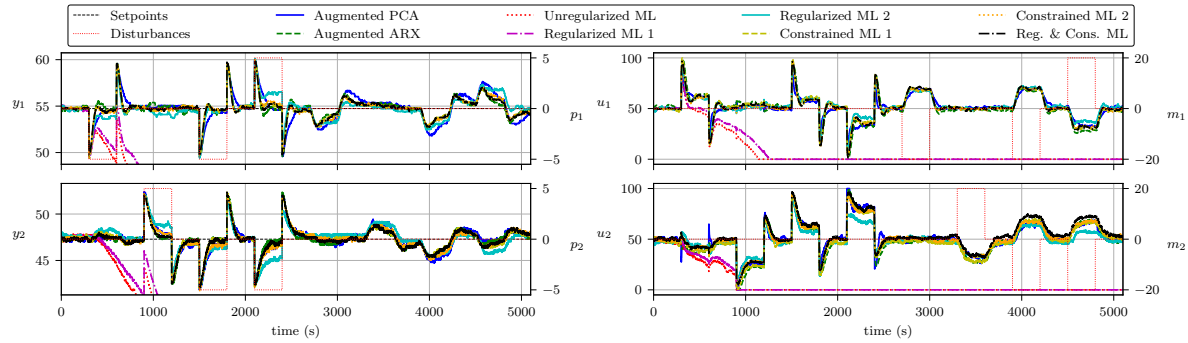


Figure 5.6: TCLab disturbance rejection tests.

ity, there are no clear minimum value of ρ to achieve this. On the other hand, the constrained ML models have stable filter eigenvalues *without regularization*, and have well-defined estimator performance guarantees based on the applied constraints.

To test offset-free control performance, we performed two sets of closed-loop experiments on offset-free MPCs designed with the models. In Figure 5.5, identical setpoint changes were applied to a TCLab running at a steady-state power output of 50%. The setpoint changes were tracked with the finite-horizon offset-free MPC design described in Chapter 2. In Figure 5.6, step disturbances in the output p_i and the input m_i are injected into a plant trying to maintain a given steady-state temperature.

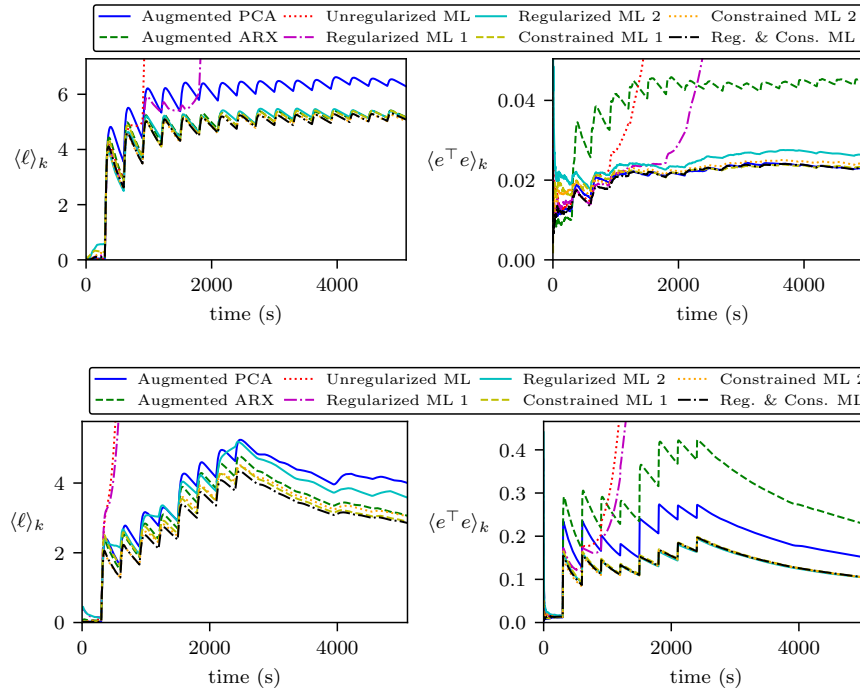


Figure 5.7: TCLab test performance for the (top) setpoint tracking and (bottom) disturbance rejection tests.

Control performance is quantified by the squared distance from the setpoint $\ell_k := |y_k - y_{\text{sp},k}|^2$. Estimation performance is quantified by the squared filter errors $e_k^\top e_k$. For any signal a_k , we define a T -sample moving average by $\langle a_k \rangle_T := T^{-1} \sum_{j=0}^{T-1} a_{k-j}$. Performance on the setpoint tracking and disturbance rejection tests is reported in Figure 5.7. The worst performing models are those with unstable filters (Unregularized ML and Regularized ML 1). These models shut off over the course of the experiment as the integrating disturbance estimates grow unbounded. The remaining classic ML model (Regularized ML 2) has slow filter eigenvalues that contribute to poor control performance on the disturbance rejection test. The augmented models (Augmented PCA/ARX) perform poorly in either control or estimation aspect on both test. The best performance is achieved by the remaining ML models, which all perform approximately the same across the tests.

To investigate the *distributional* accuracy of the models, we quantify performance with




Figure 5.8: TCLab identification index data for (left) setpoint tracking and (right) disturbance rejection tests.

the identification index $q := e^\top R_e^{-1} e$. Recall the signal $e_k \in \mathbb{R}^2$ is an i.i.d., zero-mean Gaussian process, i.e., $e_k \stackrel{\text{iid}}{\sim} \mathcal{N}(0, R_e)$, and therefore the index q_k is i.i.d. with a χ_2^2 distribution. Moreover, the moving average $\langle q_k \rangle_T$ is distributed as χ_{2T}^2/T , although it is no longer independent in time. In Figure 5.8, histograms of $\langle q \rangle_T, T \in \{1, 10, 100\}$ are plotted against their expected distribution for a few selected models (Augmented PCA/ARX, Unregularized ML, and Reg. & Cons. ML). The extreme discrepancies between the augmented models' performance index $\langle q \rangle_T$ and the reference distribution χ_{2T}^2/T are primarily due to the augmented models significantly overestimating R_e compared to the ML models,

$$\begin{aligned} \hat{R}_e^{\text{Aug. PCA}} &= \begin{bmatrix} 0.5871 & 0.3365 \\ 0.3365 & 0.2878 \end{bmatrix}, & \hat{R}_e^{\text{Aug. ARX}} &= \begin{bmatrix} 0.5084 & 0.2198 \\ 0.2198 & 0.2980 \end{bmatrix}, \\ \hat{R}_e^{\text{Unreg. ML}} &= \begin{bmatrix} 0.0106 & 0.0007 \\ 0.0007 & 0.008 \end{bmatrix}, & \hat{R}_e^{\text{Reg. Cons. ML}} &= \begin{bmatrix} 0.0107 & 0.0007 \\ 0.0007 & 0.008 \end{bmatrix}. \end{aligned}$$

The reference distribution and the ML models' $\langle q \rangle_T$ distribution diverge at large T since, due to plant-model mismatch, the filter's innovation errors are slightly autocorrelated. Frequent right-tail errors from the unregularized ML model are due to filter instability.




Figure 5.9: Training data and noise-free responses for the Eastman reactor models (Augmented HK and ML models using Augmented HK as the initial guess).

5.2.2 Eastman reactor

As stated in Section 5.1.2, the control objective of the chemical reactor is to track three setpoints (the output, a specified reactor temperature $y = T$, and the flowrates $\begin{bmatrix} u_1 & u_3 \end{bmatrix}^\top = \begin{bmatrix} F_1 & F_2 \end{bmatrix}^\top$), without offset, by controlling the three inputs (the reactant flow rates and utility temperatures $u = \begin{bmatrix} F_1 & T_H & F_2 \end{bmatrix}^\top$).⁵ See Section 5.1.2 for more details about the reactor operation. As in Section 5.2.1, we choose $n_d = p$ and consider output disturbance models $(B_d, C_d) = (0_{2 \times 1}, 1)$. This time, we use an observability canonical form Denham (1974) with $A_s = \begin{bmatrix} 0 & 1 \\ a_1 & a_2 \end{bmatrix}$ and $C_s = \begin{bmatrix} 1 & 0 \end{bmatrix}$. For the remaining model terms, we have (B_s, K_x, K_d, R_e) fully parameterized and $(D, \hat{s}_0, \hat{d}_0) = (0, 0, 0)$.

Eight reactor models were fit to closed-loop data (from Section 5.3):

1. **Augmented CCA**: a CCA model Larimore (1990) augmented with a disturbance model, as detailed in Section 4.3.⁶
- 3–5) **Unregularized ML, Regularized ML 1 to 3**: classic ML and MAP models.
- 6–8) **Constrained ML 1 to 3**: eigenvalue-constrained ML and MAP models. LMI region

⁵The flowrates are both manipulated variables and controlled variables. At steady state, we should reach the setpoints in $y = T$ and $\begin{bmatrix} u_1 & u_3 \end{bmatrix}^\top = \begin{bmatrix} F_1 & F_2 \end{bmatrix}^\top$, but $u_2 = T_H$ will not reach a predefined setpoint.

⁶This is not the same model used in Section 5.1.2, as a different input-output model is considered, although the same data is used.

Table 5.2: Eastman reactor model fitting results. * The augmented identification methods are not iterative. ** The maximum number of iterations was set at 500.

Model	Results			Configuration				
	Time (s)	Iterations	$L_N(\hat{\theta})$	Method	ρ	\mathcal{D}	ε	ε_i
Augmented CCA	0.09	N/A*	-11399.3	Sec. 4.2.2, 4.3	N/A	N/A	N/A	N/A
Unregularized ML	5.59	19	-14383.1	Algo. 1	0	\mathbb{C}	10^{-6}	N/A
Regularized ML 1	5.46	17	-14362.5	Algo. 1	0.0	\mathbb{C}	10^{-6}	N/A
Regularized ML 2	5.75	20	-14346.7	Algo. 1	0.1	\mathbb{C}	10^{-6}	N/A
Regularized ML 3	4.89	13	-14108.0	Algo. 1	1.0	\mathbb{C}	10^{-6}	N/A
Constrained ML 1	19.89	92	-13944.9	Algo. 1	0	$\mathcal{D}_1(0.3)$	10^{-6}	0.01
Constrained ML 2	16.58	73	-13941.1	Algo. 1	0	$\mathcal{D}_1(0.3)$	10^{-6}	0.02
Constrained ML 3	14.01	58	-13928.5	Algo. 1	0	$\mathcal{D}_1(0.3)$	10^{-6}	0.04




Figure 5.10: Eastman reactor models open-loop and closed-loop (filter) eigenvalues.

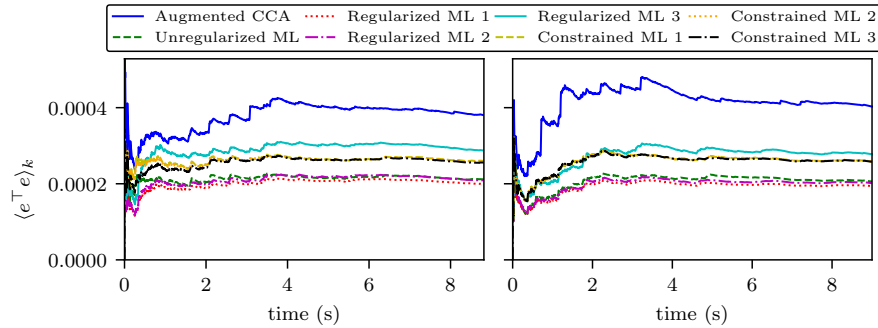


Figure 5.11: Test performance for the Eastman reactor models on the test data sets from Section 5.3.

constraints impose a lower bound on the real part of the filter eigenvalues.

Each ML model uses the augmented CCA model as the initial guess. In Figure 5.9, the closed-loop identification data and noise-free responses are presented. Computational details, the *unregularized* log-likelihood value, and model configuration details are reported in Table 5.2. The open-loop A_s and closed-loop A_K eigenvalues are plotted in Figure 5.10.

The main difference between eigenvalues of the unconstrained ML models (Unregularized ML and Regularized ML 1–3) and the constrained ML models (Constrained ML 1–3) are faster open-loop eigenvalues and closed-loop eigenvalues with possibly negative real part (Figure 5.10). For the constrained ML models, the real part of this fast filter eigenvalue is bounded from below using the LMI region constraint $\mathcal{D}_1(0.3)$. As in the TCLab case study, sufficiently high ρ is sufficient to avoid the negative eigenvalue, but there is no clear cutoff to achieve this.

The estimation performance for these filters are compared on two test data sets (from Section 5.3) in Figure 5.11. While the unconstrained models appear to have the best test performance, it is at a cost of undesirable estimate dynamics. In Figure 5.12, we plot the filter response to an initial guess equal to the eigenvector corresponding to the smallest eigenvalue of A_K . Those filters with eigenvalues having negative real parts exhibit overshoot in the estimate. The best performing filters *without* this behavior are the constrained ML models.

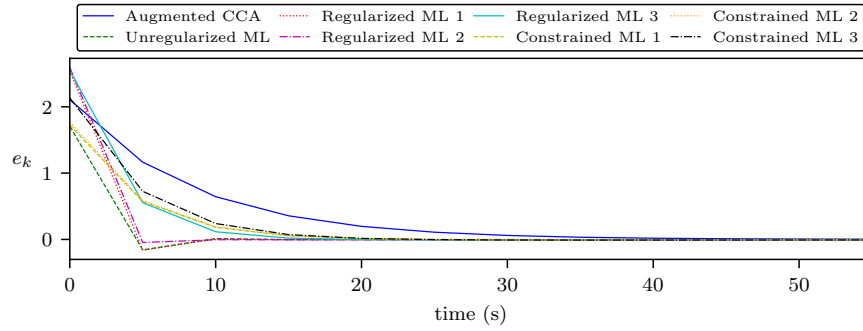


Figure 5.12: Eastman reactor models' closed-loop (filter) response to the eigenvector corresponding to the fastest eigenvalue.




Figure 5.13: Eastman reactor models simulated closed-loop test performance.

Control performance could not be compared on the real plant due to cost and safety considerations. However, the closed-loop responses can be compared in simulation. In Figure 5.13, we plot simulated responses to a setpoint change. Each simulation considers the nominal closed-loop response (i.e., plant as the model, no noise) using the infinite offset-free MPC design in Chapter 2 with $Q_y = Q_s = 1$ and $R = R_s = \text{diag}(0.01, 1, 0.01)$. The regularized ML models exhibit significant overshoot in the response, whereas the unregularized ML model and constrained ML models do not.

5.3 An industrial case study on the combined identification and offset-free control of a chemical process

In this section, we present a closed-loop re-identification of the Eastman reactor (Figure 5.2). In this case study, a closed-loop identification experiment was conducted, a model was identified using the algorithm of Section 4.3 and Section 4.2.2, and closed-loop tests were conducted with MPCs designed with the newly identified model and the previously existing model. Compared to the previous MPC design, a 38% reduction in setpoint tracking error is achieved.

5.3.1 Identification

To identify the process, we used a closed-loop experimental design based on pulses to the normal MPC setpoints. Eight setpoint pulses were applied, each lasting about 30 minutes, with 30 minute “rests” between the pulses to allow the process to settle back to the normal operating point. The setpoint pulses correspond to a full factorial design of the three controlled variables. The pulses were designed to keep the manipulated and measured variables within constraints, and they were checked against historical data to ensure production

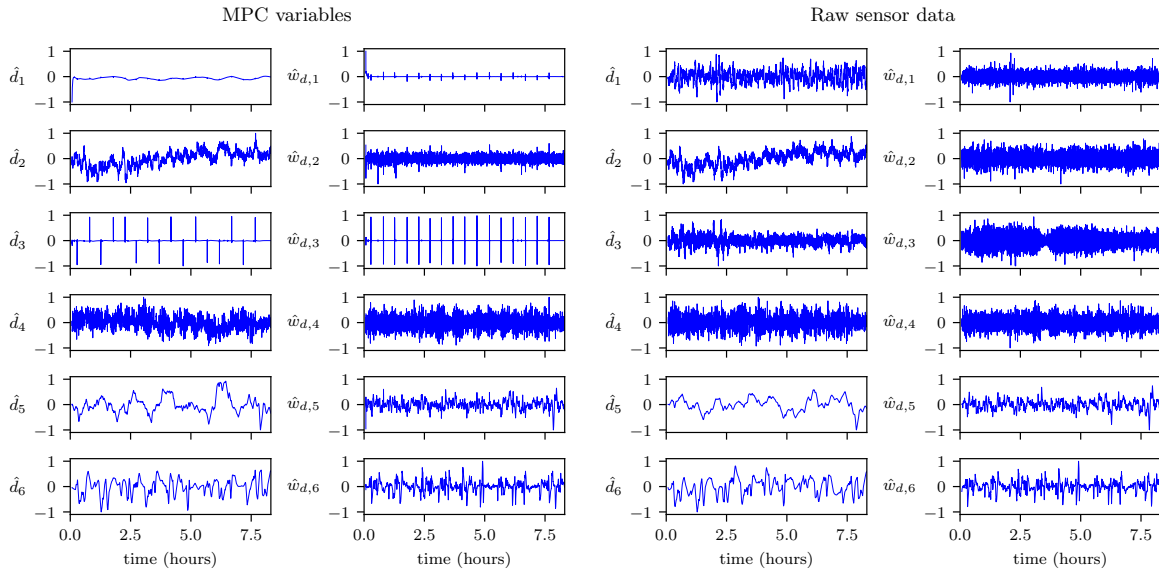


Figure 5.14: Disturbance estimates (4.20) and driving noise estimates $\hat{w}_d = \hat{d}^+ - \hat{d}$ for the unregularized models fit to (left) the MPC variables and (right) raw sensor data. To aid readability, the disturbance estimates were rescaled to have a maximum absolute deviation of 1.

would not be negatively affected. Throughout, models are fit with the algorithm of Section 4.3 and Section 4.2.2.

“Wrap-around” variables and sensor data Models were fit to two sets of process data. The first dataset was constructed from the “wrap-around” variables used on the existing MPC, and the corresponding model uses parameters $n = 20$, $f = 5$, and $p = 50$.⁷ The second dataset was constructed from the raw sensor data, and the corresponding model uses parameters $n = 15$, $f = 5$, and $p = 50$.

In Figure 5.15, for each dataset, we plot process data, setpoint changes, and long-range predictions (4.17) of the previous and new models. The disturbance estimates (4.20) and driving noise $\hat{w}_d = \hat{d}^+ - \hat{d}$ for each model is plotted in Figure 5.14. From Figure 5.14 (left), it

⁷Here we violate the assumption, used in Section 4.3, that $f \geq n$. This condition is only sufficient for producing a rank- n Hankel matrix $\mathcal{H}_{f,p}$. In practice, it is not necessary, so we used the smallest values of (f, p, n) to accurately predict system behavior.

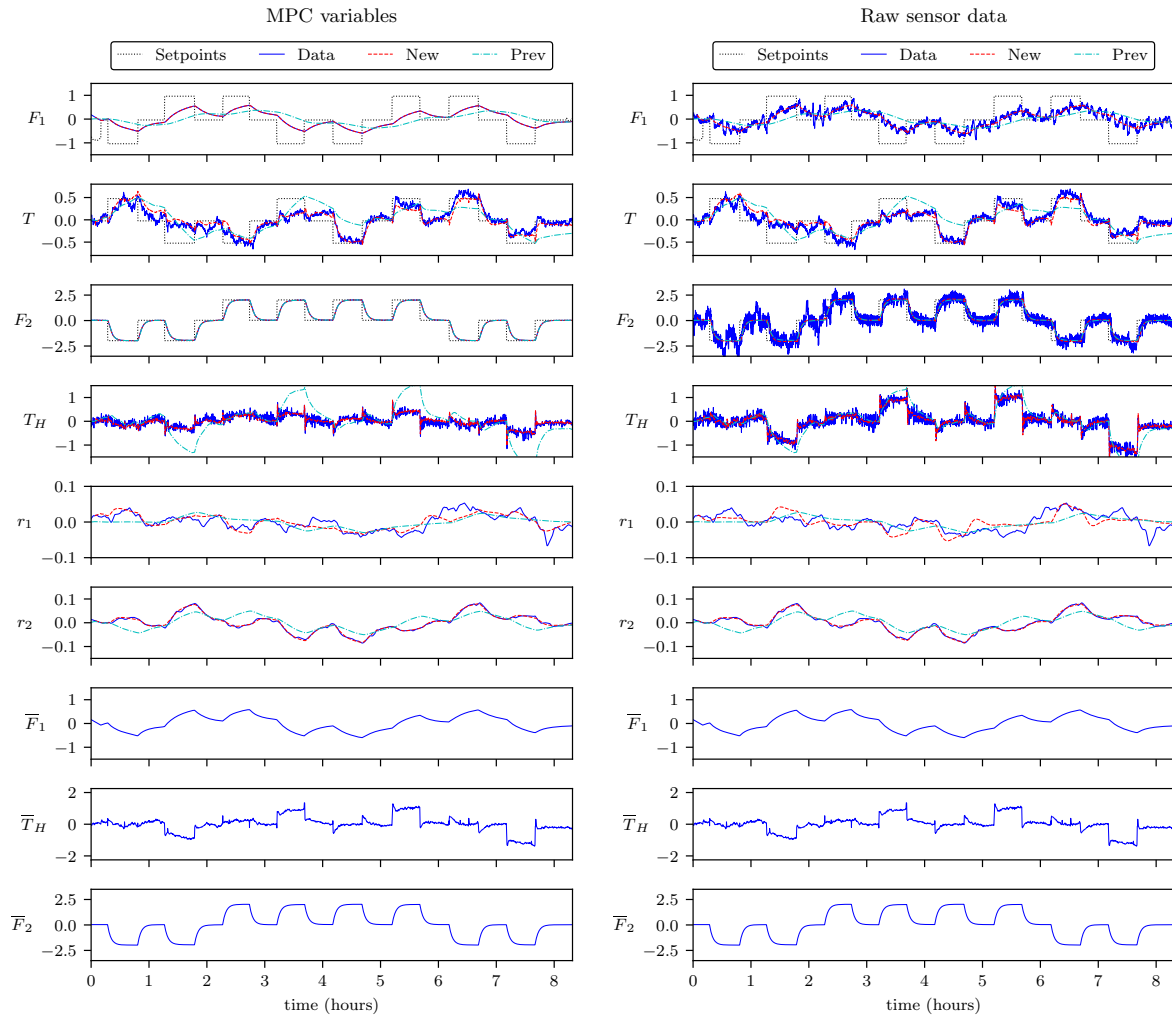


Figure 5.15: Controlled variables (top three), other measured variables (middle three), and manipulated variables (bottom three) for the closed-loop identification experiment using (left) the MPC variables ($n = 15$, $f = 5$, $p = 50$) and (right) the raw sensor data ($n = 15$, $f = 5$, $p = 50$). The dotted lines are MPC setpoints, the dot-dashed lines are the predictions of the previous MPC model, and the dashed lines are the predictions of the new model. Predictions are long-range projections based on a zero initial state (4.17).




Figure 5.16: Disturbance estimates (4.20) and driving noise estimates $\hat{w}_d = \hat{d}^+ - \hat{d}$ for the (left) unregularized model and (right) regularized model. To aid readability, the disturbance estimates were rescaled to have a maximum absolute deviation of 1.

is clear that the model fit to the MPC variables is not driven by white noise. This is to be expected; the MPC variables contain outputs that are *not* constructed from sensor data and therefore do not include upstream disturbances affecting, for example, the PID layer dynamics and offset. As the assumptions of the augmented disturbance model (1.4) are violated, we chose to continue with the model based on raw sensor data, which is clearly driven by white noise (Figure 5.14, right). It is worth pointing out that, in the experiment, the temperature failed to reach the second and fourth setpoints. This is due to plant-model mismatch in the previous MPC model, as that model incorrectly predicts that the temperature will reach the setpoint. The newly identified models do not make such predictions. Additionally, the first flowrate F_1 never reaches any of the setpoints because it has a low regulator weight relative to that of the temperature. Despite the significant noise present in the raw sensor data, the model fit to this data is no worse at predicting the outputs than the model fit to the MPC variables.

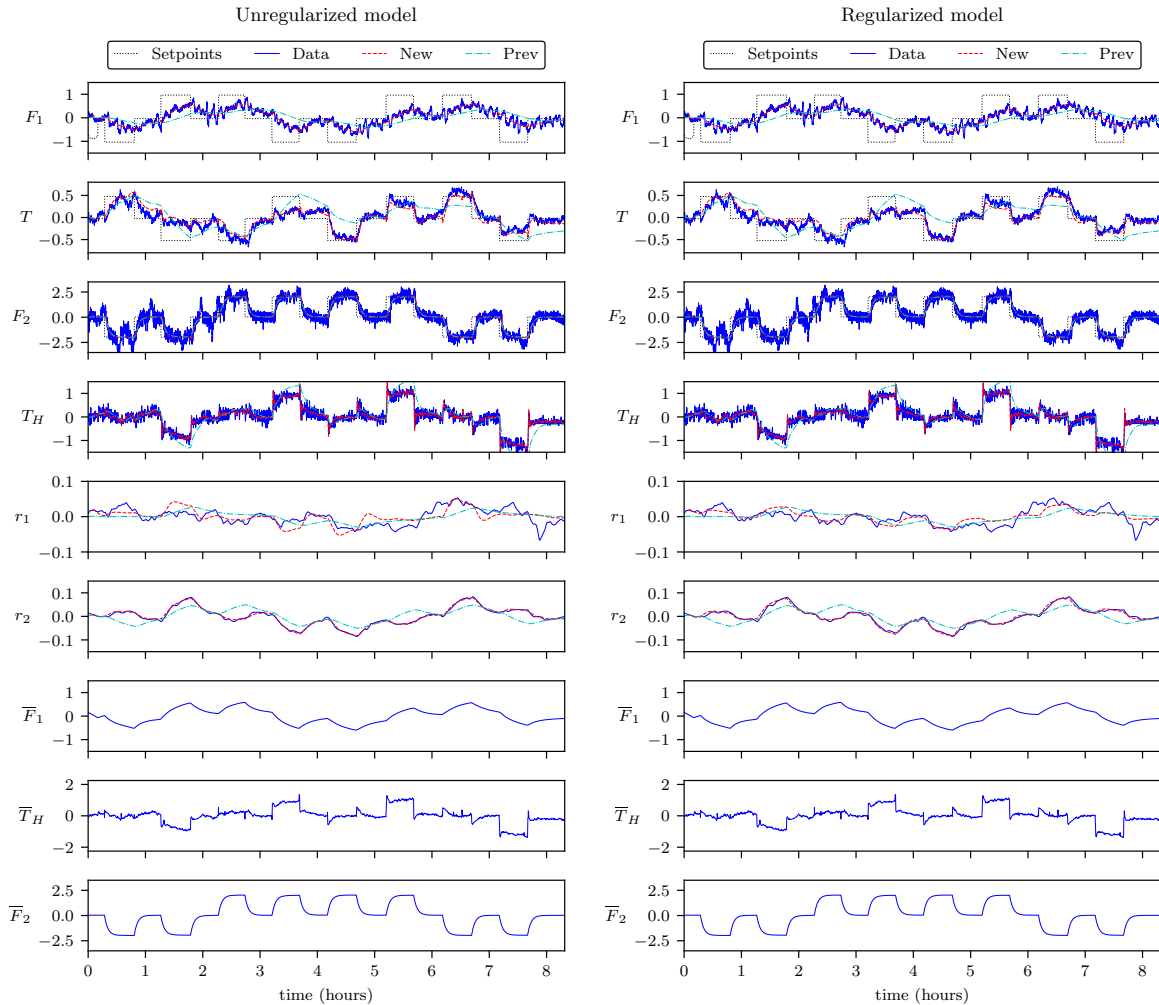


Figure 5.17: Controlled variables (top three), other measured variables (middle three), and manipulated variables (bottom three) for the closed-loop identification experiment using the (left) unregularized model and (right) regularized model. The dotted lines are MPC setpoints, the dot-dashed lines are the predictions of the previous MPC model, and the dashed lines are the predictions of the new model. Predictions are long-range projections based on a zero initial state (4.17).

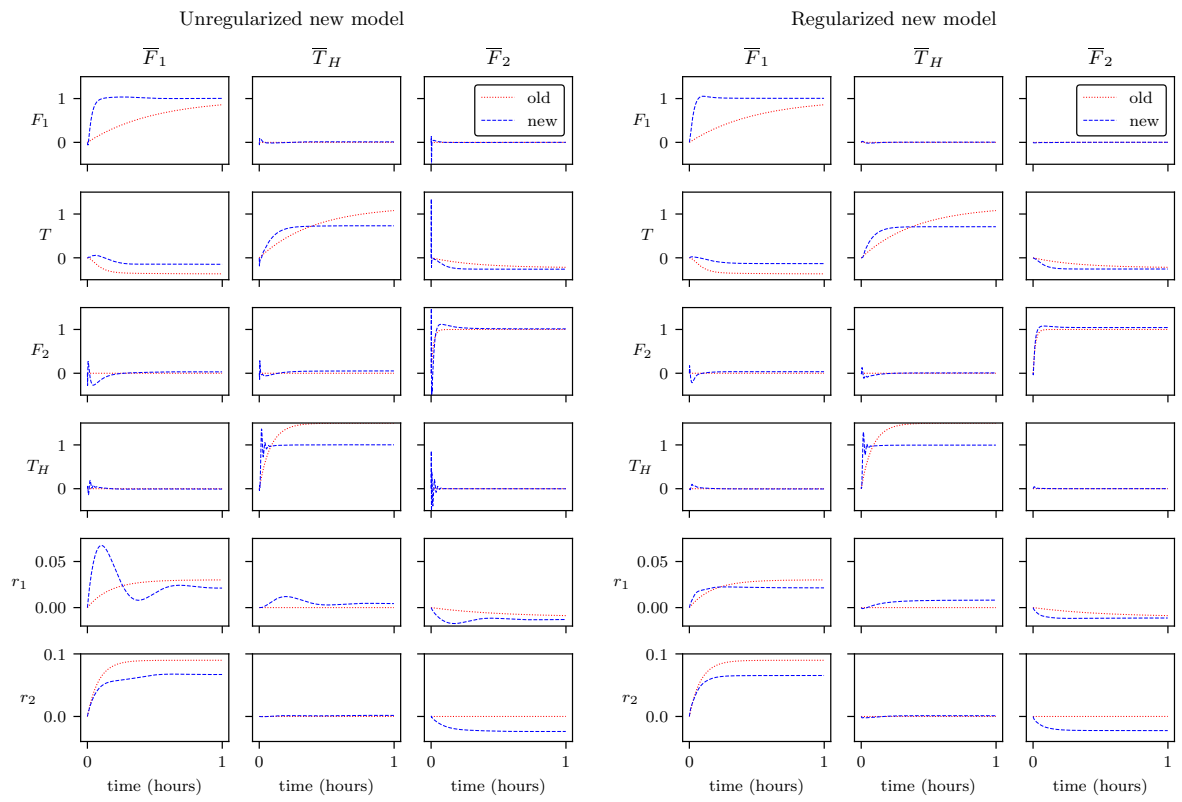


Figure 5.18: Step responses of the (left) unregularized and (right) regularized models compared to the step responses of the previous MPC model.

Model regularization Regularization is a classic technique in statistics and linear algebra used to avoid model over-fitting and ill-conditioning (Tikhonov, 1963; Hoerl and Kennard, 1970a,b). While it is less common in system identification, there is a history of its use for at least three decades (Sjöberg et al., 1993; Johansen, 1997; Chen and Ljung, 2013; Chen et al., 2014). To investigate the possibility of model over-fitting, we also used regularized estimates to produce a model. See Section 4.3 for a derivation of the regularized estimates and the meaning of the regularization parameters. A regularized model was fit to the raw measurement data using parameters $n = 15$, $f = 5$, $p = 50$,

$$\begin{aligned}\rho &= 10^{-4} \left\| \sum_{k=p}^{N-1} Z_{-p}(k) [Z_{-p}(k)]^\top \right\|^2 \\ \mu_1 &= 10^{-7} \left\| \sum_{k=p}^{N-1} Z_{-p}(k) [Z_{-p}(k)]^\top \right\|^2 \\ \mu_2 &= 10^{-4} \left\| \sum_{k=p}^{N-1} \tilde{x}(k) [\tilde{x}(k)]^\top \right\|^2.\end{aligned}$$

Process data, setpoint changes, and long-range predictions (4.17), for both previous and new models, are plotted in Figure 5.17. The disturbance estimates (4.20) and driving noise $\hat{w}_d = \hat{d}^+ - \hat{d}$ for each model is plotted in Figure 5.16.

As a sanity check of the model fits (and to tune the regularization parameters) we plotted the step responses of the unregularized and regularized models (Figure 5.18). At a first glance, the long-range predictions in Figure 5.15 (right) appear to be representative of the true process dynamics. However, when looking at the step responses of the model (Figure 5.18, left) it is clear that there are artifacts and spurious dynamics in the model fit that we speculate is due to over-fitting of the plant model to the disturbance signal in the high frequency range. Regularization takes care of these problems, creating a smoother step response (Figure 5.18, right). As such, we chose to update the MPC on the process in Figure 5.2 with the regularized

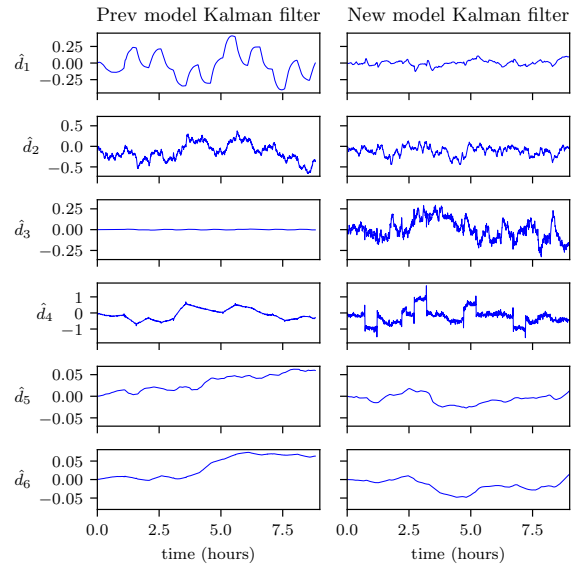


Figure 5.19: Kalman filtered disturbances from the closed-loop tests of the (left) previous and (right) new models using the MPC variables as measurements.

model.

5.3.2 Closed-loop performance

To evaluate the performance of the new MPC model, we used a closed-loop experimental design similar to the one carried out during identification. Again, eight setpoint pulses were applied, each lasting about 30 minutes, with 30 minute “rests” between the pulses to allow the process to settle back to the normal operating point. This time, however, the experiment was carried out over two separate days, switching the MPC model between the two days. Both experiments used the same infinite horizon MPC (2.18), (2.19), and (2.21) with the only difference being the model and estimator gain. It is worth pointing out that, while the new model was fit to the raw sensor data, the MPC uses the “wrap-around” variables in both experiments. As a result, there is a risk the MPC does not respond to disturbances affecting these measurements.

The MPC variables and raw sensor data from these experiments are plotted in Figure 5.21

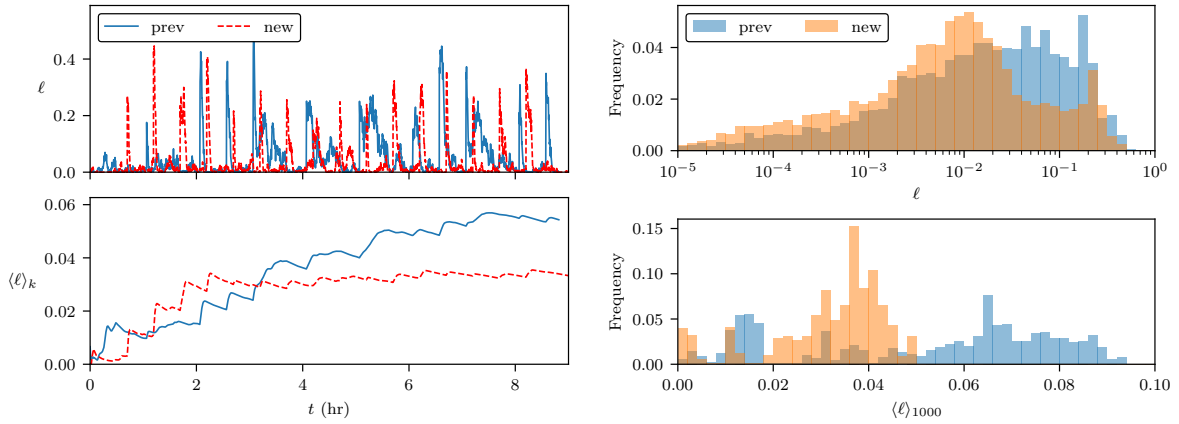


Figure 5.20: Closed-loop performance model of the previous and new models using the MPC variables as measurements. (Top left) Output tracking error $\ell(k)$, (bottom left) running average of the tracking error $\langle \ell(k) \rangle_k$, (top right) histogram of the output tracking error $\ell(k)$, (bottom right) histogram of the output tracking error moving average $\langle \ell(k) \rangle_{1000}$.

and Figure 5.22, respectively. From these plots, it appears that the F_1 valve needed servicing. However, because feedback was done with the “wrap-around” variables, there was no effect on the closed-loop performance. It is also clear that the previous MPC model continues to have difficulties reaching certain temperature setpoints, whereas the new model is confirmed to alleviate these problems. In the new model, deviations from setpoints are zero mean, so they are likely attributable to process noise and upstream disturbances. Again, both MPC implementations fail to reach F_1 setpoints as this variable has a low regulator weight relative to that of the temperature. The Kalman filtered disturbance estimates (2.18) for the previous and new models (using the MPC variables as feedback) are plotted in Figure 5.19. The new model has a much quicker filter gain. This is particularly prevalent in the \hat{d}_4 Kalman filter estimate (which corresponds to the T_H measurement), which is slow for the previous model but virtually instantaneous for the new model.

To quantify the performance of each MPC, we computed the controlled variable tracking cost,

$$\ell(k) := |Hy(k) - r_{\text{sp}}(k)|_{Q_y}^2$$




Figure 5.21: Closed-loop comparison of the (left) previous and (right) new models using the MPC variables.




Figure 5.22: Closed-loop comparison of the (left) previous and (right) new models using the raw process variables as measurements.

where $H = \begin{bmatrix} I_3 & 0 \end{bmatrix}$ and $Q_y = \text{diag}(10^{-4}, 1, 10^{-3})$, which is approximately the squared error between T and its setpoint. Tighter control will exhibit a smaller tracking cost $\ell(k)$, on average. It is known that for linear plants and linear controllers without constraints, the tracking cost $\ell(k)$ has a generalized- χ^2 distribution,⁸ but if it is time-averaged, it will approach a normal distribution (Zagrobelny et al., 2013). We define the T -lagged average at time k as

$$\langle \ell(k) \rangle_T := \frac{1}{T} \sum_{j=0}^{T-1} \ell(k-j).$$

We compare tracking costs $\ell(k)$ and time-averaged tracking costs $\langle \ell(k) \rangle_{1000}$ and $\langle \ell(k) \rangle_k$ for the previous and new models in Figure 5.20. It is immediately clear that the new model performed better than the previous model; the total average tracking cost (Figure 5.20, bottom left) is 38% lower in the new model experiment compared to that of the previous model experiment. The cost $\ell(k)$ (Figure 5.20, top right) fits the linear control assumptions on generalized- χ^2 distribution. Moreover, the time-averaged cost $\langle \ell(k) \rangle_{1000}$ (Figure 5.20, bottom right) is approaching a normal distribution, although there is some residual density near $\ell = 0$ for both experiments. These results suggest the applicability of a statistical performance monitoring scheme such as the one in Zagrobelny et al. (2013).

⁸A generalized- χ^2 random variable is generated by taking the quadratic form of a multivariate normal random variable.

Part III

Theory

Chapter 6

Indirect methods for linear control of nonlinear systems

In a wide variety of control applications, including chemical processes (Westerlund, 1981; Caveness and Downs, 2005; Raghavan et al., 2006), aerospace vehicles (Li et al., 2007; Taylor, 1985), combustion engines (Melgaard et al., 1990), nautical vehicles (Källström and Åström, 1979; Åström and Källström, 1976), and speech recognition (Digalakis et al., 1993), linear approximations of the nonlinear plant are beneficial for the convenience of linear identification relative to that of nonlinear identification and the ability to meet strict computational constraints, e.g., for on-line optimal control. Linear black-box models are particularly useful when first-principles knowledge of the plant dynamics is not available.

To show a linear controller stabilizes a nonlinear plant, the most straightforward option is to demonstrate stability of a linearization of the plant. In this chapter, we present an indirect method to show linear MPC is suitable for tracking constant setpoints for nonlinear systems. Some authors have proposed provably stable and output-tracking linear MPC designs (Limon et al., 2008; Betti et al., 2013; Falugi and Mayne, 2013), but they typically consider linear plants. In Berberich et al. (2022a), a linear MPC for nonlinear tracking was considered, although they assume access to the plant dynamic equations from which linearizations are built. This work differs in that only a bound on the linearization error is required to establish closed-loop

stability of the nonlinear system.

We use the following elementary fact about quadratics throughout, and without reference.

Lemma 6.1. *For each $b, c > 0$, we have $x^2 + 2bx - c > 0$ for all $x \in [0, c/(b + \sqrt{b^2 + c})]$. Moreover, $c/(b + \sqrt{b^2 + c})$ is decreasing in $b > 0$ and increasing in $c > 0$.*

Proof. The roots of $f(x) := x^2 + 2bx - c$ are at $x = -b \pm \sqrt{b^2 + c}$. Let

$$\gamma(b, c) := -b + \sqrt{b^2 + c} = \frac{c}{b + \sqrt{b^2 + c}} > 0. \quad (6.1)$$

Since f is strictly convex, this means $f(x) < 0$ for all $x \in (-b - \sqrt{b^2 + c}, -b + \sqrt{b^2 + c})$. But $-b - \sqrt{b^2 + c} < 0$ and $\gamma(b, c) > 0$, so $f(x) < 0$ for all $x \in [0, \gamma(b, c)]$. Finally, $\gamma(b, c)$ is clearly decreasing in $b > 0$ and increasing in $c > 0$ by the representations (6.1). \square

6.1 Lyapunov's indirect method

Before we consider any control problems, let us consider a simpler problem: how can we infer stability of a plant from a model of it and a bound on the residual error? To this end, we consider the nonlinear system

$$x^+ = f(x) \quad (6.2)$$

and the linear system

$$x^+ = Ax \quad (6.3)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $A \in \mathbb{R}^{n \times n}$.

6.1.1 Linear stability

A fundamental result in linear systems theory is the equivalence of Schur stability of the matrix A and exponential (equivalently, asymptotic) stability of the linear system (6.3). Schur

stability of the matrix A is closely tied to the existence and uniqueness of solutions to the discrete-time Lyapunov equation

$$L_A(P) := P - A^\top P A = Q \quad (6.4)$$

where $L_A : \mathbb{S}_{++}^n \rightarrow \mathbb{S}_{++}^n$ is the Lyapunov operator. With A Schur stable and $Q, P > 0$ satisfying (6.4), the Lyapunov function $V(x) := x^\top P x$ can be used to demonstrate global exponential stability of the linear system (6.3):

$$\underline{\sigma}(P)|x|^2 \leq V(x) \leq \bar{\sigma}(P)|x|^2$$

$$V(Ax) - V(x) = x^\top (A^\top P A - P)x = -x^\top Q x \leq -\underline{\sigma}(Q)|x|^2$$

for all $x \in \mathbb{R}^n$. The following propositions summarize well-known results on stability of (6.3) and the discrete Lyapunov equation (6.4) (Anderson and Moore, 1979) and properties of invertible Lyapunov operators (Gahinet et al., 1990, Prop. 2.1, Thm. 4.1, Ex. 2).

Proposition 6.2. *For any $A \in \mathbb{R}^{n \times n}$, the following statements are equivalent.*

- (a) *The linear system $x^+ = Ax$ is globally exponentially stable.*
- (b) *A is Schur stable.*
- (c) *The Lyapunov operator $L_A : \mathbb{S}_{++}^n \rightarrow \mathbb{S}_{++}^n$ defined in (6.4) is invertible.*

Moreover, if A is Schur stable, then $L_A^{-1}(Q) = \sum_{k=0}^{\infty} (A^\top)^k Q A^k$ for all $Q > 0$.

Proposition 6.3. *For each Schur stable $A \in \mathbb{R}^{n \times n}$, we have*

$$\|L_A^{-1}\| := \max_{Q>0} \frac{\|L_A^{-1}(Q)\|}{\|Q\|} = \|L_A^{-1}(I)\| = \min_{Q>0} \frac{\|L_A^{-1}(Q)\|}{\underline{\sigma}(Q)}.$$

6.1.2 Lyapunov's indirect method

Assuming f is differentiable at the origin we can define A as the Jacobian of f evaluated at the origin,

$$A := \frac{df}{dx}(0). \quad (6.5)$$

In Lyapunov's indirect method, stability of the nonlinear system (6.2) is inferred from that of the linearization (6.3) and (6.5). This can be preferable to direct stability analysis of (6.2) because, assuming A is Schur stable, (6.3) has an easily defined Lyapunov function: $V(x) := x^\top Px$ where $Q, P \in \mathbb{R}^{n \times n}$ are any positive definite matrices satisfying (6.4). The candidate Lyapunov function V still has quadratic upper and lower bounds,

$$\underline{\sigma}(P)|x|^2 \leq V(x) \leq \bar{\sigma}(P)|x|^2 \quad (6.6)$$

for all $x \in \mathbb{R}^n$, but rewriting $V(f(x)) - V(x)$ as

$$\begin{aligned} V(f(x)) - V(x) &= [f(x)]^\top P f(x) - x^\top P x \\ &= [Ax + r(x)]^\top P [Ax + r(x)] - x^\top P x \\ &= x^\top (APA^\top - P)x + 2x^\top A^\top P r(x) + [r(x)]^\top P r(x) \\ &= -x^\top Q x + 2x^\top A^\top P r(x) + [r(x)]^\top P r(x) \end{aligned}$$

where $r(x) := f(x) - Ax$, it is clear the cost decrease has ambiguous sign:

$$V(f(x)) - V(x) \leq -\underline{\sigma}(Q)|x|^2 + 2\|A\|\|P\|\|x\|\|r(x)\| + \|P\|\|r(x)\|^2 \quad (6.7)$$

for all $x \in \mathbb{R}^n$.

In Lyapunov's indirect method, we construct a positive invariant level set of V for which

the cost decrease (6.7) is bounded above by a negative definite quadratic, under the additional assumptions that the linearized system (6.3) and (6.5) is stable and f is also *continuously* differentiable at the origin. The proof is based on the continuous-time versions found in (Khalil, 2002, Thm. 4.7), however, we avoid the ε - δ formalism for the limit in the definition of the derivative and instead employ the following proposition to use \mathcal{K} function arguments throughout.

Proposition 6.4. *Suppose the function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuously differentiable at the origin and $f(0) = 0$. Then there exists $b > 0$ and $\gamma \in \mathcal{K}$ such that*

$$\frac{|f(x) - Jx|}{|x|} \leq \gamma(|x|) \quad (6.8)$$

for all $0 < |x| \leq b$, where $J := \frac{df}{dx}(0)$.

See Appendix 6.A for a proof of Proposition 6.4. With Proposition 6.4, Lyapunov's indirect method of stability through the linearization is readily proven.

Theorem 6.5 (Lyapunov's indirect method). *Suppose $f(0) = 0$ and f is continuously differentiable at the origin. Let $A := (df/dx)(0)$. If $x^+ = Ax$ is stable, then $x^+ = f(x)$ is exponentially stable in a neighborhood of the origin.*

Proof. Denote $r(x) := f(x) - Ax$, let $Q, P \in \mathbb{R}^{n \times n}$ be positive definite matrices satisfying (6.4), and take $V(x) := x^\top Px$ as the candidate Lyapunov function. Again, we have the global bounds (6.6) and (6.7). By Proposition 6.4, there exists $b > 0$ and $\gamma \in \mathcal{K}$ such that $|r(x)|/|x| \leq \gamma(|x|)$ for all $0 < |x| \leq b$. Without loss of generality, assume $\gamma \in \mathcal{K}_\infty$. For each $|x| \leq b$, the bound (6.7) gives

$$V(f(x)) - V(x) \leq -c(|x|)|x|^2 \quad (6.9)$$

where $c(\cdot) := \underline{\sigma}(Q) - (2\gamma(\cdot)\|A\| + [\gamma(\cdot)]^2)\|P\|$. From Lemma 6.1,

$$\gamma(s) < \bar{\gamma} := \frac{\underline{\sigma}(Q)/\|P\|}{\|A\| + \sqrt{\|A\|^2 + \underline{\sigma}(Q)/\|P\|}}$$

implies $c(s) > 0$. Let $\delta := b$ if $b < \gamma^{-1}(\bar{\gamma})$, but choose $\delta \in (0, \gamma^{-1}(\bar{\gamma}))$ otherwise. In either case, we have $c(\delta) > 0$, and since $c(\cdot)$ is monotonically decreasing, (6.9) implies

$$V(f(x)) - V(x) \leq -c(\delta)|x|^2 \quad (6.10)$$

for all $|x| \leq \delta$. Next, let $\rho := \underline{\sigma}(P)\delta^2 > 0$, and $X := \text{lev}_\rho V$. By (6.6), we have that $|x| \leq \sqrt{\rho/\underline{\sigma}(P)}$ implies $V(x) \leq \rho$, which in turn implies $|x| \leq \delta$. Therefore X is a neighborhood of the origin on which (6.10) holds. Moreover, $V(f(x)) \leq V(x) - c(\delta)|x|^2 \leq \rho$ for all $x \in X$ (by (6.10)), so X is positive invariant. Finally, V is an exponential Lyapunov function on X , so $x^+ = f(x)$ is exponentially stable on X . \square

Remark 6.6. To achieve the largest domain of stability, it is useful to state the (Q, P) pair satisfying (6.4) that maximizes the upper bound

$$\bar{\gamma}(Q, P) := \frac{\underline{\sigma}(Q)/\|P\|}{\|A\| + \sqrt{\|A\|^2 + \underline{\sigma}(Q)/\|P\|}}.$$

By Proposition 6.3, we have $\|P\|/\underline{\sigma}(Q) = \|\mathbf{L}_A^{-1}(Q)\|/\underline{\sigma}(Q) \geq \|\mathbf{L}_A^{-1}\|$ with equality when $Q = I$, and $\|\mathbf{L}_A^{-1}\| = \|\mathbf{L}_A^{-1}(I)\|$. Then

$$\bar{\gamma}(Q, P) \leq \gamma_1 := \frac{\|\mathbf{L}_A^{-1}\|^{-1}}{\|A\| + \sqrt{\|A\|^2 + \|\mathbf{L}_A^{-1}\|^{-1}}}$$

with equality when $Q = I$.

6.2 General indirect methods

More generally, we can consider the follow arbitrary nonlinear model:

$$x^+ = \hat{f}(x) \quad (6.11)$$

where $\hat{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$. We seek sufficient conditions for which stability of the model (6.11) implies stability of the plant (6.2) in some neighborhood of the origin. A similar approach to Lyapunov's indirect method is taken: we start with a Lyapunov function \hat{V} for the model (6.11) and show the perturbations let us construct a sublevel set of \hat{V} that is positive invariant and on which \hat{V} is a Lyapunov function for the plant (6.2).

6.2.1 Exponential stability

Suppose the model (6.11) is exponentially stable in some positive invariant neighborhood of the origin \hat{X} with an exponential Lyapunov function $\hat{V} : \hat{X} \rightarrow \mathbb{R}_{\geq 0}$ such that

$$a_1|x|^2 \leq \hat{V}(x) \leq a_2|x|^2 \quad (6.12a)$$

$$\hat{V}(\hat{f}(x)) - \hat{V}(x) \leq -a_3|x|^2 \quad (6.12b)$$

for all $x \in \hat{X}$ and some constants $a_1, a_2, a_3 > 0$. Exponential Lyapunov bounds of the form (6.12) are found in both linear systems analysis and analysis of of MPC with quadratic costs. The cost decrease inequality (6.12b) implies

$$\hat{V}(f(x)) - \hat{V}(x) \leq -a_3|x|^2 + \hat{V}(f(x)) - \hat{V}(\hat{f}(x)) \quad (6.13)$$

for all $x \in \hat{X}$. Often \hat{V} is a quadratic, in which case the following bound can be derived:

$$|\hat{V}(x) - \hat{V}(\hat{x})| \leq a_4|\hat{x}||x - \hat{x}| + a_5|x - \hat{x}|^2 \quad (6.14)$$

for all $x, \hat{x} \in \hat{X}$ and some $a_4, a_5 > 0$. The inequalities (6.12)–(6.14) will form the basis of the indirect stability analysis.

Next, we provide upper bounds on $|f(x) - \hat{f}(x)|$ and $|\hat{f}(x)|$. For $|\hat{f}(x)|$, it follows straightforwardly from (6.12) that, for all $x \in \hat{X}$,

$$|\hat{f}(x)| \leq L|x| \quad (6.15)$$

where $L := \sqrt{(a_2 - a_3)/a_1} > 0$.¹ Suppose f and \hat{f} are continuously differentiable at the origin. By triangle inequality,

$$|f(x) - \hat{f}(x)| \leq |f(x) - Ax| + |\hat{f}(x) - \hat{A}x| + \gamma_0|x|$$

for all $x \in \mathbb{R}^n$, where

$$\gamma_0 := \|A - \hat{A}\|, \quad A := \frac{df}{dx}(0), \quad \hat{A} := \frac{d\hat{f}}{dx}(0).$$

Moreover, by Proposition 6.4, there exist $b_f, \hat{b}_f > 0$ and $\gamma_f, \hat{\gamma}_f \in \mathcal{K}$ for which

$$\frac{|f(x) - Ax|}{|x|} \leq \gamma_f(|x|), \quad \frac{|\hat{f}(\hat{x}) - \hat{A}\hat{x}|}{|\hat{x}|} \leq \hat{\gamma}_f(|\hat{x}|)$$

¹From (6.12), we have $|\hat{f}(x)| \leq \sqrt{\hat{V}(\hat{f}(x))/a_1} \leq \sqrt{(\hat{V}(x) - a_3|x|^2)/a_1} \leq \sqrt{(1 - a_3/a_2)\hat{V}(x)/a_1} \leq L|x|$.

for all $0 < |x| \leq b_f$ and $0 < |\hat{x}| \leq \hat{b}_f$. Therefore, with $\gamma := \gamma_f + \hat{\gamma}_f$, we have

$$|f(x) - \hat{f}(x)| \leq [\gamma(|x|) + \gamma_0]|x| \quad (6.16)$$

for all $|x| \leq b := \min \{b_f, \hat{b}_f\}$. Finally, combining (6.13)–(6.16), we have

$$\hat{V}(f(x)) - V(x) \leq -a_6(|x|)|x|^2 \quad (6.17)$$

for all $x \in \hat{X}$ such that $|x| \leq b$, where $a_6(\cdot) := a_3 - a_4L[\gamma(\cdot) + \gamma_0] - a_5[\gamma(\cdot) + \gamma_0]^2$. This time we cannot make the coefficient $\gamma(\delta) + \gamma_0 > 0$ arbitrarily small for some $\delta > 0$ due to the error between the linearizations $|Ax - \hat{A}x|$. Without loss of generality, assume $\gamma \in \mathcal{K}_\infty$.

Then, so long as

$$\gamma_0 < \gamma_1 := \frac{-a_4L + \sqrt{a_4^2L^2 + 4a_3a_5}}{2a_5}$$

we can take $\delta \in (0, \gamma^{-1}(\gamma_1 - \gamma_0))$ such that $\delta \leq b$ to achieve $a_6(\delta) > 0$, and (6.17) implies the desired cost decrease

$$\hat{V}(f(x)) - V(x) \leq -a_6(\delta)|x|^2 \quad (6.18)$$

for all $x \in \hat{X}$ such that $|x| \leq \delta$.

In the following proposition, we complete the analysis by showing, if $\gamma_0 < \gamma_1$, then (6.17) shows exponential stability of $x^+ = f(x)$ in a positive invariant neighborhood of the origin.

Theorem 6.7. *Suppose f and \hat{f} are continuously differentiable at the origin and \hat{X} is a neighborhood of the origin such that (6.12), (6.14), and $\hat{f}(x) \in \hat{X}$ hold for all $x \in \hat{X}$. If $\gamma_0 < \gamma_1$ (as defined above), then $x^+ = f(x)$ is exponentially stable in a neighborhood of the origin.*

Proof. By assumption, we already have (6.12a) for all $x \in \hat{X}$. From the discussion above we also have, for some $b > 0$ and $\gamma \in \mathcal{K}$, the cost difference (6.17) for all $x \in \hat{X}$ such that $|x| \leq b$, where $a_6(\cdot) := a_3 - a_4L[\gamma(\cdot) + \gamma_0] - a_5[\gamma(\cdot) + \gamma_0]^2$ as before.

Since \hat{X} is a neighborhood of the origin, there exists $c > 0$ such that $|x| \leq c$ implies $x \in \hat{X}$. Let $\delta := \min\{b, c\}$ if $\min\{b, c\} < \gamma^{-1}(\gamma_1 - \gamma_0)$, but choose $\delta \in (0, \gamma^{-1}(\gamma_1 - \gamma_0))$ otherwise. Then $|x| \leq \delta$ implies $x \in \hat{X}$, $a_6(\delta) > 0$, and we have that (6.18) achieves the desired cost decrease for all $|x| \leq \delta$.

Let $\rho := a_1\delta^2$ and $X := \text{lev}_\rho \hat{V}$. Due to (6.12a), X is a neighborhood of the origin since it is a sublevel set of \hat{V} . If $x \in X$, then $\hat{V}(x) \leq \rho = a_1\delta^2$, which implies $|x| \leq \delta$ by (6.12a), and therefore $x \in \hat{X}$ by construction of δ . Therefore, we have (6.12a) and (6.17) for all $x \in X$. Moreover, by the cost decrease, $\hat{V}(f(x)) \leq V(x) - a_6(\delta)|x|^2 \leq V(x) \leq \rho$ for all $x \in X$, so $f(x) \in X$ and X is positive invariant. Finally, \hat{V} is an exponential Lyapunov function on X , so $x^+ = f(x)$ is exponentially stable on X . \square

Alternatively, suppose the exponential Lyapunov function takes the form

$$a_1|x| \leq \hat{V}(x) \leq a_2|x| \quad (6.19a)$$

$$\hat{V}(\hat{f}(x)) - \hat{V}(x) \leq -a_3|x| \quad (6.19b)$$

for all $x \in \hat{X}$ and some constants $a_1, a_2, a_3 > 0$. Moreover, assume \hat{V} is Lipschitz continuous, i.e.,

$$|\hat{V}(x) - \hat{V}(\hat{x})| \leq a_4|x - \hat{x}| \quad (6.20)$$

for all $x \in \hat{X}$ and some constant $a_4 > 0$. The bounds (6.19) and (6.20) can be found in the analysis of MPC with ℓ_1 -norm costs.

The bound (6.16) can still be constructed, and (6.15) holds with the alternate constant $L := (a_2 - a_3)/a_1 > 0$.² Combining (6.15), (6.16), (6.19b), and (6.20), we have

$$\hat{V}(\hat{f}(x)) - \hat{V}(x) \leq -a_6(|x|)|x| \quad (6.21)$$

²From (6.19), we have $|\hat{f}(x)| \leq \hat{V}(\hat{f}(x))/a_1 \leq (\hat{V}(x) - a_3|x|)/a_1 \leq (1 - a_3/a_2)\hat{V}(x)/a_1 \leq L|x|$.

for all $x \in \hat{X}$ such that $|x| \leq b$, where $a_6(\cdot) := a_3 - a_4L(\gamma(\cdot) + \gamma_0)$. So long as $\gamma_0 < \gamma_1 := a_3/(a_4L)$, we can construct a neighborhood of the origin for which (6.21) implies the desired cost decrease. This is done identically to the proof of Theorem 6.7. The result is stated in the following proposition.

Proposition 6.8. *Suppose \hat{f} and f are continuously differentiable at the origin, \hat{X} is a positive invariant neighborhood of the origin for (6.11), and (6.19) and (6.20) hold for all $x \in \hat{X}$. If $\gamma_0 < \gamma_1$ (as defined above), then $x^+ = f(x)$ is exponentially stable in a neighborhood of the origin.*

Proof. The proof is identical to that of Theorem 6.7, but with the alternate definitions $L := (a_2 - a_3)/a_1$, $\gamma_1 := a_3/(a_4L)$, $a_6(\cdot) := a_3 - a_4L\gamma(\cdot)$, and $\rho := a_1\delta$. \square

6.2.2 Asymptotic stability

Suppose the model is asymptotically stable in some positive invariant neighborhood of the origin \hat{X} . Then there exists a Lyapunov function $\hat{V} : X \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\alpha_1(|x|) \leq \hat{V}(x) \leq \alpha_2(|x|) \quad (6.22a)$$

$$\hat{V}(\hat{f}(x)) - \hat{V}(x) \leq -\alpha_3(|x|) \quad (6.22b)$$

for all $x \in \hat{X}$ and some $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ and continuous $\alpha_3 \in \mathcal{PD}$.³ Then (6.22b) implies

$$\hat{V}(f(x)) - \hat{V}(x) \leq -\alpha_3(|x|) + \hat{V}(f(x)) - \hat{V}(\hat{f}(x)) \quad (6.23)$$

for all $x \in \hat{X}$. Achieving the cost decrease bound requires $\alpha_3(|x|)$ grow sufficiently faster than $|\hat{V}(f(x)) - \hat{V}(\hat{f}(x))|$ near the origin. To this end, we consider the following assumption

³Due to (Jiang and Wang, 2002, Lem. 2.8), the function α_3 can be made class- \mathcal{K}_∞ without loss of generality. It is more convenient, however to construct a positive definite α_3 , so we use this approach throughout.

on the continuity of \hat{f} and \hat{V} :

$$|f(x) - \hat{f}(x)| \leq \alpha_4(|x|) \quad (6.24a)$$

$$|\hat{V}(x) - \hat{V}(\hat{x})| \leq \alpha_5(|x - \hat{x}|) \quad (6.24b)$$

for all $x, \hat{x} \in \hat{X}$. Substituting (6.24) into (6.23), we have

$$\hat{V}(f(x)) - \hat{V}(x) \leq -\alpha_3(|x|) + \alpha_5(\alpha_4(|x|)) =: -\alpha_6(|x|) \quad (6.25)$$

for all $x \in \hat{X}$. Again, the sign of the cost decrease is ambiguous. This time, there are no exponents to inform us of the grow rate of the functions α_3 and $\alpha_5 \circ \alpha_4$ near the origin. Instead, we simply require that $\alpha_3 - \alpha_5 \circ \alpha_4$ is positive definite in a neighborhood of the origin and construct a positive invariant sublevel set of \hat{V} within this neighborhood. We prove this in the following proposition.

Proposition 6.9. *Suppose \hat{X} is a neighborhood of the origin such that (6.22), (6.24), and $\hat{f}(x) \in \hat{X}$ hold for all $x \in \hat{X}$. If $\alpha_6 := \alpha_3 - \alpha_5 \circ \alpha_4$ is positive definite in a neighborhood of the origin, then $x^+ = f(x)$ is asymptotically stable in a neighborhood of the origin.*

Proof. Since \hat{X} contains a neighborhood of the origin, there exists $\varepsilon_1 > 0$ such that $\varepsilon_1 \mathbb{B}^n \subseteq \hat{X}$. Then we can use the lower bound (6.22a) to show $\hat{V}(x) \leq \delta_1 := \alpha_1(|\varepsilon_1|)$ implies $|x| \leq \varepsilon_1$ and therefore $x \in \hat{X}$. Next, since α_6 is positive definite in a neighborhood of the origin, there exists $\varepsilon_2 > 0$ such that $\alpha_6(s) > 0$ for all $s \in (0, \varepsilon_2]$. Again, the lower bound (6.22a) shows $\hat{V}(x) \leq \delta_2 := \alpha_1(\varepsilon_2)$ implies $|x| \leq \varepsilon_2$, and therefore $\alpha_6(|x|) \geq 0$ with equality only at the origin.

Let $\delta := \min \{ \delta_1, \delta_2 \}$ and $X := \text{lev}_\delta \hat{V}$. The upper bound (6.23) gives that $|x| \leq \alpha_2^{-1}(\delta)$ implies $\hat{V}(x) \leq \delta$ and $x \in X$, so X is a neighborhood of the origin. Next, we have, for each

$x \in X$,

$$\hat{V}(f(x)) - \hat{V}(x) \leq -\alpha_3(|x|) + \alpha_5(\alpha_4(|x|)) = -\alpha_6(|x|).$$

But the above inequality means, for each $x \in X$, that $\hat{V}(f(x)) \leq \delta$ since $\alpha_6(|x|) \geq 0$ and $\hat{V}(x) \leq \delta$, so $f(x) \in X$ and X is positive invariant. Finally, \hat{V} is a Lyapunov function on X , so $x^+ = f(x)$ is asymptotically stable on X . \square

6.3 Unconstrained linear control of nonlinear systems

Theorems 6.5 and 6.7 can readily be applied to control of nonlinear systems via linear approximations. In this section, we consider using LQR, LQG, and (constraint-free) offset-free MPC to control nonlinear systems.

6.3.1 Linear quadratic regulator

Consider the nonlinear control system

$$x^+ = f(x, u) \tag{6.26}$$

and the linear system

$$x^+ = Ax + Bu \tag{6.27}$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is continuous, $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and (A, B) stabilizable. Then there exists $K \in \mathbb{R}^{m \times n}$ such that $A_K := A + BK$ is Schur stable, and the closed-loop system $x^+ = A_K x$ is stable. Moreover, for each positive definite Q , there exists a positive definite P that uniquely solves

$$P - A_K^\top P A_K = Q. \tag{6.28}$$

With $V(x) := x^\top Px$, we have (6.6) and (6.7) for all $x \in \mathbb{R}^n$, so V is an exponential Lyapunov function for $x^+ = A_K x$.

The following propositions apply Theorems 6.5 and 6.7 to analyze linear control of the nonlinear system (6.26).

Proposition 6.10. *Suppose $f(0, 0) = 0$, f is continuously differentiable at the origin, and (6.27) is stabilizable with $A := (\partial f / \partial x)(0, 0)$ and $B := (\partial f / \partial u)(0, 0)$. For any $K \in \mathbb{R}^{m \times n}$ such that $A_K := A + BK$ is Schur stable, the closed-loop system $x^+ = f(x, Kx)$ is exponentially stable in a neighborhood of the origin.*

Proof. Let $f_{\text{cl}}(x) := f(x, Kx)$. Then $f_{\text{cl}}(0) = 0$ and, by the chain rule,

$$\frac{df_{\text{cl}}}{dx}(0) = \frac{\partial f}{\partial x}(0, 0) + \frac{\partial f}{\partial u}(0, 0)K = A_K$$

which is Schur stable by assumption. Then $x^+ = A_K x$ is a stable linearization of $x^+ = f_{\text{cl}}(x) = f(x, Kx)$, and the result follows by Theorem 6.5. \square

Proposition 6.11. *Suppose $f(0, 0) = 0$, f is continuously differentiable at the origin, and (6.27) is stabilizable with $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$. Let $\bar{A} := (\partial f / \partial x)(0, 0)$ and $\bar{B} := (\partial f / \partial u)(0, 0)$, $K \in \mathbb{R}^{m \times n}$ such that $A_K := A + BK$ is Schur stable, and $\bar{A}_K := \bar{A} + \bar{B}K$. If $\gamma_0 := \|A_K - \bar{A}_K\| < \gamma_1 := -\|A_K\| + \sqrt{\|A_K\|^2 + \|L_{A_K}^{-1}\|^{-1}}$, then $x^+ = f(x, Kx)$ is exponentially stable in a neighborhood of the origin.*

Proof. Let $f_{\text{cl}}(x) := f(x, Kx)$. Then $f_{\text{cl}}(0) = 0$ and $(df_{\text{cl}}/dx)(0) = \bar{A}_K$ as in the proof of Proposition 6.10. Let Q be positive definite and denote the unique positive definite solution to (6.28) by P . Then we have (6.6) and (6.7), $|A_K x| \leq \|A_K\| |x|$, and

$$\begin{aligned} |V(f_{\text{cl}}(x)) - V(A_K x)| &= |2[f_{\text{cl}}(x) - A_K x]^\top P A_K x + |f_{\text{cl}}(x) - A_K x|_P^2| \\ &\leq 2\|P\| \|A_K x\| |f_{\text{cl}}(x) - A_K x| + \|P\| |f_{\text{cl}}(x) - A_K x|^2 \end{aligned}$$

for all $x \in \mathbb{R}^n$. By Theorem 6.7, if

$$\gamma_0 < \tilde{\gamma}_1(Q, P) := \frac{\underline{\sigma}(Q)/\|P\|}{\|A_K\| + \sqrt{\|A_K\|^2 + \underline{\sigma}(Q)/\|P\|}}$$

then $x^+ = f(x, Kx)$ is exponentially stable in a neighborhood of the origin. But, as discussed in Remark 6.6, $\|P\| = \|L_A^{-1}(Q)\| \geq \underline{\sigma}(Q)\|L_A^{-1}\|$ with equality when $Q = I$ (Gahinet et al., 1990). Therefore $\tilde{\gamma}_1(Q, P) \leq \gamma_1$, with equality when $Q = I$, so we have stability at the origin whenever $\gamma_0 < \gamma_1$. \square

Remark 6.12. While we could have started the proof of Proposition 6.11 with $Q = I$, it would not have been clear that this choice maximizes the allowed error margin γ_1 . It is also worth pointing out that $K \in \mathbb{R}^{m \times n}$ can be designed to maximize γ_1 . If A_K is diagonalizable, then $\|L_{A_K}^{-1}\| \leq \kappa^2(V)/(1 - \rho^2(A_K))$, where $A_K = V\Lambda V^{-1}$ is the eigenvalue decomposition of A_K (Gahinet et al., 1990, Thm. 5.4). Since the function $\psi(a, b) := -a + \sqrt{a^2 + b^{-1}}$ is decreasing in both a and b , for all $a, b > 0$, it is clear that maximization of γ_1 requires simultaneous minimization of the spectral norm $\|A_K\|$, spectral radius $\rho(A_K)$, and condition number $\kappa(V)$.

6.3.2 Linear quadratic Gaussian regulator

In the absence of full-state observation, state estimation becomes necessary. Consider the partially-observed nonlinear system

$$x^+ = f(x, u), \quad y = h(x) \quad (6.29)$$

and its linear approximation

$$x^+ = Ax + Bu, \quad y = Cx. \quad (6.30)$$

The state is estimated with a linear filter

$$\hat{x}^+ = A\hat{x} + Bu + L(y - C\hat{x}) \quad (6.31)$$

where the observer gain L is chosen such that $A_L := A - LC$ is Schur stable. To stabilize the system, we use the feedback law $u = K\hat{x}$ where K is chosen such that $A_K := A + BK$ is Schur stable. With the error $e := x - \hat{x}$, the joint linear estimate-error system is

$$\begin{bmatrix} \hat{x} \\ e \end{bmatrix}^+ = \begin{bmatrix} A_K & LC \\ 0 & A_L \end{bmatrix} \begin{bmatrix} \hat{x} \\ e \end{bmatrix} \quad (6.32)$$

which is exponentially stable because

$$\mathbf{A} := \begin{bmatrix} A_K & LC \\ 0 & A_L \end{bmatrix}$$

is Schur stable.⁴ Moreover, for each positive definite \mathbf{Q} , there is a unique positive definite solution \mathbf{P} to the discrete Lyapunov equation

$$\mathbf{P} - \mathbf{A}^\top \mathbf{P} \mathbf{A} = \mathbf{Q}. \quad (6.33)$$

and $V(\hat{x}, e) := \begin{bmatrix} \hat{x}^\top & e^\top \end{bmatrix} \mathbf{P} \begin{bmatrix} \hat{x}^\top & e^\top \end{bmatrix}^\top$ is an exponential Lyapunov function for the linear estimate-error system (6.32).

It is likely the plant and model states are poor approximations of each other, even when the input-output behavior is similar. To align the plant and model states, we consider invertible linear transformations of the plant state, i.e., $\tilde{x} := Tx$ where $T \in \mathbb{R}^{n \times n}$ is nonsingular, which

⁴A block triangular matrix with Schur stable diagonal blocks is itself Schur stable.

gives

$$\tilde{x}^+ = f_T(\tilde{x}, u) := Tf(T^{-1}\tilde{x}, u), \quad y = h_T(\tilde{x}) := h(T^{-1}\tilde{x}). \quad (6.34)$$

Combining (6.31) and (6.34) gives the joint estimate-error system

$$\begin{bmatrix} \hat{x}^+ \\ e^+ \end{bmatrix} = \mathbf{f}_T(\hat{x}, e) := \begin{bmatrix} A_K \hat{x} + L(h_T(\hat{x} + e) - C\hat{x}) \\ f_T(\hat{x} + e, K\hat{x}) - A_K \hat{x} - L(h_T(\hat{x} + e) - C\hat{x}) \end{bmatrix} \quad (6.35)$$

where $e := \tilde{x} - \hat{x}$. Then $\mathbf{f}_T(0, 0) = (0, 0)$, and, by the chain rule, $(d\mathbf{f}_T/d(\hat{x}, e))(0, 0) = \overline{\mathbf{A}}(T)$

where

$$\overline{\mathbf{A}}(T) := \begin{bmatrix} A_K + L\overline{C}T^{-1} - LC & L\overline{C}T^{-1} \\ T\overline{A}T^{-1} + T\overline{B}K - A_K - L\overline{C}T^{-1} + LC & T\overline{A}T^{-1} - L\overline{C}T^{-1} \end{bmatrix} \quad (6.36)$$

and $\overline{A} := (\partial f/\partial x)(0, 0)$, $\overline{B} := (\partial f/\partial u)(0, 0)$, and $\overline{C} := (dh/dx)(0)$.

The following two corollaries specialize Theorems 6.5 and 6.7 to consider stability of the estimate-error system (6.35) under certain similarity transformations $T \in \mathbb{R}^{n \times n}$.

Proposition 6.13. *Suppose the following conditions hold.*

- (a) $f(0, 0) = 0$, $h(0) = 0$, f and h are continuously differentiable at the origin.
- (b) The system (6.30) is stabilizable and detectable.
- (c) There exists $T \in \mathbb{R}^{n \times n}$ nonsingular such that $A = T\overline{A}T^{-1}$, $B = T\overline{B}$, and $C = \overline{C}T^{-1}$ (with $(\overline{A}, \overline{B}, \overline{C})$ as defined above).

Then there exist $K \in \mathbb{R}^{m \times n}$ and $L \in \mathbb{R}^{n \times p}$ such that $A_K := A + BK$ and $A_L := A - LC$ are Schur stable, and the closed-loop system (6.35) is exponentially stable in a neighborhood of the origin.

Proof. First, there exist a pair of gains (K, L) for which A_K and A_L are Schur stable due to stabilizability and detectability of the linearization. From (6.35) and (6.36) and condition

(c), we have $\mathbf{f}_T(0, 0) = (0, 0)$ and $(d\mathbf{f}_T/d(\hat{x}, e))(0, 0) = \mathbf{A}$ (as defined above) which is Schur stable. Then (6.32) is a stable linearization of (6.35), and the result follows by Theorem 6.5. \square

Proposition 6.14. *Suppose the following conditions hold.*

(a) $f(0, 0) = 0, h(0) = 0, f$ and h are continuously differentiable at the origin.

(b) (A, B) and (A, C) are stabilizable and detectable, respectively.

Then there exist $K \in \mathbb{R}^{m \times n}$ and $L \in \mathbb{R}^{n \times p}$ such that $A_K := A + BK$ and $A_L := A - LC$ are Schur stable, and if

$$\gamma_0 := \inf_{T \in \mathbb{R}^{n \times n}, \det(T) \neq 0} \|\mathbf{A} - \bar{\mathbf{A}}(T)\| < \gamma_1 := -\|\mathbf{A}\| + \sqrt{\|\mathbf{A}\|^2 + \|\mathbf{L}_{\mathbf{A}}^{-1}\|^{-1}}$$

(with $(\mathbf{A}, \bar{\mathbf{A}}(T))$ as defined above) then there exists $T \in \mathbb{R}^{n \times n}$ nonsingular such that (6.35) is exponentially stable in a neighborhood of the origin.

Proof. First, Schur stability of A_K and A_L for some pair (K, L) follow from stabilizability and detectability of the linear approximation. Let $\gamma^* \in (\gamma_0, \gamma_1)$ and $T \in \mathbb{R}^{n \times n}$ be a nonsingular transformation such that $\gamma^* \geq \|\mathbf{A} - \bar{\mathbf{A}}(T)\|$.

Let $\mathbf{x} := (\hat{x}, e)$ throughout. As in the proof of Proposition 6.13, we have, from (6.35) and (6.36), that $\mathbf{f}_T(0) = (0)$ and $(d\mathbf{f}_T/d\mathbf{x})(0) = \bar{\mathbf{A}}(T)$. Let (\mathbf{Q}, \mathbf{P}) be positive definite matrices satisfying (6.33). Then, with the candidate Lyapunov function $V(\mathbf{x}) := \mathbf{x}^\top \mathbf{P} \mathbf{x}$ for each $\mathbf{x} \in \mathbb{R}^{2n}$, we have $|\mathbf{A}\mathbf{x}| \leq \|\mathbf{A}\|\|\mathbf{x}\|$ and

$$\underline{\sigma}(\mathbf{P})|\mathbf{x}|^2 \leq V(\mathbf{x}) \leq \bar{\sigma}(\mathbf{P})|\mathbf{x}|^2,$$

$$|V(\mathbf{f}_T(\mathbf{x})) - V(\mathbf{A}\mathbf{x})| \leq 2\|\mathbf{P}\|\|\mathbf{A}\mathbf{x}\|\|\mathbf{f}_T(\mathbf{x}) - \mathbf{A}\mathbf{x}\| + \|\mathbf{P}\|\|\mathbf{f}_T(\mathbf{x}) - \mathbf{A}\mathbf{x}\|^2$$

as in the proof of Proposition 6.10. By Theorem 6.7, if

$$\|\mathbf{A} - \bar{\mathbf{A}}(T)\| < \tilde{\gamma}_1(\mathbf{Q}) := -\|\mathbf{A}\| + \sqrt{\|\mathbf{A}\|^2 + \underline{\sigma}(\mathbf{Q})/\|\mathbf{P}\|}$$

then $\mathbf{x}^+ = \mathbf{f}_T(\mathbf{x})$ is exponentially stable in a neighborhood of the origin. But

$$\|\mathbf{P}\| = \|\mathbf{L}_{\mathbf{A}}^{-1}(\mathbf{Q})\| \geq \underline{\sigma}(\mathbf{Q})\|\mathbf{L}_{\mathbf{A}}^{-1}\|$$

with equality when $\mathbf{Q} = I$ (Gahinet et al., 1990). Therefore $\tilde{\gamma}_1(\mathbf{Q}) \leq \gamma_1$ with equality when $\mathbf{Q} = I$, and we have stability at the origin since $\gamma_0 \leq \|\mathbf{A} - \overline{\mathbf{A}}(T)\| \leq \gamma^* < \gamma_1$. \square

Remark 6.15. Both Propositions 6.13 and 6.14 imply that $|\hat{x}(k)| \rightarrow 0$ and $|\tilde{x}(k)| \rightarrow 0$ as $k \rightarrow \infty$. But this means $|x(k)| \rightarrow 0$ because $x = T^{-1}\tilde{x}$, so the plant is also stabilized to the origin. In fact, we have convergence for *any* nonsingular transformation, not just the one chosen in Propositions 6.13 and 6.14. In either proposition, the chosen transformation simply minimizes the different between the plant linearization and approximate linear model linearizations, making it easy to invoke Theorems 6.5 and 6.7.

Remark 6.16. It is unrealistic to expect the plant (6.29) states are adequately linearized as written. Therefore, we should realistically consider all plant realizations

$$x^+ = f_\varphi(x, u) := \varphi(f(\varphi^{-1}(x), u)), \quad y = h_\varphi := h(\varphi^{-1}(x))$$

where $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is any sufficiently smooth⁵ function such that $\varphi(0) = 0$. But with $T := (d\varphi/dx)(0)$, we still have the derivative $(df_T/d(\hat{x}, e))(0, 0) = \overline{\mathbf{A}}(T)$ where (6.36). So to show stabilization in some neighborhood of the origin, it suffices to consider only linear transformations $\varphi(x) := Tx$ where T is nonsingular. While it is outside of the scope of this work, a nonlinear φ satisfying $(d\varphi/dx)(0) = T$ could be fine-tuned to maximize the guaranteed size of the basin of attraction by shrinking the \mathcal{K} -function bound implied by Proposition 6.4.

⁵We conjecture that continuous differentiability at the origin is a sufficient smoothness condition for φ . Then every pair (f_φ, h_φ) is continuously differentiable at the origin so long as the original pair (f, h) is.

6.3.3 Linear offset-free MPC

In most control applications, we select a few input/output signals to track at steady state. We do not necessarily know the steady state at which (6.29) achieves these reference signals, so the controller must correct for disturbances and plant-model mismatch to completely eliminate offset from the reference signals. Consider the steady-state reference

$$r_s = \begin{bmatrix} H_y & H_u \end{bmatrix} \begin{bmatrix} h(x_s) \\ u_s \end{bmatrix} \quad (6.37a)$$

where $r_s \in \mathbb{R}^m$ is the reference and $(x_s, u_s) \in \mathbb{R}^n \times \mathbb{R}^m$ are steady-state targets satisfying

$$x_s = f(x_s, u_s). \quad (6.37b)$$

Typically the columns of $\begin{bmatrix} H_y & H_u \end{bmatrix}$ are chosen as elementary vectors so that r_s represents a few of the system inputs and outputs. To ensure the setpoint r_s is reachable, we make the following assumption.

Assumption 6.17. The setpoint r_s and input u have the same dimension, i.e., $n_r = m$. Given the setpoint $r_s \in \mathbb{R}^m$, there exists a unique steady state $(x_s, u_s) \in \mathbb{R}^n \times \mathbb{R}^m$ satisfying (6.37).

The nonlinear plant (6.29) is approximated by the augmented linear system

$$x^+ = Ax + B_d d + Bu \quad (6.38a)$$

$$d^+ = d \quad (6.38b)$$

$$y = Cx + C_d d \quad (6.38c)$$

where $d \in \mathbb{R}^{n_d}$ is a disturbance state intended to correct of the effect of plant-model mismatch. The following assumption guarantees, for the augmented system (6.38), detectability of the states and disturbances, and reachability of the setpoints.

Assumption 6.18. The unaugmented system (6.30) is detectable, $n_r = m$, $n_d = p$, and the following rank conditions hold:

$$\text{rank} \begin{bmatrix} A - I & B_d \\ C & C_d \end{bmatrix} = n + p, \quad (6.39)$$

$$\text{rank} \begin{bmatrix} A - I & B \\ H_y C & H_u \end{bmatrix} = n + m. \quad (6.40)$$

In (Pannocchia and Rawlings, 2003, Lem. 1) it is shown that detectability of (6.38) requires detectability of the unaugmented system (6.30) and a rank condition. We restate this lemma below, and note that it shows Assumption 6.18 implies detectability of (6.38).

Lemma 6.19 ((Pannocchia and Rawlings, 2003, Lem. 1)). *The augmented system (6.38) is detectable if and only if the unaugmented system (6.30) is detectable and (6.39) holds.*

Under Assumption 6.18, we have the existence of a unique steady state $(x_s, u_s) \in \mathbb{R}^n \times \mathbb{R}^m$ that solve

$$\begin{bmatrix} A - I & B \\ H_y C & H_u \end{bmatrix} \begin{bmatrix} x_s \\ u_s \end{bmatrix} = \begin{bmatrix} -B_d d \\ -H_y C_d d + r_s \end{bmatrix} \quad (6.41)$$

for any disturbance $d \in \mathbb{R}^{n_d}$ and setpoint $r_s \in \mathbb{R}^{n_r}$. Moreover, by Assumption 6.18 and Lemma 6.19, there exist gains (K, L_x, L_d) such that $A_K := A + BK$ and

$$A_L := \begin{bmatrix} A - L_x C & B_d - L_x C_d \\ -L_d C & I - L_d C_d \end{bmatrix} \quad (6.42)$$

are Schur stable. We estimate the state and disturbance (\hat{x}, \hat{d}) with the following filter:

$$\hat{x}^+ = A\hat{x} + B_d\hat{d} + Bu + L_x(y - C\hat{x} - c_d\hat{d}) \quad (6.43a)$$

$$\hat{d}^+ = \hat{d} + L_d(y - C\hat{x} - c_d\hat{d}). \quad (6.43b)$$

To steer (6.38) to the setpoint r_s , linear feedback will not do; instead we require an affine feed-

back that stabilizes the steady state (x_s, u_s) such that (6.41) is satisfied. With the disturbance estimate \hat{d} , we define the steady-state targets as

$$\begin{bmatrix} \hat{x}_t \\ \hat{u}_t \end{bmatrix} := \begin{bmatrix} A - I & B \\ H_y C & H_u \end{bmatrix}^{-1} \begin{bmatrix} -B_d \hat{d} \\ -H_y C_d \hat{d} + r_s \end{bmatrix} = \begin{bmatrix} T_x \\ T_u \end{bmatrix} \hat{d} + \begin{bmatrix} R_x \\ R_u \end{bmatrix} r_s \quad (6.44)$$

where existence of the inverse follows from Assumption 6.17. With these targets, we use the control law

$$u = K(\hat{x} - \hat{x}_t) + \hat{u}_t. \quad (6.45)$$

Combining (6.38) and (6.43)–(6.45) gives the joint system

$$\begin{bmatrix} \Delta \hat{x} \\ e_x \\ e_d \end{bmatrix}^+ = \begin{bmatrix} A_K & | & L_x [C & C_d] \\ 0 & | & A_L \end{bmatrix} \begin{bmatrix} \Delta \hat{x} \\ e_x \\ e_d \end{bmatrix} \quad (6.46)$$

where $\Delta \hat{x} := \hat{x} - \hat{x}_t$, $e_x := x - \hat{x}$ and $e_d := d - \hat{d}$. Similarly to the LQG, we have that (6.46) is stable because

$$\mathbf{A} := \begin{bmatrix} A_K & | & L_x [C & C_d] \\ 0 & | & A_L \end{bmatrix}$$

is Schur stable.

As in Section 6.3.2, the plant and model states are likely unaligned even when the input-output behavior is similar. This time, even the steady states that reach the setpoint r_s can be different. We define $(x_s, u_s) \in \mathbb{R}^n \times \mathbb{R}^m$ as the unique steady state such that (6.37) holds, and define $(\hat{x}_s, \hat{d}_s) \in \mathbb{R}^n \times \mathbb{R}^m$ as a steady state satisfying

$$\begin{bmatrix} A - I & B_d \\ C & C_d \end{bmatrix} \begin{bmatrix} \hat{x}_s \\ \hat{d}_s \end{bmatrix} = \begin{bmatrix} -B u_s \\ -h(x_s) + r_s \end{bmatrix} \quad (6.47)$$

which exist and are unique under Assumption 6.18. Since (\hat{x}_s, \hat{d}_s) depend on (x_s, u_s) , there is no way to know the state and disturbance (model) setpoints (\hat{x}_s, \hat{d}_s) beforehand. Instead, the

disturbance estimate \hat{d} is intended to “integrate” the model error and converge to \hat{d}_s .

Consider affine transformations of the form $\tilde{x} := T(x - x_s) + \hat{x}_s$. Under such a transformation, the steady state for the new system is always (\hat{x}_s, u_s) . Define the family of (6.29) realizations

$$\tilde{x}^+ = f_T(\tilde{x}, u) := T[(f(T^{-1}(\tilde{x} - \hat{x}_s) + x_s), u) - x_s] + \hat{x}_s \quad (6.48a)$$

$$y = h_T(\tilde{x}) := h(T^{-1}(\tilde{x} - \hat{x}_s) + x_s). \quad (6.48b)$$

Let $\Delta\hat{x} := \hat{x} - \hat{x}_t$, $e_x := \tilde{x} - \hat{x}$, and $e_d := \hat{d}_s - \hat{d}$ denote the state estimate in deviation variables, the state estimate error, and the disturbance estimate error, respectively. Then we have the closed-loop system

$$\begin{bmatrix} \Delta\hat{x} \\ e_x \\ e_d \end{bmatrix}^+ = \mathbf{f}_T(\Delta\hat{x}, e_x, e_d) := \begin{bmatrix} A_K\Delta\hat{x} + L_x\varepsilon \\ f_T(\tilde{x}, u) - \hat{x}_t - A_K\Delta\hat{x} - L_x\varepsilon \\ e_d - L_d\varepsilon \end{bmatrix} \quad (6.49a)$$

with

$$\hat{x}_t = \hat{x}_s - T_x e_d \quad (6.49b)$$

$$\tilde{x} = \hat{x}_s - T_x e_d + \Delta\hat{x} + e_x \quad (6.49c)$$

$$u = u_s + K\Delta\hat{x} - T_u e_d \quad (6.49d)$$

$$\varepsilon = h_T(\tilde{x}) - h(x_s) - C\Delta\hat{x} + (CT_x + C_d)e_d \quad (6.49e)$$

and we have used the facts $h(x_s) = C\hat{x}_s + C_d\hat{d}_s$ and

$$\begin{bmatrix} \hat{x}_s - \hat{x}_t \\ u_s - \hat{u}_t \end{bmatrix} = \begin{bmatrix} T_x \\ T_u \end{bmatrix} (\hat{d}_s - \hat{d}) = \begin{bmatrix} T_x \\ T_u \end{bmatrix} e_d. \quad (6.50)$$

Then $\mathbf{f}_T(0, 0, 0) = (0, 0, 0)$ by Assumption 6.17, and by liberal application of the chain rule,

it can be shown that

$$\frac{d\mathbf{f}_T}{d(\Delta\hat{x}, e_x, e_d)}(0, 0, 0) = \bar{\mathbf{A}}(T) \quad (6.51)$$

where

$$\bar{\mathbf{A}}(T) := \mathbf{A} + \begin{bmatrix} L_x\Delta C(T) & L_x\Delta C(T) & -L_x\Delta C(T)T_x \\ \Delta A_K(T) - L_x\Delta C(T) & \Delta A_L(T) & \Delta A(T)T_x \\ \Delta C(T) & -L_d\Delta C(T) & \Delta C(T)T_x \end{bmatrix} \quad (6.52a)$$

and

$$\Delta A(T) := T\bar{A}T^{-1} - A, \quad \Delta A_K(T) := \Delta A(T) + \Delta B(T)K, \quad (6.52b)$$

$$\Delta B(T) := T\bar{B} - B, \quad \Delta A_L(T) := \Delta A(T) - L\Delta C(T), \quad (6.52c)$$

$$\Delta C(T) := \bar{C}T^{-1} - C \quad (6.52d)$$

are the transform-dependent model parameters, and

$$\bar{A} := \frac{\partial f}{\partial x}(x_s, u_s), \quad \bar{B} := \frac{\partial f}{\partial u}(x_s, u_s), \quad \bar{C} := \frac{dh}{dx}(x_s) \quad (6.52e)$$

are the derivatives at the setpoint. While the expression (6.52) appears cumbersome, notice that, for any nonsingular T for which $A = T\bar{A}T^{-1}$, $B = T\bar{B}$, and $C = \bar{C}T^{-1}$, the second term drops out entirely and we have $\bar{\mathbf{A}}(T) = \mathbf{A}$, which is Schur stable. Thus, if this similarity transformation holds or is sufficiently approximated, then the origin is stable.

Proposition 6.20. *If the origin of (6.49) is stable, then, for sufficiently small $(\Delta\hat{x}, e_x, e_d)$, we have*

$$\begin{aligned} (x(k), u(k)) &\rightarrow (x_s, u_s) & (\hat{x}(k), \hat{d}(k)) &\rightarrow (\hat{x}_s, \hat{d}_s) \\ (\hat{x}_t(k), \hat{u}_t(k)) &\rightarrow (\hat{x}_s, u_s) & r(k) &\rightarrow r_s \end{aligned}$$

as $k \rightarrow \infty$, where $r(k) := H_y h(x(k)) + H_u u(k)$.

Proof. Stability of the origin and $(\Delta\hat{x}, e_x, e_d)$ in its basin of attraction implies $\Delta\hat{x}(k) \rightarrow 0$, $e_x(k) \rightarrow 0$, and $e_d(k) \rightarrow 0$. The last limit implies $\hat{d}(k) \rightarrow \hat{d}_s$, so $(\hat{x}_t(k), \hat{u}_t(k)) \rightarrow (\hat{x}_s, u_s)$ by (6.50). Moreover, $\hat{x}(k) = \Delta\hat{x}(k) + \hat{x}_t(k) \rightarrow \hat{x}_s$ and $\tilde{x}(k) = \hat{x}(k) + e_x(k) \rightarrow \hat{x}_s$ by the definitions of $(\Delta\hat{x}, e_x)$, and $x(k) = T^{-1}(\tilde{x}(k) - \hat{x}_s) + x_s \rightarrow x_s$ by the affine transformation. Finally, by continuity of h at x_s , we have $r(k) \rightarrow H_y h(x_s) + H_u u_s = r_s$. \square

The following two propositions apply Theorems 6.5 and 6.7 to stability of the setpoint tracking system (6.49).

Proposition 6.21. *Let $r_s \in \mathbb{R}^{n_r}$, suppose Assumptions 6.17 and 6.18 hold, and let $T \in \mathbb{R}^{n \times n}$ be a nonsingular matrix for which $A = T\bar{A}T^{-1}$, $B = T\bar{B}$, and $C = \bar{C}T^{-1}$ (where $(\bar{A}, \bar{B}, \bar{C})$ are defined as above). Then there exist gains $(K, L_x, L_d) \in \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times p} \times \mathbb{R}^{n_a \times p}$ such that $A_K := A + BK$ and (6.42) are Schur stable, and (6.49) is exponentially stable in a neighborhood of the origin.*

Proof. By Assumption 6.18 and Lemma 6.19, the gains (K, L_x, L_d) such that A_K and A_L are Schur stable exist. Moreover, by the assumptions and the discussion above, we have the closed-loop system (6.49), for which $\mathbf{f}_T(0, 0, 0) = (0, 0, 0)$ and by the chain rule (6.51). But, by assumption $\bar{\mathbf{A}}(T) = \mathbf{A}$, so (6.46) is a linearization of (6.49), and moreover, (6.46) is stable, so (6.49) must be exponentially stable in a neighborhood of the origin by Theorem 6.5. \square

Proposition 6.22. *Let $r_s \in \mathbb{R}^{n_r}$, suppose Assumptions 6.17 and 6.18 hold. Then there exist gains $(K, L_x, L_d) \in \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times p} \times \mathbb{R}^{n_a \times p}$ such that $A_K := A + BK$ and (6.42) are Schur stable, and if*

$$\gamma_0 := \inf_{T \in \mathbb{R}^{n \times n}, \det(T) \neq 0} \|\mathbf{A} - \bar{\mathbf{A}}(T)\| < \gamma_1 := -\|\mathbf{A}\| + \sqrt{\|\mathbf{A}\|^2 + \|\mathbf{L}_{\mathbf{A}}^{-1}\|^{-1}}$$

(with $(\mathbf{A}, \bar{\mathbf{A}}(T))$ as defined above) then there exists a nonsingular matrix $T \in \mathbb{R}^{n \times n}$ such that (6.49) is exponentially stable in a neighborhood of the origin.

Proof. Schur stability of A_K and (6.42) for some (K, L_x, L_d) follow from Assumptions 6.17 and 6.18. Let $\gamma^* \in (\gamma_0, \gamma_1)$ and $T \in \mathbb{R}^{n \times n}$ be a nonsingular transformation such that $\gamma^* \geq \|\mathbf{A} - \bar{\mathbf{A}}(T)\|$. The remainder of the proof follows identically to that of Proposition 6.14 with the slight modification that $\mathbf{x} := (\Delta \hat{x}, e_x, e_d)$ throughout. \square

6.4 Other considerations

Throughout we have assumed the plant and model have the same state dimension. However, it is important to note the prior results do not require this to be true. Consider the plant

$$x_p^+ = f(x_p, u) \qquad y = h(x_p) \qquad (6.53)$$

where $x_p \in \mathbb{R}^{n_p}$ is the plant state.

Under-modeled state Suppose $n < n_p$. Then we can take

$$\tilde{x} := \begin{bmatrix} x \\ 0 \end{bmatrix} \in \mathbb{R}^{n_p}, \quad \tilde{A} := \begin{bmatrix} A & \\ & 0 \end{bmatrix} \in \mathbb{R}^{n_p \times n_p}, \quad \tilde{B} := \begin{bmatrix} B \\ 0 \end{bmatrix} \in \mathbb{R}^{n_p \times m}, \quad \tilde{C} := [C \ 0] \in \mathbb{R}^{p \times n_p}$$

and consider the extended linear model

$$\tilde{x}^+ = \tilde{A}\tilde{x} + \tilde{B}u, \qquad y = \tilde{C}\tilde{x}.$$

Over-modeled state Suppose $n > n_p$. Then we can take

$$\tilde{x} := \begin{bmatrix} x_p \\ 0 \end{bmatrix} \in \mathbb{R}^n, \quad \tilde{f}(\tilde{x}, u) := \begin{bmatrix} f(x_p, u) \\ 0 \end{bmatrix} \in \mathbb{R}^n, \quad \tilde{h}(\tilde{x}) := h(x_p)$$

and consider the extended plant

$$\tilde{x}^+ = \tilde{f}(\tilde{x}, u), \quad y = \tilde{h}(\tilde{x}).$$

Appendices

6.A Derivatives and \mathcal{K} -functions

In this appendix, we prove Proposition 6.4. It is a direct extension of (Rawlings and Risbeck, 2015, Props. 5, 13), who show the equivalence of ε - δ and \mathcal{K} -function definitions of continuity. The results of Rawlings and Risbeck (2015) assume existence of the function value at the limit point, which does not hold for the following definition of the Jacobian:

$$\lim_{|h| \rightarrow 0} \frac{|f(x_0 + h) - f(x_0) - Jh|}{|h|} = 0$$

where $J := \frac{df}{dx}(x_0)$ (Rudin, 1976, Defn. 9.11). Therefore it is necessary to modify the results of Rawlings and Risbeck (2015) to accommodate Proposition 6.4. First, however, we borrow the following \mathcal{K} -function lower bounding result from (Rawlings and Risbeck, 2015, Prop. 4).

Proposition 6.23 ((Rawlings and Risbeck, 2015, Prop. 4)). *Let $\delta : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ be a nondecreasing function. Then there exists $\alpha \in \mathcal{K}$ such that $\alpha(\varepsilon) \leq \delta(\varepsilon)$ for all $\varepsilon > 0$.*

Then we can show the equivalence of ε - δ definition of the limit and a \mathcal{K} -function overbound on an *excluded* neighborhood of the limit point.

Proposition 6.24. *The function $V : \mathbb{R}^n \rightarrow \mathbb{R}^m$ has the limit V_0 at x_0 if and only if there exists $b > 0$ and $\gamma \in \mathcal{K}$ such that*

$$|V(x) - V_0| \leq \gamma(|x - x_0|) \tag{6.54}$$

for all $0 < |x - x_0| \leq b$.

Proof. Without loss of generality, assume $x_0 = 0$ and $V_0 = 0$.

(\Rightarrow) Suppose the limit exists. Then for each $\varepsilon > 0$, there exists $\delta(\varepsilon) > 0$ such that $|V(x)| \leq \varepsilon$ for all $0 < |x| \leq \delta(\varepsilon)$. Clearly $\delta(\varepsilon)$ can be made a nondecreasing function (Rawlings and Risbeck, 2015, Prop. 11), so by Proposition 6.23 there exists $\alpha \in \mathcal{K}$ that lower bounds δ , i.e., $\alpha(\varepsilon) \leq \delta(\varepsilon)$ for all $\varepsilon > 0$. Let $\gamma := \alpha^{-1}$ denote the inverse of α on $[0, b]$ where $b := \sup_{\varepsilon > 0} \delta(\varepsilon)$. For each $\varepsilon > 0$, choose $x \in \mathbb{R}^n$ such that $|x| \leq \alpha(\varepsilon)$. Then $|x| \leq \delta(\varepsilon)$ by construction of α , and $|V(x)| \leq \varepsilon = \alpha^{-1}(|x|)$.

(\Leftarrow) Suppose $b > 0$ and $\gamma \in \mathcal{K}$ exist such that (6.54) hold for all $|x| \leq b$. Let $\delta := \gamma^{-1}$ denote the inverse of γ on $[0, b]$. Then, for each $\varepsilon > 0$, we have that $|x| \leq \delta(\varepsilon) = \gamma^{-1}(\varepsilon)$ implies $|V(x)| \leq \gamma(|x|) \leq \gamma(\gamma^{-1}(\varepsilon)) = \varepsilon$. \square

Proposition 6.24 differs from (Rawlings and Risbeck, 2015, Props. 5, 13) in that we do not require $V(x_0)$ to equal V_0 , and in fact, we do not require a finite value for V at x_0 at all. Finally, we can prove Proposition 6.4 using Proposition 6.24.

Proof of Proposition 6.4. By the definition of the derivative, $|f(x) - Jx|/|x| \rightarrow 0$ as $|x| \rightarrow 0$. By Proposition 6.24, there exist $b > 0$ and $\gamma \in \mathcal{K}$ satisfying (6.8) for all $0 < |x| \leq b$. \square

Chapter 7

Stability of model predictive control despite plant-model mismatch

Plant-model mismatch is an ever-present challenge in model predictive control (MPC) practice. In industrial implementations, the main driver of MPC performance is model quality (Qin and Badgwell, 2003; Darby and Nikolaou, 2012). There has been recent progress on improving model quality and MPC performance through disturbance modeling and estimator tuning (Kuntz and Rawlings, 2022, 2024b; Simpson et al., 2024), simultaneous state and parameter estimation (Baumgärtner et al., 2022; Muntwiler et al., 2023; Schiller and Müller, 2023), and even direct data-driven MPC design (Berberich et al., 2021, 2022a,b), to name a few methods. However, there is not yet a sharp theoretical understanding of the robustness of MPC to plant-model mismatch.

Before discussing MPC robustness, let us first define *robustness*. In the stability literature, *robust asymptotic stability* has been used to refer to both (i) input-to-state stability (ISS) and (ii) asymptotic stability despite disturbances. To avoid confusion, we reserve the term *robust asymptotic stability* for (i) and use *strong asymptotic stability* to refer to (ii). The latter term is borrowed from the differential inclusion literature (Clarke et al., 1998) (see Jiang and Wang (2001); Kellett and Teel (2005) for discrete-time definitions). Note that some authors use the term *uniform asymptotic stability* to refer to (ii) (Jiang and Wang, 2001), but we do not use

this term to avoid confusion with the time-varying case. When such properties are given by a nominal MPC,¹ we call it *inherently robust* or *inherently strongly stabilizing*. Robust and strong exponential stability are defined similarly.

It is well-known that MPC is stabilizing under certain assumptions on the terminal ingredients (Rawlings et al., 2020, Ch. 2). To achieve robust stability in the presence of disturbances (parameter errors, estimation errors, exogenous perturbations), a disturbance model can be included. The simplest manner of handling disturbances is with feedback. For MPC this would require future knowledge of the disturbance trajectory, or at least a forecast of it, to implement the controller. While this is a strong requirement, it would confer strong stability rather than robust stability. Alternatively, a disturbance model may be included. Several MPC variants include disturbance models in their design, such as offset-free (Pannocchia et al., 2015), stochastic (McAllister, 2022), tube-based (Rawlings et al., 2020, Ch. 3), and min-max MPC (Limon et al., 2006). For a survey of these methods, see (Rawlings et al., 2020, Ch. 1, 3).

Even in the absence of a disturbance model, a wide range of nominal MPC designs are inherently robust to disturbances. Continuity of the control law was first proven to be a sufficient condition for inherent robustness (De Nicolao et al., 1996; Scokaert et al., 1997). Later, Grimm et al. (2004) proved continuity of the optimal value function is sufficient for inherent robustness, and stated MPC examples with discontinuous optimal value functions that are nominally stable but otherwise not robust to disturbances. A special class of time-varying terminal constraints were proven to confer robust stability to nominal MPC by Grimm et al. (2007), and to suboptimal MPC by Lazar and Heemels (2009). In Pannocchia et al. (2011); Allan et al. (2017), the inherent robustness of optimal and suboptimal MPC, using a class of time-invariant terminal constraints, was proven. With the same terminal constraints, the

¹By *nominal* MPC, we mean any MPC designed without a disturbance model, possibly admitting parameter errors. This includes not only standard nonlinear MPC, but also suboptimal, offset-free, and (some) data-driven MPC.

inherent stochastic robustness (in probability, expectation, and distribution) of nominal MPC was shown by McAllister and Rawlings (2022b,a, 2024). Lastly, direct data-driven MPC was shown to be inherently robust to noisy data by Berberich et al. (2022a).

If the origin remains a steady state under mismatch (e.g., for some kinematic and inventory problems), we might expect strong asymptotic stability. In unconstrained linear optimal control problems (LQR/LQG), the margin of stability (maximum perturbation to the open-loop gain that still gives a closed-loop system) is always nonzero. However, it is important to note that there is no guaranteed relative value of this margin below which the closed loop is stable, save a few exceptional cases such as a single input, or with diagonally-weighted stage costs (Doyle, 1978; Lehtomaki et al., 1981; Zhang and Fu, 1996). Examples are shown by Doyle (1978); Zhang and Fu (1996) in which arbitrarily small perturbations to the gain matrix destabilize the system. These examples use *multiplicative disturbances* that, while persistent in the aforementioned papers, do not need to be time-invariant for the results to hold. The disturbances treated in the MPC literature are typically *additive disturbances* entering the states and measurements (Rawlings et al., 2020, Ch. 3). In the multiplicative case, borrowing from knowledge of linear systems, we should expect strong exponential stability. However, in the additive case, we should expect only robust exponential stability.

To the best of our knowledge, the inherent strong stability of nominal MPC to plant-model mismatch has been discussed by only Santos and Biegler (1999); Santos et al. (2008). In these papers, the magnitude of plant-model mismatch is assumed to be upper bounded by a power law in the magnitude of the state, and for unconstrained systems exhibiting sufficiently small error bounds, the nominal MPC is shown to stabilize the plant to the origin. While the papers consider exact penalty functions for constraint handling, there is no guarantee of recursive feasibility.

In this chapter, we extend the work of Santos et al. (2008) to include input constraints and stabilizing terminal constraints. In Section 7.1, we define the system, state the MPC problem

and assumptions, review nominal MPC stability, and present a motivating example exhibiting both robust and strong stability under plant-model mismatch. In Section 7.2, we formally define robust and strong stability and review the relevant Lyapunov theory. In Section 7.3, we review inherent robustness of MPC. In Section 7.4, we present the main results. For MPC with quadratic costs, it is shown that the closed loop is strongly exponentially stable under (i) a fixed steady state, (ii) a mild differentiability condition, and (iii) the standard MPC assumptions used by Pannocchia et al. (2011); Allan et al. (2017). For MPC with general, positive definite cost functions, we show a *joint \mathcal{K} -function* bound holds on the increase in the optimal value function, but strong stability is only implied if this bound decays sufficiently quickly near the origin. To illustrate the main results, we present three examples in Section 7.5. The first example is a continuous yet nondifferentiable system with a general cost MPC that is not strongly stable, demonstrating inherent strong stability is not a guaranteed property of nonlinear MPC. The second example is a nondifferentiable system for which the quadratic cost MPC is strongly stabilizing. In the third and final example, we use the upright pendulum problem to showcase several types of plant-model mismatch that are covered by the main results, namely, discretization errors, unmodeled dynamics, and errors in estimated parameters. We conclude the chapter and discuss future work in Section 7.6.

7.1 Problem statement

7.1.1 System of interest

Consider the following discrete-time plant:

$$x^+ = f(x, u, \theta) \tag{7.1}$$

where $x \in \mathbb{R}^n$ is the plant state, $u \in \mathbb{R}^m$ is the plant input, and $\theta \in \mathbb{R}^{n_\theta}$ is an *unknown* parameter vector. We denote the parameter estimate by $\hat{\theta} \in \mathbb{R}^{n_\theta}$ and the modeled system by

$$x^+ = f(x, u, \hat{\theta}). \quad (7.2)$$

We assume the parameter estimate is time-invariant, while the parameter vector itself may be time-varying. For simplicity, let $\hat{\theta} = 0$ and denote the model as

$$x^+ = \hat{f}(x, u) := f(x, u, 0). \quad (7.3)$$

Let $\hat{\phi}(k; x, \mathbf{u})$ denote the solution to (7.3) at time k , given an initial state x and a sufficiently long input sequence \mathbf{u} .

In this chapter, we study the behavior of an MPC designed with the model (7.2), but applied to the plant (7.1). We adopt a user-oriented perspective in this analysis: while the model is fixed (e.g., via system identification or prior knowledge), the plant behavior is unknown and possibly changing over time as equipment or the environment changes. Under the assumption $\hat{\theta} = 0$, θ takes the role of an estimate residual. In the language of inherent robustness, the model (7.3) is the nominal system, and the plant (7.1) is the uncertain system.

7.1.2 Nominal MPC and basic assumptions

We consider an MPC problem with control constraints $u \in \mathbb{U} \subseteq \mathbb{R}^m$, a horizon length of $N \in \mathbb{I}_{>0}$, a stage cost $\ell : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$, a terminal constraint $\mathbb{X}_f \subseteq \mathbb{R}^n$, and a terminal cost $V_f : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$. For an initial state $x \in \mathbb{R}^n$, we define the set of admissible (x, \mathbf{u}) pairs

(7.4), admissible input sequences (7.5), and admissible initial states (7.6) by

$$\mathcal{Z}_N := \{ (x, \mathbf{u}) \in \mathbb{R}^n \times \mathbb{U}^N \mid \hat{\phi}(N; x, \mathbf{u}) \in \mathbb{X}_f \} \quad (7.4)$$

$$\mathcal{U}_N(x) := \{ \mathbf{u} \in \mathbb{U}^N \mid (x, \mathbf{u}) \in \mathcal{Z}_N \} \quad (7.5)$$

$$\mathcal{X}_N := \{ x \in \mathbb{R}^n \mid \mathcal{U}_N(x) \text{ is nonempty} \}. \quad (7.6)$$

For each $(x, \mathbf{u}) \in \mathbb{R}^{n+Nm}$, we define the MPC objective by

$$V_N(x, \mathbf{u}) := \sum_{k=0}^{N-1} \ell(\hat{\phi}(k; x, \mathbf{u}), u(k)) + V_f(\hat{\phi}(N; x, \mathbf{u})) \quad (7.7)$$

and for each $x \in \mathcal{X}_N$, we define the MPC problem by

$$V_N^0(x) := \min_{\mathbf{u} \in \mathcal{U}_N(x)} V_N(x, \mathbf{u}). \quad (7.8)$$

According to the convention of Rockafellar and Wets (1998) for infeasible problems, we take

$$V_N^0(x) := \infty \text{ for all } x \notin \mathcal{X}_N.$$

Throughout, we use the standard assumptions for inherent robustness of MPC from Allan et al. (2017).

Assumption 7.1 (Continuity). The functions $f : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}^n$, $\ell : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$, and $V_f : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ are continuous and $\hat{f}(0, 0) = 0$, $\ell(0, 0) = 0$, and $V_f(0) = 0$.

Assumption 7.2 (Constraint properties). The set \mathbb{U} is compact and contains the origin. The set \mathbb{X}_f is defined by $\mathbb{X}_f := \text{lev}_{c_f} V_f$ for some $c_f > 0$.

Assumption 7.3 (Terminal control law). There exists a terminal control law $\kappa_f : \mathbb{X}_f \rightarrow \mathbb{U}$ such that

$$V_f(\hat{f}(x, \kappa_f(x))) \leq V_f(x) - \ell(x, \kappa_f(x)), \quad \forall x \in \mathbb{X}_f.$$

Assumption 7.4 (Stage cost bound). There exists a function $\alpha_1 \in \mathcal{K}_\infty$ such that

$$\ell(x, u) \geq \alpha_1(|(x, u)|), \quad \forall (x, u) \in \mathbb{R}^n \times \mathbb{U}. \quad (7.9)$$

Quadratic stage and terminal costs are of particular interest in this work. Throughout, we call an MPC satisfying the following assumption a *quadratic cost MPC*.

Assumption 7.5 (Quadratic cost). We have

$$\ell(x, u) := |x|_Q^2 + |u|_R^2, \quad V_f(x) := |x|_{P_f}^2 \quad (7.10)$$

for all $(x, u) \in \mathbb{R}^n \times \mathbb{R}^m$ and positive definite Q , R , and P_f .

Under Assumptions 7.1 and 7.2, the existence of solutions to (7.8) follows from (Rawlings et al., 2020, Prop. 2.4). We denote any such solution by $\mathbf{u}^0(x) = (u^0(0; x), \dots, u^0(N-1; x))$, denote the optimal state sequence by $\hat{x}^0(k; x) := \hat{\phi}(k; x, \mathbf{u}^0(x))$ for each $k \in \mathbb{I}_{0:N}$, and define the MPC control law $\kappa_N : \mathcal{X}_N \rightarrow \mathbb{U}$ by $\kappa_N(x) := u^0(0; x)$. It is also useful to define the following suboptimal input sequence:

$$\tilde{\mathbf{u}}(x) := (u^0(1; x), \dots, u^0(N-1; x), \kappa_f(\hat{x}^0(N; x))).$$

Consider the *modeled* closed-loop system

$$x^+ = \hat{f}_c(x) := \hat{f}(x, \kappa_N(x)). \quad (7.11)$$

From Assumptions 7.1 to 7.4, it can be shown $x^+ = \hat{f}_c(x)$ is asymptotically stable in \mathcal{X}_N with the Lyapunov function V_N^0 (Rawlings et al., 2020, Thm. 2.19). For completeness, we include a sketch of the proof in Appendix 7.A.1.

Theorem 7.6 (Thm. 2.19 of Rawlings et al. (2020)). *Suppose Assumptions 7.1 to 7.4 hold. Then*

- (a) \mathcal{X}_N is positive invariant for $x^+ = \hat{f}_c(x)$;
- (b) there exists $\alpha_2 \in \mathcal{K}_\infty$ such that, for each $x \in \mathcal{X}_N$,

$$\alpha_1(|x|) \leq V_N^0(x) \leq \alpha_2(|x|) \quad (7.12a)$$

$$V_N^0(\hat{f}_c(x)) \leq V_N^0(x) - \alpha_1(|x|); \quad (7.12b)$$

- (c) and $x^+ = \hat{f}_c(x)$ is asymptotically stable on \mathcal{X}_N .

Similarly, it is shown in (Rawlings et al., 2020, Sec. 2.5.5) that, under Assumptions 7.1 to 7.3 and 7.5, the quadratic cost MPC *exponentially* stabilizes the closed-loop system (7.11) on any sublevel set of the optimal value function $\mathcal{S} := \text{lev}_\rho V_N^0$. Note that, because V_N^0 is only defined on \mathcal{X}_N , we have $\mathcal{S} \subseteq \mathcal{X}_N$ by the definition of the sublevel set. For completeness, we restate the conclusion of (Rawlings et al., 2020, Sec. 2.5.5) in the theorem below and include a sketch of the proof in Appendix 7.A.1.

Theorem 7.7 (Sec. 2.5.5 of Rawlings et al. (2020)). *Suppose Assumptions 7.1 to 7.3 and 7.5 hold.*

Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. Then

- (a) \mathcal{S} is positive invariant for $x^+ = \hat{f}_c(x)$;
- (b) there exists a constant $c_2 > 0$ such that

$$c_1|x|^2 \leq V_N^0(x) \leq c_2|x|^2 \quad (7.13a)$$

$$V_N^0(\hat{f}_c(x)) \leq V_N^0(x) - c_1|x|^2 \quad (7.13b)$$

for each $x \in \mathcal{S}$, where $c_1 := \underline{\sigma}(Q)$; and

- (c) $x^+ = \hat{f}_c(x)$ is exponentially stable on \mathcal{S} .

To show strong stability of the MPC with mismatch, we eventually require one or both of the following assumptions.

Assumption 7.8 (Steady state). The origin is a steady state, uniformly in $\theta \in \mathbb{R}^{n_\theta}$, i.e., $f(0, 0, \theta) = 0$ for all $\theta \in \mathbb{R}^{n_\theta}$.

Assumption 7.9 (Differentiability). The function $f(\cdot, \cdot, \theta)$ is continuously differentiable for each $\theta \in \mathbb{R}^{n_\theta}$.

Remark 7.10. Assumption 7.8 limits our results to problems where the steady state is known and fixed (e.g., path-planning and inventory problems). If the steady state depends on θ , i.e., $x_s(\theta) = f(x_s(\theta), u_s(\theta), \theta)$, we can still work with deviation variables $(\delta x, \delta u) := (x - x_s(\theta), u - u_s(\theta))$, but (i) we have to estimate the steady-state pair $(x_s(\theta), u_s(\theta))$ (e.g., via an integrating disturbance model (Rawlings et al., 2020, Ch. 1)), and (ii) we only achieve strong stability in the case where the steady-state map is continuous, the parameters are asymptotically constant, and the estimation errors converge.

7.1.3 Motivating example

We close this section with a motivating example exhibiting many types of stability under persistent mismatch. Recall from the introduction we define *robust stability* as an ISS property for parameter errors, and *strong stability* as convergence to the origin despite mismatch. While precise definitions are given in Section 7.2, these informal definitions suffice for the example.

Consider the scalar system

$$x^+ = f(x, u, \theta) := x + (1 + \theta)u. \quad (7.14)$$

The plant (7.14) is a prototypical integrating system, such as a storage tank or vehicle on a track, with an uncertain input gain. As usual the system is modeled with $\hat{\theta} = 0$,

$$x^+ = \hat{f}(x, u) := f(x, u, 0) = x + u. \quad (7.15)$$

We define a nominal MPC with $\mathbb{U} := [-1, 1]$, $\ell(x, u) := (1/2)(x^2 + u^2)$, $V_f(x) := (1/2)x^2$, $\mathbb{X}_f := [-1, 1]$, and $N := 2$. Notice that the terminal set can be reached in $N = 2$ moves if and only if $|x| \leq 3$, so we have the steerable set $\mathcal{X}_2 = [-3, 3]$. *Without* the terminal constraint (i.e., $\mathbb{X}_f = \mathbb{R}$), the optimal control sequence is

$$\mathbf{u}^0(x) = \begin{cases} (-3x/5, -x/5), & |x| \leq 5/3 \\ (-\text{sgn}(x), -x/2 + \text{sgn}(x)/2), & 5/3 < |x| \leq 3 \end{cases}$$

and the control law is $\kappa_2(x) := -\text{sat}(3x/5)$ (Rawlings et al., 2020, p. 104). However, the optimal input sequence gives

$$\hat{x}^0(2; x) = \begin{cases} x/5, & |x| \leq 5/3 \\ x/2 - \text{sgn}(x)/2, & 5/3 < |x| \leq 3 \end{cases}$$

so the terminal constraint $\mathbb{X}_f = [-1, 1]$ is automatically satisfied for all $|x| \leq 3$. Therefore $\kappa_2(x) = -\text{sat}(3x/5)$ is also the control law of the problem *with* the terminal constraint.

In Figure 7.1 we plot contours of the cost difference $\Delta V_2^0(x, \theta) := V_2^0(f(x, \kappa_2(x), \theta)) - V_2^0(x)$, and in Figure 7.2, we plot closed-loop trajectories and the cost difference $\Delta V_2^0(\cdot, \theta)$ for several values of θ . The system is strongly stable for all $-1 < \theta < 7/3$ as the cost difference is negative definite. When $\theta < -1$, the entire cost difference curve is positive definite, so the trajectories become unbounded. This is because the disturbance cancels out the effect of the controller and drives the system in the opposite direction. On the other hand, when $\theta > 7/3$, the cost difference curve is only positive definite near the origin, but negative elsewhere, so the trajectories remain bounded for all time, although they do not converge to the origin. In this case, high parameter values push the system in the same direction as the input, and input saturation moderates the effect of overshoot at high parameter values. We point out the

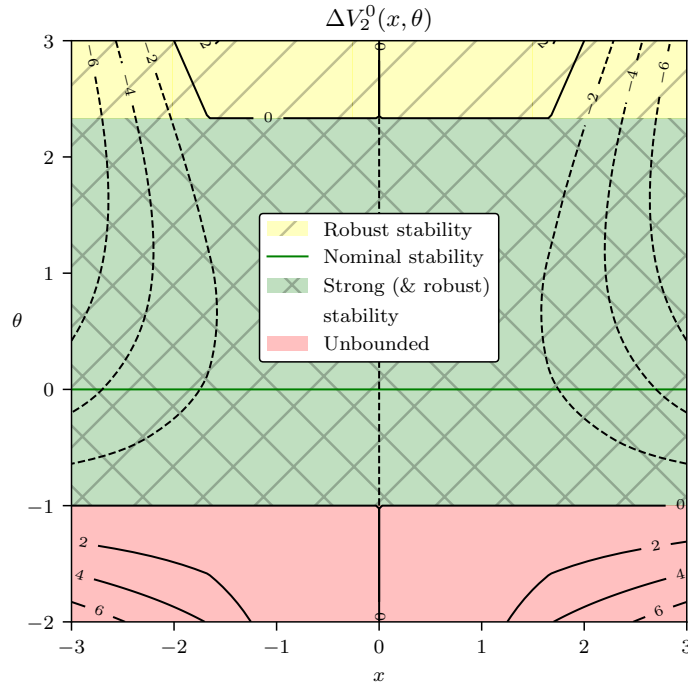


Figure 7.1: Contours of the cost difference as a function of the initial state x and the parameter θ .

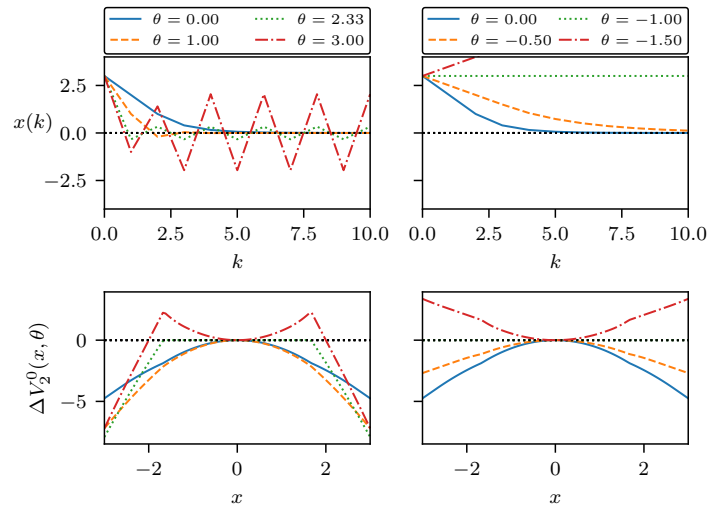


Figure 7.2: For (left) positive and (right) negative values of θ , the (top) closed-loop trajectories with initial state $x = 3$, and (bottom) cost differences as a function of x , along with the nominal values.

existing literature on inherent robustness is not sufficient to predict strong stability whenever $-1 < \theta < 7/3$.

7.2 Robust and strong stability

Consider the closed-loop system

$$x^+ = f_c(x, \theta) := f(x, \kappa_N(x), \theta), \quad \theta \in \Theta \quad (7.16)$$

where $\Theta \subseteq \mathbb{R}^{n_\theta}$. Let $\phi_c(k; x, \boldsymbol{\theta})$ denote solutions to (7.16) at time k , given an initial state $x \in \mathcal{X}_N$ and a sufficiently long parameter sequence $\theta \in \Theta$. If $\Theta := \{\theta \in \mathbb{R}^{n_\theta} \mid |\theta| \leq \delta\}$, it is convenient to write (7.16) as $x^+ = f_c(x, \theta), |\theta| \leq \delta$.

In this section, we review stability definitions and results for (7.16). For brevity, asymptotic and exponential definitions and results are consolidated into the same statement. We define robustly positive invariant (RPI) sets as follows.

Definition 7.11 (Robust positive invariance). A set $X \subseteq \mathbb{R}^n$ is *robustly positive invariant* for the system $x^+ = f_c(x, \theta), \theta \in \Theta$ if $f_c(x, \theta) \in X$ for all $x \in X$ and $\theta \in \Theta$.

7.2.1 Robust stability

We define robust asymptotic stability (RAS) similarly to input-to-state stability (ISS) from Jiang and Wang (2001). Likewise, we define robust exponential stability (RES) similarly to input-to-state exponential stability (ISES) from Grüne et al. (1999).

Definition 7.12 (Robust stability). A system $x^+ = f_c(x, \theta), \theta \in \Theta$ is *robustly asymptotically stable* (in a RPI set $X \subseteq \mathbb{R}^n$) if there exists $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that

$$|\phi_c(k; x, \boldsymbol{\theta})| \leq \beta(|x|, k) + \gamma(\|\boldsymbol{\theta}\|_{0:k-1}) \quad (7.17)$$

for all $k \in \mathbb{I}_{\geq 0}$, $x \in X$, and $\theta \in \Theta^k$. If, additionally, $\beta(s, k) = cs\lambda^k$ for some $c > 0$ and $\lambda \in (0, 1)$, we say $x^+ = f_c(x, \theta), \theta \in \Theta$ is *robustly exponentially stable* (in X).

Definition 7.13 (ISS/ISES Lyapunov function). A function $V : X \rightarrow \mathbb{R}_{\geq 0}$ is an *ISS Lyapunov function* (in an RPI set $X \subseteq \mathbb{R}^n$, for the system $x^+ = f_c(x, \theta), \theta \in \Theta$) if there exists functions $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_{\infty}$ and $\sigma \in \mathcal{K}$ such that

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|) \tag{7.18a}$$

$$V(f_c(x, \theta)) \leq V(x) - \alpha_3(|x|) + \sigma(|\theta|). \tag{7.18b}$$

for all $x \in X$ and $\theta \in \Theta$. If, additionally, $\alpha_i(\cdot) := a_i(\cdot)^b$ for some $a_i, b > 0$ and each $i \in \{1, 2, 3\}$, we say V is an *ISES Lyapunov function* (in X , for $x^+ = f_c(x, \theta), \theta \in \Theta$).

The result below is a generalization of (Allan et al., 2017, Prop. 19) to include general disturbance sets and the exponential case. For completeness, we provide the proof of the exponential case in Appendix 7.A.2.

Theorem 7.14 (ISS/ISES Lyapunov theorem). *The system $x^+ = f_c(x, \theta), \theta \in \Theta$ is RAS (RES) in an RPI set $X \subseteq \mathbb{R}^n$ if it admits an ISS (ISES) Lyapunov function in X .*

Remark 7.15. Whereas (Allan et al., 2017, Prop. 19) only considers disturbance sets of the form $\Theta := \{\theta \in \mathbb{R}^{n_{\theta}} \mid |\theta| \leq \delta\}$ for some $\delta > 0$, it is trivial to modify the proof to use a general constraint set.

7.2.2 Strong stability

We take strong asymptotic stability (SAS) as a time-invariant version of the conclusion of (Jiang and Wang, 2002, Prop. 2.2). Strong exponential stability (SES) is defined similarly.

Definition 7.16 (Strong stability). A system $x^+ = f_c(x, \theta)$, $\theta \in \Theta$ is *strongly asymptotically stable* (in a RPI set $X \subseteq \mathbb{R}^n$) if there exists $\beta \in \mathcal{KL}$ such that

$$|\phi_c(k; x, \theta)| \leq \beta(|x|, k)$$

for all $k \in \mathbb{I}_{\geq 0}$, $x \in X$, and $\theta \in \Theta^k$. If, additionally, $\beta(s, k) := cs\lambda^k$ for all $s \geq 0$ and $k \in \mathbb{I}_{\geq 0}$, and some $c > 0$ and $\lambda \in (0, 1)$, we say $x^+ = f_c(x, \theta)$, $\theta \in \Theta$ is *strongly exponentially stable* (in X).

Definition 7.17 (Lyapunov function). A function $V : X \rightarrow \mathbb{R}_{\geq 0}$ is a *Lyapunov function* (in a RPI set $X \subseteq \mathbb{R}^n$, for the system $x^+ = f_c(x, \theta)$, $\theta \in \Theta$), if there exist functions $\alpha_1, \alpha_2 \in \mathcal{K}_{\infty}$ and a continuous function $\sigma \in \mathcal{PD}$ such that

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|) \tag{7.19a}$$

$$V(f_c(x, \theta)) \leq V(x) - \sigma(|x|) \tag{7.19b}$$

for all $x \in X$ and $\theta \in \Theta$. If, additionally, $\alpha_i(\cdot) := a_i(\cdot)^b$ for some $a_i, b > 0$ and each $i \in \mathbb{I}_{1:3}$, we say V is an *exponential Lyapunov function* (in X , for $x^+ = f_c(x, \theta)$, $\theta \in \Theta$).

The following Lyapunov theorem combines from (Allan et al., 2017, Prop. 13) and (Pannocchia et al., 2011, Lem. 15).

Theorem 7.18. *The system $x^+ = f_c(x, \theta)$, $\theta \in \Theta$ is SAS (SES) in a RPI set $X \subseteq \mathbb{R}^n$ if it admits a Lyapunov function (an exponential Lyapunov function) in X .*

Remark 7.19. In (Allan et al., 2017, Prop. 13), the Lyapunov function requires a class- \mathcal{K}_{∞} bound rather than a continuous class- \mathcal{PD} bound. However, it is shown in (Jiang and Wang, 2002, Lem. 2.8) that a continuous function $\sigma \in \mathcal{PD}$ suffices.

7.3 Inherent robustness of MPC

Assumptions 7.1 to 7.4 are in fact sufficient to show inherent robustness of the nominal MPC. The theorem below can be viewed as a minor generalization of the results in (Rawlings et al., 2020, Sec. 3.2.4), or as a special case of (Allan et al., 2017, Thm. 21).

Theorem 7.20 (Sec. 3.2.4 of Rawlings et al. (2020)). *Suppose Assumptions 7.1 to 7.4 hold. Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. Then there exist $\delta > 0$, $\alpha_2 \in \mathcal{K}_\infty$, and $\sigma \in \mathcal{K}$ such that*

$$\alpha_1(|x|) \leq V_N^0(x) \leq \alpha_2(|x|) \quad (7.20a)$$

$$V_N^0(f_c(x, \theta)) \leq V_N^0(x) - \alpha_1(|x|) + \sigma(|\theta|) \quad (7.20b)$$

for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$, and the system $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is RAS in the RPI set \mathcal{S} .

For completeness, we include a proof of Theorem 7.20 in Appendix 7.A.3. Before moving on, we note that a key step of the proof of Theorem 7.20 and the main results is to establish the following robust descent property:

$$V_N^0(f_c(x, \theta)) \leq V_N^0(x) - \ell(x, \kappa_N(x)) + V_N(f_c(x, \theta), \tilde{\mathbf{u}}(x)) - V_N(\hat{f}_c(x), \tilde{\mathbf{u}}(x)). \quad (7.21)$$

In (Rawlings et al., 2020, Sec. 3.2.4), it is shown that (7.21) can be achieved on any sublevel set of V_N^0 and a sufficiently small neighborhood $|\theta| \leq \delta$. We restate this in the following proposition.

Proposition 7.21 (Sec. 3.2.4 of Rawlings et al. (2020)). *Suppose Assumptions 7.1 to 7.4 hold. Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. Then there exists $\delta > 0$ such that (7.21) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$ and \mathcal{S} is RPI for $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$.*

With quadratic costs (Assumption 7.5), Assumptions 7.1 to 7.3 also imply inherent *exponential* robustness of MPC. This can also be ascertained from (Rawlings et al., 2020, Sec. 3.2.4),

or considered as a special case of (Pannocchia et al., 2011, Thm. 18). A proof of Theorem 7.22, which follows similarly to that of Theorem 7.20, is included in Appendix 7.A.3.

Theorem 7.22 (Sec. 3.2.4 of Rawlings et al. (2020)). *Suppose Assumptions 7.1 to 7.3 and 7.5 hold. Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. There exist $\delta, c_2 > 0$ and $\sigma \in \mathcal{K}$ such that*

$$c_1|x|^2 \leq V_N^0(x) \leq c_2|x|^2 \quad (7.22a)$$

$$V_N^0(f_c(x, \theta)) \leq V_N^0(x) - c_1|x|^2 + \sigma(|\theta|) \quad (7.22b)$$

for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$, where $c_1 := \underline{\sigma}(Q)$, and the system $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is RES in the RPI set \mathcal{S} .

7.4 Stability of MPC despite mismatch

In this section, we investigate two approaches to guarantee strong stability of the closed-loop system (7.16). First, we take a direct approach and assume the existence of an ISS Lyapunov function that achieves a certain maximum increase due to mismatch. In general, an additional scaling condition is required for the mismatch term, although it is automatically satisfied for quadratic cost MPC. Second, we construct error bounds that imply the maximum Lyapunov increase for V_N^0 via the standard MPC assumptions (Assumptions 7.1 to 7.5) and one or both of Assumptions 7.8 and 7.9.

7.4.1 Maximum Lyapunov increase

We begin with the direct approach. The goal here is not (necessarily) to provide the means to check if a given MPC is strongly stabilizing, but to (i) identify a set of conditions for which an ISS Lyapunov function also guarantees strong stability and (ii) provide a path towards proving certain classes of nominal MPCs are strongly stabilizing.

Asymptotic case

For inherent robustness, a maximum increase of the form (7.20b) is proven for the optimal value function V_N^0 . However, since the perturbation term $\sigma(|\theta|)$ is uniform in $|x|$, strong stability is not demonstrated for nonzero θ . Under Assumption 7.8, we might assume the perturbation vanishes in either of the limits $|x| \rightarrow 0$ or $|\theta| \rightarrow 0$. In this sense, the perturbation should be class- \mathcal{K} in $|x|$ whenever $|\theta|$ is fixed, and vice versa. We call these functions *joint \mathcal{K} -functions* or *\mathcal{K}^2 -functions* and define them as follows.

Definition 7.23 (Class \mathcal{K}^2). The class of *joint \mathcal{K} -functions*, denoted \mathcal{K}^2 is the class of continuous functions $\gamma : \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$ such that $\gamma(s, \cdot), \gamma(\cdot, s) \in \mathcal{K}$ for all $s > 0$.

To achieve strong stability, we assume the existence of an ISS Lyapunov function with a \mathcal{K}^2 -function perturbation term, rather than the standard \mathcal{K} -function perturbation term. Moreover, we require the perturbation to decay faster than the nominal cost decrease in the limit $|x| \rightarrow 0$ so that the descent property of Definition 7.17 is achieved for sufficiently small θ .

Assumption 7.24 (Maximum Lyapunov increase). There exists a l.s.c. function $V : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}_{\geq 0}$ such that, for each $\rho > 0$, there exist $\delta_0 > 0$, $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$, and $\gamma_V \in \mathcal{K}^2$ such that

- (a) $\mathcal{S} := \text{lev}_\rho V \subseteq \mathcal{X}_N$;
- (b) for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$, we have

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|) \tag{7.23a}$$

$$V(f_c(x, \theta)) \leq V(x) - \alpha_3(|x|) + \gamma_V(|x|, |\theta|); \tag{7.23b}$$

- (c) and there exists $\tau > 0$ such that

$$\limsup_{s \rightarrow 0^+} \frac{\gamma_V(s, \tau)}{\alpha_3(s)} < 1. \tag{7.24}$$

With Assumption 7.24, we have our first main result.

Theorem 7.25. *Suppose Assumption 7.24 holds with $V : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}_{\geq 0}$. For each $\rho > 0$, there exists $\delta > 0$ for which $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SAS in the RPI set $\mathcal{S} := \text{lev}_\rho V$.*

To prove Theorem 7.25, we require a preliminary result related to the ability of a given \mathcal{K}^2 -function to lower bound another given \mathcal{K} -function (see Appendix 7.A.4 for proof).

Proposition 7.26. *Let $\alpha \in \mathcal{K}_\infty$ and $\gamma \in \mathcal{K}^2$. If there exists $\tau > 0$ such that*

$$\limsup_{s \rightarrow 0^+} \frac{\gamma(s, \tau)}{\alpha(s)} < 1$$

then, for each $\sigma > 0$, there exists $\delta > 0$ such that $\gamma(s, t) < \alpha(s)$ for all $s \in (0, \sigma]$ and $t \in [0, \delta]$.

Finally, we prove Theorem 7.25.

Proof of Theorem 7.25. By Assumption 7.24(a,b) there exists $\delta_0 > 0$, $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$, and $\gamma_V \in \mathcal{K}^2$ such that $\mathcal{S} \subseteq \mathcal{X}_N$ and (7.23) holds for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$. Let $\varepsilon_0 := \sup_{x \in \mathcal{S}} |x| > 0$.² By Assumption 7.24(c) and Proposition 7.26, there exists $\delta_1 > 0$ such that $\alpha_3(s) > \gamma_V(s, t)$ for all $s \in (0, \varepsilon_0]$ and $t \in [0, \delta_1]$. With $\delta := \min \{ \delta_0, \delta_1 \}$, the function

$$\sigma(s) := \begin{cases} \alpha_3(s) - \gamma_V(s, \delta), & 0 \leq s \leq \varepsilon_0 \\ \alpha_3(\varepsilon_0) - \gamma_V(\varepsilon_0, \delta), & s > \varepsilon_0 \end{cases}$$

is both class- \mathcal{PD} and continuous. By (7.23b), we have

$$V(f_c(x, \theta)) - V(x) \leq -\alpha_3(|x|) + \gamma_V(|x|, \delta) = -\sigma(|x|)$$

²If $\mathcal{S} = \{0\}$, the conclusion would hold trivially, so we can assume $\mathcal{S} \neq \{0\}$ without loss of generality.

for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$. Moreover, $V(x) \leq \rho$ implies

$$V(f_c(x, \theta)) \leq V(x) - \sigma(|x|) \leq \rho$$

so $\mathcal{S} = \text{lev}_\rho V$ must be RPI. Finally, $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SAS in \mathcal{S} by Theorem 7.18. \square

Remark 7.27. One might naïvely assume that the closed-loop system (7.16) is SAS under only Assumption 7.24(a,b). However, if the scaling condition Assumption 7.24(c) does not hold, then it may be the case that we cannot shrink t small enough to make $\alpha_3(\cdot) - \gamma_V(\cdot, t)$ positive definite in a sufficiently large neighborhood of the origin, let alone any neighborhood at all. Thus Assumption 7.24(a,b) alone are insufficient to show V is a Lyapunov function for the closed-loop system (7.16). This is illustrated in the example of Section 7.5.1 and in the following examples.

Example 7.28. Let $\alpha_3(s) := s^2$, $\gamma_V(s, t) := st$, and $L := \limsup_{s \rightarrow 0^+} \frac{\gamma_V(s, t)}{\alpha_3(s)}$. Then $\alpha_3 \in \mathcal{K}_\infty$ and $\gamma_V \in \mathcal{K}^2$, but $L = \lim_{s \rightarrow 0^+} t/s = \infty$ for each $t > 0$. In fact, since $\sigma_t(s) := \alpha_3(s) - \gamma_V(s, t) = s^2 - st$, σ_t is negative definite near the origin for each $t > 0$.

Example 7.29. Let $\alpha_3(s) := s$, $\gamma_V(s, t) := \frac{2st}{s+t}$, and $L := \limsup_{s \rightarrow 0^+} \frac{\gamma_V(s, t)}{\alpha_3(s)}$. Then $\alpha_3 \in \mathcal{K}_\infty$ and $\gamma_V \in \mathcal{K}^2$, but $L = \lim_{s \rightarrow 0^+} \frac{2t}{s+t} = 2$ for each $t > 0$. Moreover, since $\sigma_t(s) := \alpha_3(s) - \gamma_V(s, t) = s - \frac{2st}{s+t} = \frac{s^2 - st}{s+t}$, σ_t is negative definite near the origin for each $t > 0$.

Remark 7.30. While Assumption 7.4 implies (7.23) can be satisfied with $\alpha_3 := \alpha_1$, it may be the case that (7.24) is not satisfied. For example, suppose in some neighborhood of the origin, that $\ell(x, u) := |x|^2 + |u|$, $\kappa_N(x) := -x$, and (f, ℓ, V_f) are Lipschitz on compact sets. Then $\gamma_V(s, t) := Lst$, $\alpha_1(s) := s^2$, and $\alpha_3(s) := s^2 + s$ satisfy (7.9), (7.23b), and (7.32) for some $L > 0$. While $\limsup_{s \rightarrow 0^+} \gamma_V(s, t)/\alpha_1(s) = \infty$ for each $t > 0$, we have $\limsup_{s \rightarrow 0^+} \gamma_V(s, t)/\alpha_3(s) = Lt$ and therefore (7.24) holds for any $\tau \in [0, 1/L)$.

Remark 7.31. To achieve Assumption 7.24(a), it is necessary to have $V(x) = \infty$ for all $x \notin \mathcal{X}_N$. Under Assumptions 7.1 to 7.4, this is automatically achieved by the optimal value function V_N^0 , since, according to the convention of Rockafellar and Wets (1998), we have $V_N^0(x) = \infty$ for infeasible problems.

Remark 7.32. A restricted version of Assumption 7.8 is automatically satisfied under Assumption 7.24(b). To see this, we set $x = 0$ in (7.23) to give $f_c(0, \theta) = f(0, \kappa_N(0), \theta) = 0$ for all $|\theta| \leq \delta$ and some $\delta > 0$. If, additionally, Assumptions 7.1, 7.2, and 7.4 are satisfied, we have

$$\tilde{\alpha}_1(|(x, \kappa_N(x))|) \leq \tilde{\alpha}_1(|(x, \kappa_N(x))|) \leq V_N^0(x) \leq \tilde{\alpha}_2(|x|)$$

for some $\tilde{\alpha}_1, \tilde{\alpha}_2 \in \mathcal{K}_\infty$, which implies $\kappa_N(0) = 0$, so $f(0, 0, \theta) = 0$ for all $|\theta| \leq \delta$.

Exponential case

To achieve strong *exponential* stability, Assumption 7.24 is strengthened to require power law versions of the bounds in (7.23). Since identical exponents are required, the scaling condition Assumption 7.24(c) is automatically satisfied.

Assumption 7.33 (Max. Lyapunov incr. (exp.)). There exists a l.s.c. function $V : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}_{\geq 0}$ such that, for each $\rho > 0$, there exist $\delta_0, a_1, a_2, a_3, b > 0$ and $\sigma_V \in \mathcal{K}_\infty$ satisfying

- (a) $\mathcal{S} := \text{lev}_\rho V \subseteq \mathcal{X}_N$; and
- (b) for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$, we have

$$a_1|x|^b \leq V(x) \leq a_2|x|^b \tag{7.25a}$$

$$V(f_c(x, \theta)) \leq V(x) - a_3|x|^b + \sigma_V(|\theta|)|x|^b. \tag{7.25b}$$

With Assumption 7.33, we have our second main result.

Theorem 7.34. *Suppose Assumption 7.33 holds with $V : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}_{\geq 0}$. For each $\rho > 0$, there exists $\delta > 0$ for which $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SES in the RPI set $\mathcal{S} := \text{lev}_\rho V$.*

Proof. Assumption 7.33 gives $\delta_0, a_1, a_2, a_3, b > 0$ such that $\mathcal{S} \subseteq \mathcal{X}_N$ and (7.25) holds for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$. Let $\delta_1 \in (0, \sigma_V^{-1}(a_3))$ and $\delta := \min \{ \delta_0, \delta_1 \} > 0$. Then, by (7.25b),

$$V(f_c(x, \theta)) - V(x) \leq -[a_3 - \sigma_V(\delta)]|x|^b = -a_4|x|^b$$

for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$, where $a_4 := a_3 - \sigma_V(\delta) \geq a_3 - \sigma_V(\delta_1) > 0$. But this means that $V(x) \leq \rho$ implies

$$V(f_c(x, \theta)) \leq V(x) - a_4|x|^b \leq \rho$$

so $\mathcal{S} = \text{lev}_\rho V$ must be RPI. Finally, $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SES in \mathcal{S} by Theorem 7.18. \square

Remark 7.35. Remark 7.31 also applies to Assumption 7.33(a): we require $V(x) = \infty$ for all $x \notin \mathcal{X}_N$.

Remark 7.36. A restricted version of Assumption 7.8 is automatically satisfied under Assumption 7.33(b). Setting $x = 0$ in (7.25) gives $f_c(0, \theta) = f(0, \kappa_N(0), \theta) = 0$ for all $|\theta| \leq \delta$ and some $\delta > 0$. If, additionally, Assumptions 7.1, 7.2, and 7.5 are satisfied, we have

$$c_1|(x, \kappa_N(x))|^2 \leq c_1|(x, \kappa_N(x))|^2 \leq V_N^0(x) \leq c_2|x|^2$$

for some $c_1, c_2 > 0$, which implies $\kappa_N(0) = 0$, so $f(0, 0, \theta) = 0$ for all $|\theta| \leq \delta$.

7.4.2 Error bounds

While the maximum Lyapunov increases (7.23b) and (7.25b) are difficult to verify directly, they are in fact satisfied for the optimal value function (i.e., $V := V_N^0$) under fairly general conditions. To show this, however, we require bounds on the error due to mismatch.

Model error bounds

Stability of MPC under mismatch was first investigated by Santos and Biegler (1999); Santos et al. (2008), who considered, for a fixed parameter $\theta \in \mathbb{R}^{n_\theta}$, the following power law bound:

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq c|x| \quad (7.26)$$

for some $c > 0$ and all $(x, u) \in \mathbb{R}^n \times \mathbb{R}^m$. However, the bound (7.26) does not account for changing or unknown $\theta \in \mathbb{R}^{n_\theta}$ and is uniform in $u \in \mathbb{R}^m$, thus ruling out the motivating example from Section 7.1.3. To handle the former issue, we can take $c = \sigma_f(|\theta|)$ for some $\sigma_f \in \mathcal{K}_\infty$. For the latter issue, it suffices to either replace $|x|$ with $|(x, u)|$, i.e.,

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq \sigma_f(|\theta|)|(x, u)| \quad (7.27)$$

or consider a bound on the closed-loop error, i.e.,

$$|f_c(x, \theta) - \hat{f}_c(x)| \leq \tilde{\sigma}_f(|\theta|)|x| \quad (7.28)$$

for all $x \in \mathcal{S}$, $u \in \mathbb{U}$, and $\theta \in \mathbb{R}^{n_\theta}$, where $\sigma_f, \tilde{\sigma}_f \in \mathcal{K}_\infty$ and $\mathcal{S} \subseteq \mathbb{R}^n$ is an appropriately chosen compact set.

For illustrative purposes, consider the following examples of the bounds (7.27). Note that, for a robustly exponentially stabilizing MPC with quadratic costs (satisfying (7.22)), the control law satisfies $|\kappa_N(x)| \leq \sqrt{c_2/\underline{\sigma}(R)}|x|$, so (7.27) implies (7.28).

Example 7.37. The linear system $x^+ = Ax + Bu$ achieves (7.27) with θ defined as the vectorization of $\begin{bmatrix} A & B \end{bmatrix}$ and $\sigma_f(\cdot) = (\cdot) \in \mathcal{K}_\infty$. More generally, we could consider arbitrary parameterizations of (A, B) that are continuous at $\theta = 0$, i.e., $x^+ = A(\theta)x + B(\theta)u$ where $\bar{\sigma}\left(\begin{bmatrix} A(\theta) & B(\theta) \end{bmatrix} - \begin{bmatrix} A(0) & B(0) \end{bmatrix}\right) \leq \sigma_f(|\theta|)$ and $\sigma_f \in \mathcal{K}_\infty$ is guaranteed by Proposition 7.49

in Appendix 7.A.1.

Example 7.38. Consider the discretized pendulum system

$$x^+ = f(x, u, \theta) := \begin{bmatrix} x_1 + \Delta x_2 \\ x_2 + \Delta(\theta_1 \sin x_1 - \theta_2 x_2 + \theta_3 u) \end{bmatrix}$$

where $\theta \in \mathbb{R}_{>0}^3$ is a vector of lumped parameters and $\Delta > 0$ is the sample time. For a real pendulum system, the discretization will introduce numerical errors, but since the errors are $O(\Delta^2)$, we may assume $\Delta > 0$ is sufficiently small so that they can be safely ignored. For this system we have

$$|f(x, u, \theta) - \hat{f}(x, u, \hat{\theta})| \leq \Delta |\theta - \hat{\theta}| |(x, u)|.$$

where $\hat{\theta} \in \mathbb{R}_{>0}^3$. Shifting θ by $-\hat{\theta}$ gives the bound (7.27).

In the following propositions, we derive the bounds (7.27) and (7.28) using Taylor's theorem and Assumptions 7.1 to 7.3, 7.5, 7.8, and 7.9 (see Appendix 7.A.4 for proofs).

Proposition 7.39. *Suppose Assumptions 7.1, 7.2, 7.8, and 7.9 hold. For each compact set $\mathcal{S} \subseteq \mathbb{R}^n$, there exists $\sigma_f \in \mathcal{K}_\infty$ such that (7.27) holds for all $x \in \mathcal{S}$, $u \in \mathbb{U}$, and $\theta \in \mathbb{R}^{n_\theta}$.*

Proposition 7.40. *Suppose Assumptions 7.1 to 7.3, 7.5, 7.8, and 7.9 hold. For each compact set $\mathcal{S} \subseteq \mathcal{X}_N$, there exists $\tilde{\sigma}_f \in \mathcal{K}_\infty$ such that (7.28) holds for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$.*

More generally, we could consider \mathcal{K}^2 -function bounds,

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq \gamma_f(|(x, u)|, |\theta|) \quad (7.29)$$

$$|f_c(x, \theta) - \hat{f}_c(x)| \leq \tilde{\gamma}_f(|x|, |\theta|) \quad (7.30)$$

for all $x \in \mathcal{S}$ and $\theta \in \Theta$, where $\gamma_f, \tilde{\gamma}_f \in \mathcal{K}^2$, and $\mathcal{S} \subseteq \mathbb{R}^n$ and $\Theta \subseteq \mathbb{R}^{n_\theta}$ are appropriately

chosen compact sets. In the following propositions, we derive the bounds (7.29) and (7.30) using Assumptions 7.1 to 7.3, 7.5, and 7.8 (see Appendix 7.A.4 for proofs).

Proposition 7.41. *Suppose Assumptions 7.1, 7.2, and 7.8 hold. For any compact sets $\mathcal{S} \subseteq \mathbb{R}^n$ and $\Theta \subseteq \mathbb{R}^{n_\theta}$, there exists $\gamma_f \in \mathcal{K}^2$ satisfying (7.29) for all $x \in \mathcal{S}$, $u \in \mathbb{U}$, and $\theta \in \Theta$.*

Proposition 7.42. *Suppose Assumptions 7.1 to 7.4 and 7.8 hold. For any compact sets $\mathcal{S} \subseteq \mathcal{X}_N$ and $\Theta \subseteq \mathbb{R}^{n_\theta}$, there exists $\tilde{\gamma}_f \in \mathcal{K}^2$ satisfying (7.30) for all $x \in \mathcal{S}$ and $\theta \in \Theta$.*

Suboptimal cost error bounds

Ultimately, we require a maximum Lyapunov increase of the form (7.23b) or (7.25b). The robust descent property (7.21) suggests a path through imposing an error bound on the suboptimal cost function $V_N(f_c(x, \theta), \tilde{\mathbf{u}}(x))$, i.e.,

$$|V_N(f_c(x, \theta), \tilde{\mathbf{u}}(x)) - V_N(\hat{f}_c(x), \tilde{\mathbf{u}}(x))| \leq \sigma_V(|\theta|)|x|^2 \quad (7.31)$$

where $\sigma_V \in \mathcal{K}_\infty$. In Proposition 7.43, we establish (7.31) under Assumptions 7.1 to 7.3, 7.5, 7.8, and 7.9 (see Appendix 7.A.4 for proof).

Proposition 7.43. *Suppose Assumptions 7.1 to 7.3, 7.5, 7.8, and 7.9 hold and let $\mathcal{S} \subseteq \mathcal{X}_N$ be compact. Then there exists $\sigma_V \in \mathcal{K}_\infty$ such that (7.31) holds for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$.*

Similarly, we can derive a \mathcal{K}^2 -function version of (7.31) under Assumptions 7.1 to 7.4 and 7.8 (see Appendix 7.A.4 for proof).

Proposition 7.44. *Suppose Assumptions 7.1 to 7.4 and 7.8 hold. Let $\mathcal{S} \subseteq \mathcal{X}_N$ and $\Theta \subseteq \mathbb{R}^{n_\theta}$ be compact. Then there exists $\gamma_V \in \mathcal{K}^2$ such that, for each $x \in \mathcal{S}$ and $\theta \in \Theta$,*

$$|V_N(f_c(x, \theta), \tilde{\mathbf{u}}(x)) - V_N(\hat{f}_c(x), \tilde{\mathbf{u}}(x))| \leq \gamma_V(|x|, |\theta|). \quad (7.32)$$

7.4.3 Stability despite mismatch

General costs

Finally, we are in a position to construct a maximum Lyapunov increase (7.23b) or (7.25b). For general costs, this is accomplished in the following proposition.

Proposition 7.45. *Suppose Assumptions 7.1 to 7.4 and 7.8 hold. Then Assumption 7.24(a,b) hold with $V := V_N^0$.*

Proof. Let $\rho > 0$, $\mathcal{S} := \text{lev}_\rho V_N^0$, and $V := V_N^0$. Then $\mathcal{S} \subseteq \mathcal{X}_N$ trivially. Since V_N^0 is l.s.c. (Bertsekas and Shreve, 1978, Lem. 7.18), \mathcal{S} is closed. By Theorem 7.20, there exists $\alpha_2 \in \mathcal{K}_\infty$ satisfying (7.23a) for all $x \in \mathcal{S}$. Then $|x| \leq \alpha_1^{-1}(V(x)) \leq \alpha_1^{-1}(\rho)$ for all $x \in \mathcal{S}$, so \mathcal{S} is compact.

By Proposition 7.21, there exists $\delta_0 > 0$ such that \mathcal{S} is RPI for $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta_0$ and (7.21) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$. Moreover, for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$, (7.32) holds for some $\gamma_V \in \mathcal{K}^2$ by Proposition 7.44. Finally, combining (7.9), (7.21), and (7.32) gives (7.23b) with $\alpha_3 := \alpha_1$. \square

Assumption 7.24(a,b) alone do not guarantee strong stability. However, we can strengthen the hypothesis of Proposition 7.45 with a scaling requirement to guarantee strong stability.

Corollary 7.46. *Suppose Assumptions 7.1 to 7.4 and 7.8 hold. Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. Then (7.23) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$ with $V := V_N^0$ and some $\delta_0 > 0$, $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$, and $\gamma_V \in \mathcal{K}^2$. If, additionally, there exists $\tau > 0$ satisfying (7.24), then there exists $\delta > 0$ such that $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SES in the RPI set \mathcal{S} .*

Proof. The first part follows from Proposition 7.45, and the second part follows from Theorem 7.25. \square

Quadratic costs

For quadratic costs, we construct (7.25b) in the following proposition.

Proposition 7.47. *Suppose Assumptions 7.1 to 7.3, 7.5, 7.8, and 7.9 hold. Then Assumption 7.33 holds with $b := 2$ and $V := V_N^0$.*

Proof. Let $\rho > 0$, $V := V_N^0$, and $\mathcal{S} := \text{lev}_\rho V$. Since Assumption 7.5 implies Assumption 7.4, we have from the first paragraph of the proof of Proposition 7.45 that \mathcal{S} is compact.

Theorem 7.22 also implies (7.25a) holds for all $x \in \mathcal{S}$, with $a_1, a_2 > 0$ and $b := 2$. By Proposition 7.21, there exists $\delta_0 > 0$ such that \mathcal{S} is RPI for $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta_0$ and (7.21) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$. Moreover, for each $x \in \mathcal{S}$ and $|\theta| \leq \delta_0$, (7.31) holds for some $\sigma_V \in \mathcal{K}_\infty$ by Proposition 7.44, and combining (7.21) and (7.32) gives (7.25b). \square

Our third and final main result is an immediate corollary to Theorem 7.25 and Proposition 7.47.

Corollary 7.48. *Suppose Assumptions 7.1 to 7.4, 7.8, and 7.9 holds. For each $\rho > 0$, there exists $\delta > 0$ for which $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SES in the RPI set $\mathcal{S} := \text{lev}_\rho V_N^0$.*

Proof. By Proposition 7.47, Assumption 7.33 holds with $V := V_N^0$, and by Theorem 7.34, there exists $\delta > 0$ for which \mathcal{S} is RPI and $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is SES in \mathcal{S} . \square

7.5 Examples

In this section, we illustrate the nuances of Assumptions 7.24 and 7.33 through several examples. First, we consider a non-differentiable system that satisfies Assumption 7.24(a,b) but not Assumption 7.24(c), and is not SAS. Second, we consider a non-differentiable example that nonetheless satisfies Assumption 7.33 and is therefore SES. Finally, we consider the inverted pendulum system to showcase how the nominal MPC handles different types of mismatch.




Figure 7.3: Contours of the cost difference for the MPC of (7.33).

Notably, we consider (i) discretization errors, (ii) unmodeled dynamics, and (iii) incorrectly estimated input gains.

7.5.1 Strong asymptotic stability counterexample

Consider the scalar system

$$x^+ = f(x, u, \theta) := \sigma(x + (1 + \theta)u) \quad (7.33)$$

where σ is the *signed square root* defined as $\sigma(y) := \text{sgn}(y)\sqrt{|y|}$ for each $y \in \mathbb{R}$. We define a nominal MPC with $\mathbb{U} := [-1, 1]$, $\ell(x, u) := x^2 + u^2$, $V_f(x) := 4x^2$, $\mathbb{X}_f := [-1, 1]$, and $N := 1$.

In Appendix 7.B, it is shown the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 3$ is RES on $\mathcal{X}_1 = [-2, 2]$ with the nominal control law $\kappa_1(x) := -\text{sat}(x)$. Additionally, it is shown




Figure 7.4: For (left) nonnegative and (right) nonpositive values of θ , the (top) closed-loop trajectories for the MPC of (7.33) with initial state $x = 2$, and (bottom) cost differences of the same MPC as a function of x .

Assumption 7.24(a,b) is satisfied with $V := V_1^0$, and (7.23b) holds for all $x \in \mathcal{S} := \text{lev}_2 V_1^0 = [-1, 1]$ and $|\theta| \leq \delta_0 := 3$ with $\alpha_3(s) := 2s^2$, and $\gamma_V(s, t) := st + 4\sqrt{st}$. But this implies $\lim_{s \rightarrow 0^+} \gamma_V(s, t)/\alpha_3(s) = \infty$ for each $t > 0$, so Assumption 7.24(c) is not satisfied.

However, Assumption 7.24 is only sufficient, not necessary, for establishing strong stability. But we have $V_1^0(x) = 2x^2$ and

$$\begin{aligned} \Delta V_1^0(x, \theta) &:= V_1^0(f(x, \kappa_1(x), \theta)) - V_1^0(x) \\ &= 2[\sigma(\theta x)]^2 - 2x^2 = 2(|\theta| - |x|)|x| > 0. \end{aligned}$$

for each $0 < |x| < |\theta| \leq 1$, so the state always gets pushed out of $(-|\theta|, |\theta|)$ unless it starts at the origin or $\theta = 0$. In other words, the MPC only provides inherent robustness, not strong stability, even though Assumption 7.24(a,b) is satisfied.

In Figure 7.3, we plot contours of the cost difference $\Delta V_1^0(x, \theta)$, and in Figure 7.4 we plot closed-loop trajectories and the cost difference curve $\Delta V_1^0(\cdot, \theta)$ for several values of θ . Only with $\theta = 0$ does the trajectory converge to the origin and the cost difference curve remain




Figure 7.5: Contours of the cost difference for the MPC of (7.34).

negative definite. For each $\theta \neq 0$, the cost difference is positive definite near the origin, and the trajectory does not converge to the origin.

7.5.2 Non-differentiable yet strongly exponential stable

Consider the scalar system

$$x^+ = f(x, u, \theta) := x + (1/2)\gamma(x) + (1 + \theta)u \quad (7.34)$$

where $\gamma : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\gamma(x) := \begin{cases} 0, & x = 0, \\ |x| \sin(2\pi/x), & x \neq 0. \end{cases}$$

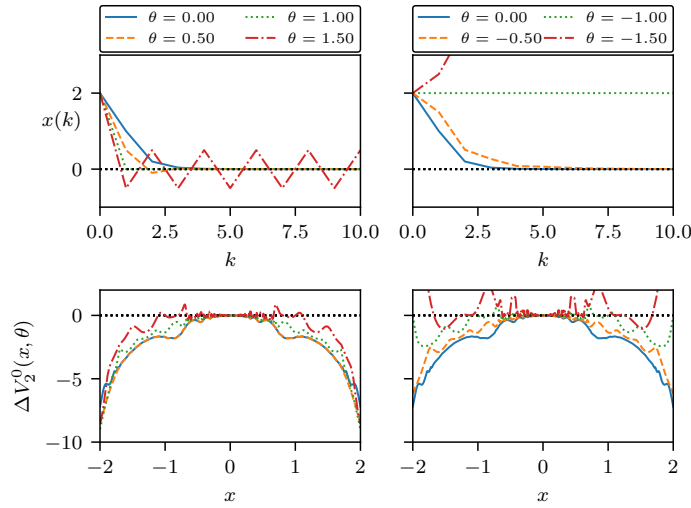


Figure 7.6: For (left) nonnegative and (right) nonpositive values of θ , the (top) closed-loop trajectories for the MPC of (7.34) with initial state $x = 2$, and (bottom) cost differences of the same MPC as a function of x .

While the function γ is continuous, it is not differentiable at the origin. We define a nominal MPC with $\mathbb{U} := [-1, 1]$, $\ell(x, u) := x^2 + u^2$, $V_f(x) := 4x^2$, $\mathbb{X}_f := [-1, 1]$, and $N := 1$.

In Appendix 7.B.2, we show the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 1$ is RES on $\mathcal{X}_1 = [-2, 2]$ with the nominal control law $\kappa_1(x) := -\text{sat}((4/5)x + (2/5)\gamma(x))$. Moreover, it is shown that Assumption 7.33 is satisfied, and by Theorem 7.34 (and its proof), the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta := 0.5$ is SES on $\mathcal{X}_1 = [-2, 2]$.

To establish a clearer picture of robust and strong stability for the closed-loop, we plot in Figure 7.5 contours of the cost difference $\Delta V_1^0(x, \theta) := V_1^0(f(x, \kappa_1(x), \theta)) - V_1^0(x)$, and in Figure 7.6 closed-loop trajectories and the cost difference curve $\Delta V_1^0(\cdot, \theta)$ for several values of θ . For θ between $\theta_0 \approx 0.57$ and $\theta_1 \approx 1.08$, the closed-loop system is strongly stable, with trajectories converging to the origin, and a negative definite cost difference curve. Outside of this range but with $\theta \in [-1, 1.5]$, the closed-loop system is still robustly stable, with a cost difference curve of ambiguous sign but trajectories converging to a neighborhood of the origin. Finally, for $\theta < -1$, trajectories are unbounded because \mathcal{X}_1 is not RPI.

7.5.3 Upright pendulum

Consider the nondimensionalized pendulum system

$$\dot{x} = F(x, u, \theta) := \begin{bmatrix} x_2 \\ \sin x_1 - \theta_1^2 x_2 + (\hat{k} + \theta_2)u \end{bmatrix} \quad (7.35)$$

where $x_1, x_2 \in \mathbb{R}$ are the angle and angular velocity, $u \in [-1, 1]$ is the (signed and normalized) motor voltage, $\theta_1 \in \mathbb{R}$ is an air resistance factor, $\hat{k} > 0$ is the estimated gain of the motor, and $\theta_2 \in \mathbb{R}$ is the error in the motor gain estimate. Let $\psi(t; x, u, \theta)$ denote the solution to the differential equation (7.35) at time $t \geq 0$ given an initial condition $x(0) = x$, constant input signal $u(t) = u$, and parameters θ . We model the continuous-time system (7.35) as

$$x^+ = f(x, u, \theta) := x + \Delta F(x, u, \theta) + \theta_3 r(x, u, \theta) \quad (7.36)$$

where r is a residual function given by

$$r(x, u, \theta) := \int_0^\Delta [F(\psi(t; x, u, \theta), u, \theta) - F(x, u, \theta)] dt.$$

Assuming a zero-order hold on the input u , the system (7.35) is discretized (exactly) as (7.36) with $\theta_3 = 1$. Since we model the system with $\theta = 0$ as

$$x^+ = \hat{f}(x, u) := f(x, u, 0) = x + \Delta \begin{bmatrix} x_2 \\ \sin x_1 + \hat{k}u \end{bmatrix} \quad (7.37)$$

we do not need access to r to design the nominal MPC.

For the following simulations, let the model gain be $\hat{k} = 5 \text{ rad/s}^2$, the sample time be $\Delta = 0.1 \text{ s}$, and define a nominal MPC with $N := 20$, $\mathbb{U} := [-1, 1]$, $\ell(x, u) := |x|^2 + u^2$, $V_f(x) := |x|_{P_f}^2$, $\mathbb{X}_f := \text{lev}_{c_f} V_f$, and $c_f := \underline{\sigma}(P_f)/8$, where $P_f = \begin{bmatrix} 31.133\dots & 10.196\dots \\ 10.196\dots & 10.311\dots \end{bmatrix}$ is shown, in Appendix 7.B.3, to satisfy Assumption 7.3 with the terminal law $\kappa_f(x) := -2x_1 - 2x_2$.

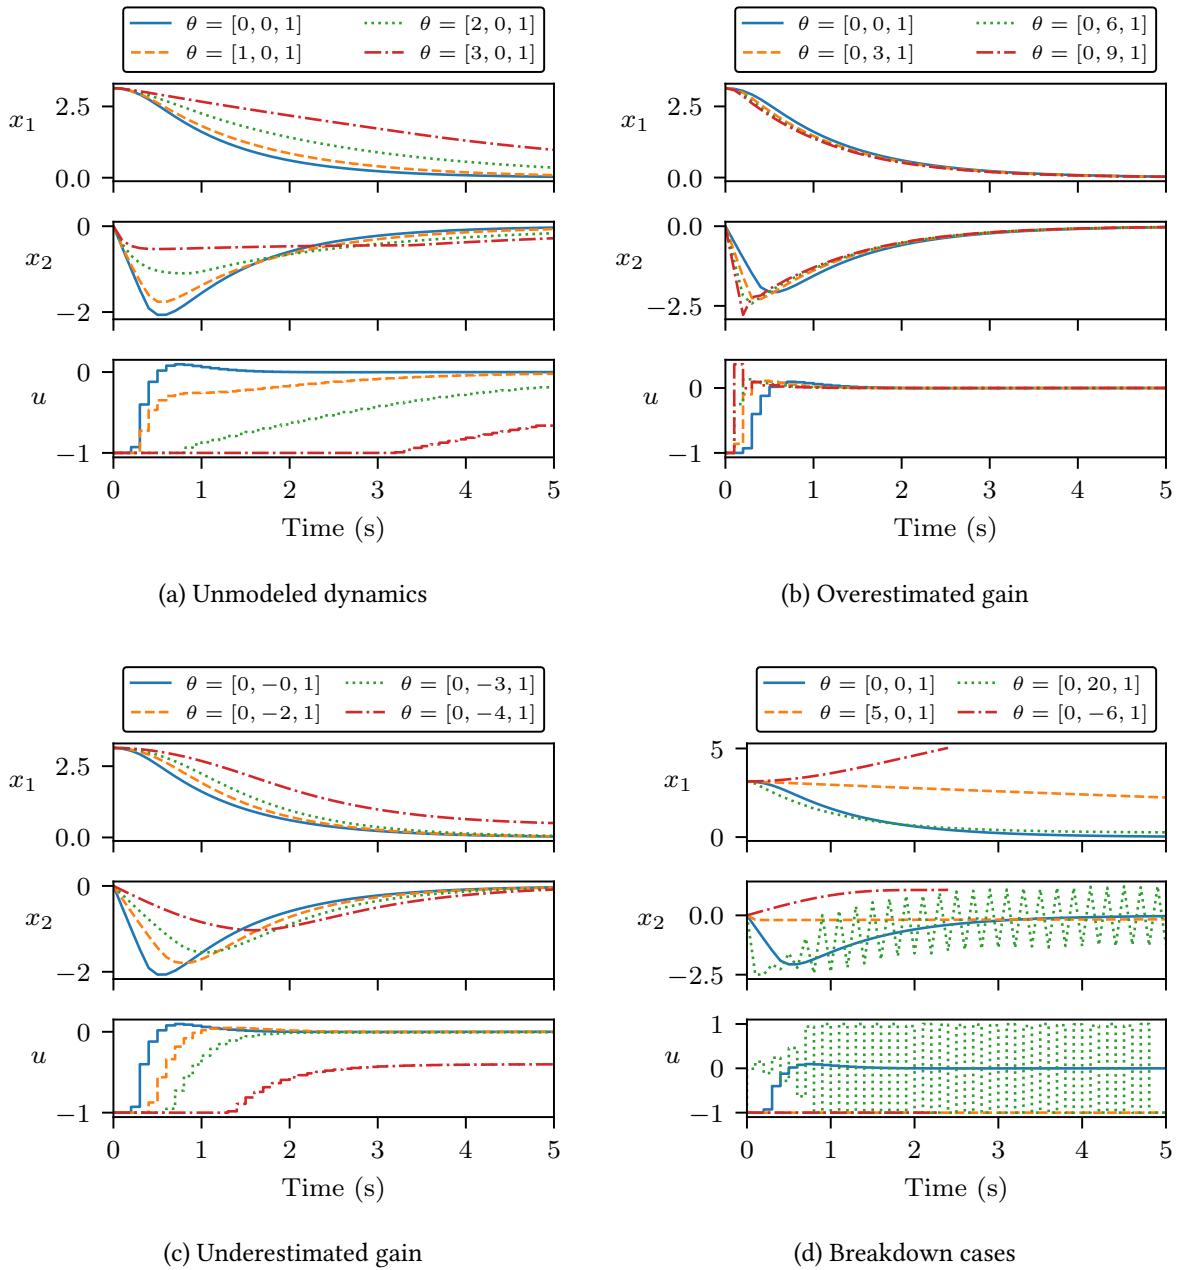


Figure 7.7: Simulated closed-loop trajectories for the MPC of (7.35) from the resting position $x(0) = (\pi, 0)$ to the upright position $x_s = (0, 0)$ for various values of $(\theta_1, \theta_2) \in \mathbb{R}^2$.

Assumptions 7.1, 7.2, 7.5, and 7.8 are satisfied trivially, and Assumption 7.9 is satisfied since continuous differentiability of F implies continuous differentiability of ψ (and therefore also r and f) (Hale, 1980, Thm. 3.3). Thus, the conclusion of Corollary 7.48 holds for some $\delta > 0$, and if we can take $\delta > 1$, the nominal MPC is inherently strongly stabilizing with $\begin{bmatrix} \theta_1 & \theta_2 \end{bmatrix}$ sufficiently small.

In Figure 7.7, we simulate the closed-loop system $x^+ = f(x, \kappa_{20}(x), \theta)$ for some fixed $\begin{bmatrix} \theta_1 & \theta_2 & 1 \end{bmatrix}^\top \in \mathbb{R}^3$. Note that all of these simulations include discretization errors. Figure 7.7a showcases the ability of MPC to handle unmodeled dynamics (i.e., a missing air resistance term). In Figure 7.7b, the gain of the motor is increased until the nominal controller is severely underdamped. In Figure 7.7c, the gain of the motor is decreased until the motor cannot overcome the force of gravity and strong stability is not achieved. In Figure 7.7d, we plot cases where the errors as made so extreme as to prevent stability.

7.6 Conclusion

We establish conditions under which MPC is strongly stabilizing despite plant-model mismatch in the form of parameter errors. Namely, it suffices to assume the existence of a Lyapunov function with a maximum increase, suitably bounded level sets, and a scaling condition (Assumptions 7.24 and 7.33). While we are not able to show the assumptions hold in general, when the MPC has quadratic costs it is possible to show that continuous differentiability of the dynamics implies strong stability (Theorem 7.34). When the \mathcal{K}^2 -function bound is not properly scaled, the MPC may not be stabilizing, as illustrated in the examples. In this sense, while MPC is not inherently stabilizing under mismatch *in general*, there is a common class of cost functions (quadratic costs) for which nominal MPC is inherently stabilizing under mismatch.

Several questions about the strong stability of MPC remain unanswered. While quadratic

costs are used in many control problems, it may be possible to generalize Corollary 7.48 to other useful classes of stage costs, such as q -norm costs, or costs with exact penalty functions for soft state constraints. We propose the direct approach to strong exponential stability (Assumption 7.33 and Theorem 7.34) provides a path to generalizing Corollary 7.48 to other classes of stage costs, output feedback, or semidefinite costs. We note that the Assumptions 7.24 and 7.33 are dependent on the horizon length. This leaves the possibility that some MPC problems are strongly stabilizing at smaller horizon lengths but only inherently robust at longer horizon lengths, or vice versa. However, this remains to be seen. Nonlinear MPC is computationally difficult to implement online. Therefore it would be worth extending this work to include the suboptimal MPC algorithm from Allan et al. (2017) using the approach therein.

While systems with fixed and known setpoints are a useful and interesting class of problems, many systems have setpoints that must be tracked that may change based on the value of the parameters. Offset-free MPC may be used to accommodate the effect of mismatch on the setpoints. As we discussed in Chapter 2, theory on nonlinear offset-free MPC is fairly limited, typically relying on stability of the closed-loop system to guarantee offset-free performance (Pannocchia et al., 2015). In the subsequent chapter, we use the tools developed in this chapter to extend the offset-free MPC theory by establishing closed-loop stability and guaranteed offset-free performance for tracking random, asymptotically constant setpoints subject to plant-model mismatch.

Appendices

7.A Additional proofs

7.A.1 Nominal MPC stability

In this appendix, we provide sketches of the MPC stability results Theorems 7.6 and 7.7. First, the lower bound $V_N^0(x) \geq \alpha_1(|x|)$ follows immediately from Assumption 7.4. Next, consider the following proposition from Allan et al. (2017).

Proposition 7.49 (Prop. 20 of Allan et al. (2017)). *Let $C \subseteq D \subseteq \mathbb{R}^n$, with C compact, D closed, and $f : D \rightarrow \mathbb{R}^m$ continuous. Then there exists $\alpha \in \mathcal{K}_\infty$ such that $|f(x) - f(y)| \leq \alpha(|x - y|)$ for all $x \in C$ and $y \in D$.*

Under Assumptions 7.1 to 7.4, we can establish the following bounds via Proposition 7.49,³

$$V_f(x) \leq \alpha_f(|x|), \quad \forall x \in \mathbb{X}_f \quad (7.38)$$

$$V_N^0(x) \leq \alpha_2(|x|), \quad \forall x \in \mathcal{X}_N \quad (7.39)$$

for some $\alpha_f, \alpha_2 \in \mathcal{K}_\infty$. To establish the cost difference bound, first note that, under Assump-

³Equation (7.38) follows immediately from Proposition 7.49 and Assumptions 7.1 and 7.2. For (7.39), see (Rawlings et al., 2020, Prop. 2.16).

tions 7.2 and 7.3, we have

$$V_f(\hat{f}(x, \kappa_f(x))) \leq V_f(x) - \ell(x, \kappa_f(x)) \leq c_f$$

for all $x \in \mathbb{X}_f$. Therefore \mathbb{X}_f is positive invariant for $x^+ = \hat{f}(x, \kappa_f(x))$. But this means \mathcal{X}_N is positively invariant because, for each $x \in \mathcal{X}_N$, $\tilde{\mathbf{u}}(x)$ steers the system into \mathbb{X}_f in $N - 1$ moves and keeps it there, meaning $\hat{f}_c(x) \in \mathcal{X}_N$. Finally, Assumption 7.3 implies

$$V_N^0(\hat{f}_c(x)) \leq V_N(\hat{f}_c(x), \tilde{\mathbf{u}}(x)) \leq V_N^0(x) - \ell(x, \kappa_N(x)) \quad (7.40)$$

for all $x \in \mathcal{X}_N$ (Rawlings et al., 2020, pp. 116–117). Therefore $V_N^0(\hat{f}_c(x)) \leq V_N^0(x) - \alpha_1(|x|)$ by Assumption 7.4.

Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$. As noted in the main text, we have $\mathcal{S} \subseteq \mathcal{X}_N$ by definition of the sublevel set. Assumptions 7.2 and 7.5 implies $\underline{\sigma}(P_f)|x|^2 \leq V_f(x) \leq c_f$ for all $x \in \mathbb{X}_f$, so we have $|x| \leq \varepsilon := \sqrt{c_f/\underline{\sigma}(P_f)}$ for all $x \in \mathbb{X}_f$. Then with $c_2 := \max\{\bar{\sigma}(P_f), \rho/\varepsilon^2\}$, we can write

$$V_N^0(x) \leq \begin{cases} V_f(x) \leq \bar{\sigma}(P_f)|x|^2 \leq c_2|x|^2, & |x| \leq \varepsilon, \\ \rho \leq c_2\varepsilon^2 \leq c_2|x|^2, & |x| \geq \varepsilon. \end{cases}$$

for each $x \in \mathcal{S}$. Finally, V_N^0 is an exponential Lyapunov function in \mathcal{S} for $x^+ = \hat{f}_c(x)$.

7.A.2 Lyapunov proofs

In this appendix, we prove some of the Lyapunov results of Section 7.2.

Proof of Theorem 7.14 (exponential case). The case where an ISS Lyapunov function implies RAS for a system is covered by (Allan et al., 2017, Prop. 19), so we only consider the ISES/RES case.

Let $X \subseteq \mathbb{R}^n$ be RPI and suppose $V : X \rightarrow \mathbb{R}_{\geq 0}$ is an ISES Lyapunov function, both for the system $x^+ = f_c(x, \theta)$, $\theta \in \Theta$. Then there exist $a_1, a_2, a_3, b > 0$ satisfying (7.18) for all $x \in X$, where $\alpha_i(\cdot) := a_i(\cdot)^b$ for each $i \in \{1, 2, 3\}$. Suppose, without loss of generality, that $a_3 < a_2$. Then (7.18) can be rewritten

$$\begin{aligned} V(f_c(x, \theta)) &\leq V(x) - a_3|x|^b + \sigma(|\theta|) \\ &\leq V(x) - \frac{a_3}{a_2}V(x) + \sigma(|\theta|) \\ &= \lambda_0 V(x) + \sigma(|\theta|) \end{aligned}$$

for all $x \in X$ and $\theta \in \Theta$, where $\lambda_0 := 1 - \frac{a_3}{a_2} \in (0, 1)$. Since X is RPI, this implies

$$\begin{aligned} V(\phi_c(k; x, \boldsymbol{\theta}_{0:k-1})) &\leq \lambda_0^k V(x) + \sum_{i=1}^k \lambda_0^{i-1} \sigma(|\theta(k-i)|) \\ &\leq \lambda_0^k V(x) + \left(\sum_{i=1}^k \lambda_0^{i-1} \right) \max_{i \in \mathbb{I}_{0:k-1}} \sigma(|\theta(i)|) \\ &\leq \lambda_0^k V(x) + \frac{\max_{i \in \mathbb{I}_{0:k-1}} \sigma(|\theta(i)|)}{1 - \lambda_0} \\ &= \lambda_0^k V(x) + \frac{\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{1 - \lambda_0} \\ &= a_2|x|^b \lambda_0^k + \frac{\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{1 - \lambda_0} \end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$, $x \in X$, and $\boldsymbol{\theta} \in \Theta^k$. If $b \geq 1$, then, by the triangle inequality for the b -norm,

we have

$$\begin{aligned}
|\phi_c(k; x, \boldsymbol{\theta}_{0:k-1})| &\leq \left(\frac{V(\phi_c(k; x, \boldsymbol{\theta}_{0:k-1}))}{a_1} \right)^{1/b} \\
&\leq \frac{1}{a_1^{1/b}} \left(a_2 |x|^b \lambda_0^k + \frac{\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{1 - \lambda_0} \right)^{1/b} \\
&\leq \left(\frac{a_2}{a_1} \right)^{1/b} |x| (\lambda_0^b)^k + \left(\frac{\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{a_1(1 - \lambda_0)} \right)^{1/b} \\
&\leq c |x| \lambda^k + \gamma(\|\boldsymbol{\theta}_{0:k-1}\|)
\end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$, $x \in X$, and $\boldsymbol{\theta} \in \Theta^k$, where $\lambda := \lambda_0^{1/b} \in (0, 1)$, $c := \left(\frac{a_2}{a_1} \right)^{1/b} > 0$, and $\gamma(\cdot) := \left(\frac{\sigma(\cdot)}{a_1(1 - \lambda_0)} \right)^{1/b} \in \mathcal{K}$. On the other hand, if $b \in (0, 1)$, then $1/b \geq 1$, so by convexity of $(\cdot)^{1/b}$, we have

$$\begin{aligned}
|\phi_c(k; x, \boldsymbol{\theta}_{0:k-1})| &\leq \left(\frac{2}{a_1} \right)^{1/b} \left(\frac{1}{2} a_2 |x|^b \lambda_0^k + \frac{1}{2} \frac{\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{1 - \lambda_0} \right)^{1/b} \\
&\leq \frac{1}{2} \left(\frac{2a_2}{a_1} \right)^{1/b} |x| (\lambda_0^b)^k + \frac{1}{2} \left(\frac{2\sigma(\|\boldsymbol{\theta}_{0:k-1}\|)}{a_1(1 - \lambda_0)} \right)^{1/b} \\
&\leq c |x| \lambda^k + \gamma(\|\boldsymbol{\theta}_{0:k-1}\|)
\end{aligned}$$

for all $k \in \mathbb{I}_{\geq 0}$, $x \in X$, and $\boldsymbol{\theta} \in \Theta^k$, where $\lambda := \lambda_0^{1/b} \in (0, 1)$, $c := \frac{1}{2} \left(\frac{2a_2}{a_1} \right)^{1/b} > 0$, and $\gamma(\cdot) := \frac{1}{2} \left(\frac{2\sigma(\cdot)}{a_1(1 - \lambda_0)} \right)^{1/b} \in \mathcal{K}$. In either case, (7.17) is satisfied with $\beta(s, k) := cs\lambda^k$ and $\gamma \in \mathcal{K}$ for some $c > 0$ and $\lambda \in (0, 1)$. \square

7.A.3 Proofs of inherent robustness results

This appendix contains proofs of the inherent robustness results from Section 7.3. From Proposition 7.49, we have the following proposition.

Proposition 7.50. *Suppose Assumptions 7.1 and 7.2 holds and let $\tilde{V}_f(\cdot, \cdot) := V_f(\hat{\phi}(N; \cdot, \cdot))$. Then, for any compact set $\mathcal{S} \subseteq \mathcal{X}_N$, there exist $\alpha_a, \alpha_b, \alpha_\theta \in \mathcal{K}_\infty$ such that, for each $x \in \mathcal{S}$ and*

$\theta \in \mathbb{R}^{n_\theta}$,

$$|\tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) - \tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x))| \leq \alpha_a(|x^+ - \hat{x}^+|) \quad (7.41)$$

$$|V_N(x^+, \tilde{\mathbf{u}}(x)) - V_N(\hat{x}^+, \tilde{\mathbf{u}}(x))| \leq \alpha_b(|x^+ - \hat{x}^+|) \quad (7.42)$$

$$|f_c(x, \theta) - \hat{f}_c(x)| \leq \alpha_\theta(|\theta|) \quad (7.43)$$

where $x^+ := f_c(x, \theta)$ and $\hat{x}^+ := \hat{f}_c(x)$.

Proof. Assumptions 7.1 and 7.2 guarantee $\tilde{\mathbf{u}}(x)$ is well-defined for all $x \in \mathcal{X}_N$ (Rawlings et al., 2020, Prop. 2.4). Define $C_0 := \mathcal{S} \times \mathbb{U} \times \{0\}$ and $C_1 := \mathcal{S} \times \mathbb{U}^N$. Then C_0 and C_1 are compact, f is continuous, and \tilde{V}_f and V_N are continuous as they are finite compositions of continuous functions. By Proposition 7.49, there exist $\alpha_a, \alpha_b, \alpha_\theta \in \mathcal{K}_\infty$ such that

$$|\tilde{V}_f(x^+, \mathbf{u}) - \tilde{V}_f(\hat{x}^+, \hat{\mathbf{u}})| \leq \alpha_a(|(x^+ - \hat{x}^+, \mathbf{u} - \hat{\mathbf{u}})|)$$

$$|V_N(x^+, \mathbf{u}) - V_N(\hat{x}^+, \hat{\mathbf{u}})| \leq \alpha_b(|(x^+ - \hat{x}^+, \mathbf{u} - \hat{\mathbf{u}})|)$$

$$|f(x, u, \theta) - \hat{f}(\hat{x}, \hat{u})| \leq \alpha_\theta(|(x - \hat{x}, u - \hat{u}, \theta)|)$$

for all $(\hat{x}, \hat{u}, 0) \in C_0$, $(\hat{x}^+, \hat{\mathbf{u}}) \in C_1$, $(x, u, \theta) \in \mathbb{R}^{n+m+n_\theta}$, and $(x^+, \mathbf{u}) \in \mathbb{R}^{n+Nm}$. Specializing the above inequalities to $x = \hat{x}$, $\hat{x}^+ = \hat{f}_c(x)$, $x^+ = f_c(x, \theta)$, $u = \hat{u} = \kappa_N(x)$, and $\mathbf{u} = \hat{\mathbf{u}} = \tilde{\mathbf{u}}(x)$ gives (7.41)–(7.43) for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$. \square

Next, we can prove Proposition 7.21.

Proof of Proposition 7.21. First, we have $\alpha_1, \alpha_2, \alpha_f \in \mathcal{K}_\infty$ satisfying the bounds (7.9), (7.12), and (7.38)–(7.40) from the assumptions and Theorem 7.6 (and its proof). Next, we let $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$, and define $\tilde{V}_f(\cdot, \cdot) := V_f(\hat{\phi}(N; \cdot, \cdot))$, $x^+ := f_c(x, \theta)$, and $\hat{x}^+ := \hat{f}_c(x)$, throughout. By Proposition 7.50, there exist $\alpha_a, \alpha_b, \alpha_\theta \in \mathcal{K}_\infty$ satisfying the bounds (7.41)–(7.43).

(a)—*Robust feasibility*: By nominal feasibility, we have $\hat{x}^0(N; \hat{x}^+) \in \mathbb{X}_f$ and therefore $V_f(\hat{x}^0(N; x)) \leq c_f$. But $\hat{\phi}(N; x^+, \tilde{\mathbf{u}}(x)) = \hat{f}(\hat{x}^0(N; x), \kappa_f(\hat{x}^0(N; x)))$ and therefore

$$\begin{aligned} \tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) &= V_f(\hat{\phi}(N; \hat{x}^+, \tilde{\mathbf{u}}(x))) \\ &= V_f(\hat{f}(\hat{x}^0(N; x), \kappa_f(\hat{x}^0(N; x)))) \\ &\leq V_f(\hat{x}^0(N; x)) - \alpha_1(|\hat{x}^0(N; x)|) \end{aligned}$$

where the inequality follows from Assumptions 7.3 and 7.4. If $V_f(\hat{x}^0(N; x)) \geq c_f/2$, then $|\hat{x}^0(N; x)| \geq \alpha_f^{-1}(c_f/2)$ and $\tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) \leq c_f - \alpha_1(\alpha_f^{-1}(c_f/2))$. On the other hand, if $V_f(\hat{x}^0(N; x)) < c_f/2$, then $\tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) < c_f/2$. In summary,

$$\tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) \leq c_f - \gamma_1$$

where $\gamma_1 := \min \{ c_f/2, \alpha_1(\alpha_f^{-1}(c_f/2)) \} > 0$. Combining the above inequality with (7.41) and (7.43) gives

$$\tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) \leq c_f - \gamma_1 + \alpha_a(\alpha_\theta(|\theta|)).$$

Therefore, so long as $|\theta| \leq \delta_1 := \alpha_\theta^{-1}(\alpha_a^{-1}(\gamma_1))$, we have

$$V_f(\hat{\phi}(N; x^+, \tilde{\mathbf{u}}(x))) = \tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) \leq c_f$$

which implies $\hat{\phi}(N; x^+, \tilde{\mathbf{u}}(x)) \in \mathbb{X}_f$, and therefore $(x^+, \tilde{\mathbf{u}}(x)) \in \mathcal{Z}_N$.

(b)—*Descent property*: Suppose $|\theta| \leq \delta_1$. Then $(x^+, \tilde{\mathbf{u}}(x)) \in \mathcal{Z}_N$ by part (a), so the inequality $V_N^0(x^+) \leq V_N(x^+, \tilde{\mathbf{u}}(x))$ follows by optimality. Combining this inequality with the nominal descent property (7.40) gives the robust descent property (7.21).

(c)—*Positive invariance of \mathcal{S}* : Suppose again that $|\theta| \leq \delta_1$. Then the inequality (7.43) holds

from part (b), and combining it with (7.21) and (7.42) gives

$$V_N^0(x^+) \leq V_N^0(x) - \alpha_1(|x|) + \alpha_b(\alpha_\theta(|\theta|)). \quad (7.44)$$

If $V_N^0(x) \geq \rho/2$, then $|x| \geq \alpha_2^{-1}(\rho/2)$ and $V_N^0(x^+) \leq \rho - \alpha_1(\alpha_2^{-1}(\rho/2)) + \alpha_b(\alpha_\theta(|\theta|))$. On the other hand, if $V_N^0(x) < \rho/2$, then $V_N^0(x^+) < \rho/2 + \alpha_b(\alpha_\theta(|\theta|))$. Then

$$V_N^0(x^+) \leq \rho - \gamma_2 + \alpha_b(\alpha_\theta(|\theta|))$$

where $\gamma_2 := \min \{ \rho/2, \alpha_1(\alpha_2^{-1}(\rho/2)) \} > 0$. Therefore $V_N^0(x^+) \leq \rho$ and $x^+ \in \mathcal{S}$ so long as $|\theta| \leq \delta := \min \{ \delta_1, \delta_2 \}$ where $\delta_2 := \alpha_\theta^{-1}(\alpha_b^{-1}(\gamma_2))$. \square

Finally, Theorem 7.20 follows from Propositions 7.21 and 7.50 by combining the inequalities (7.21), (7.42), and (7.43).

Proof of Theorem 7.20. From Theorem 7.6, there exists $\alpha_2 \in \mathcal{K}_\infty$ such that (7.12a) holds for all $x \in \mathcal{S} \subseteq \mathcal{X}_N$, where $\alpha_1 \in \mathcal{K}_\infty$ is from Assumption 7.4. By Proposition 7.50, there exist $\alpha_b, \alpha_\theta \in \mathcal{K}$ such that (7.42) and (7.43) hold for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$. By Proposition 7.21, there exists $\delta > 0$ such that (7.21) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$, and \mathcal{S} is RPI for $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$. As in the proof of Proposition 7.21, we can combine (7.21), (7.42), and (7.43) to give (7.44) for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$, which is the desired cost decrease bound with $\sigma := \alpha_b \circ \alpha_\theta \in \mathcal{K}$. Thus, part (a) is established, and part (b) follows by Theorem 7.14. \square

Proof of Theorem 7.22. All the conditions of Theorems 7.7 and 7.20 are satisfied. Thus, there exists $c_2 > 0$ such that (7.13) holds for all $x \in \mathcal{S}$ with $c_1 := \underline{\sigma}(Q) > 0$. Moreover, we can substitute $\alpha_1(\cdot) := c_1 |\cdot|^2$ and $\alpha_2(\cdot) := c_2 |\cdot|^2$ into the proof of Theorem 7.20 to construct $\delta > 0$ and $\sigma \in \mathcal{K}$ such that (7.22) holds for all $x \in \mathcal{S}$ and $|\theta| \leq \delta$. Therefore, by Theorem 7.14, $x^+ = f_c(x, \theta)$, $|\theta| \leq \delta$ is ISES in \mathcal{S} . \square

7.A.4 Proofs of strong stability results

In this appendix we prove strong stability results from Section 7.3.

Quadratic cost MPC

We first consider results pertaining to strong stability of the quadratic cost MPC (Propositions 7.39, 7.40 and 7.43). Note that several preliminary results are required.

Proposition 7.51. *Suppose Assumptions 7.1 to 7.3 and 7.5 hold. Let $\rho > 0$ and $\mathcal{S} := \text{lev}_\rho V_N^0$.*

There exist $c_x, c_u > 0$ such that

$$|\hat{x}^0(k; x)| \leq c_x |x|, \quad \forall x \in \mathcal{S}, k \in \mathbb{I}_{0:N}. \quad (7.45)$$

$$|u^0(k; x)| \leq c_u |x|, \quad \forall x \in \mathcal{S}, k \in \mathbb{I}_{0:N-1}. \quad (7.46)$$

Proof. By Theorem 7.22, we have the upper bound (7.22a) for all $x \in \mathcal{S}$ and some $c_2 > 0$. Moreover, since Q, R, P_f are positive definite, we can write, for each $x \in \mathcal{S}$ and $k \in \mathbb{I}_{0:N-1}$,

$$\begin{aligned} \underline{\sigma}(Q) |\hat{x}^0(k; x)|^2 &\leq |\hat{x}^0(k; x)|_Q^2 \leq V_N^0(x) \leq c_2 |x|^2 \\ \underline{\sigma}(P_f) |\hat{x}^0(N; x)|^2 &\leq |\hat{x}^0(N; x)|_{P_f}^2 \leq V_N^0(x) \leq c_2 |x|^2 \\ \underline{\sigma}(R) |u^0(k; x)|^2 &\leq |u^0(k; x)|_R^2 \leq V_N^0(x) \leq c_2 |x|^2. \end{aligned}$$

Thus, with $c_x := \max \{ \sqrt{c_2 / \underline{\sigma}(Q)}, \sqrt{c_2 / \underline{\sigma}(P_f)} \}$ and $c_u := \sqrt{c_2 / \underline{\sigma}(R)}$, we have (7.45) and (7.46). \square

Proof of Proposition 7.39. Let $z := (x, u)$. By Proposition 7.49, for each $i \in \mathbb{I}_{1:n}$, there exists $\sigma_i \in \mathcal{K}_\infty$ such that

$$\left| \frac{\partial f_i}{\partial z}(z, \theta) - \frac{\partial \hat{f}_i}{\partial z}(\tilde{z}) \right| \leq \sigma_i(|(z - \tilde{z}, \theta)|) \quad (7.47)$$

for all $z, \tilde{z} \in \mathcal{S} \times \mathbb{U}$ and $\theta \in \mathbb{R}^{n_\theta}$. Next, let \mathcal{Z} denote the convex hull of $\mathcal{S} \times \mathbb{U}$. Then $tz \in \mathcal{Z}$ for all $t \in [0, 1]$ and $z \in \mathcal{Z}$. By Taylor's theorem (Apostol, 1974, Thm. 12.14), for each $i \in \mathbb{I}_{1:n}$ and $(z, \theta) \in \mathcal{Z} \times \Theta$, there exists $t_i(z, \theta) \in (0, 1)$ such that

$$f_i(z, \theta) - \hat{f}_i(z) = \left(\frac{\partial f_i}{\partial z}(t_i(z, \theta)z, \theta) - \frac{\partial \hat{f}_i}{\partial z}(t_i(z, \theta)z) \right) z. \quad (7.48)$$

Combining (7.47) and (7.48) gives, for each $(z, \theta) \in \mathcal{S} \times \mathbb{U} \times \mathbb{R}^{n_\theta}$,

$$|f(z, \theta) - \hat{f}(z)| \leq \sum_{i=1}^n |f_i(z, \theta) - \hat{f}_i(z)| \leq \sum_{i=1}^n \sigma_i(|\theta|)|z|$$

and therefore (7.27) holds with $\sigma_f := \sum_{i=1}^n \sigma_i$. \square

Proof of Proposition 7.40. By Proposition 7.51, there exists $c_u > 0$ such that

$$|\kappa_N(x)| = |u^0(0; x)| \leq c_u|x|$$

for all $x \in \mathcal{S}$. Moreover, by Proposition 7.39, there exists $\sigma_f \in \mathcal{K}_\infty$ such that

$$\begin{aligned} |f_c(x, \theta) - \hat{f}_c(x)| &\leq \sigma_f(|\theta|)|(x, \kappa_N(x))| \leq \sigma_f(|\theta|)(|x| + |\kappa_N(x)|) \\ &\leq \sigma_f(|\theta|)(|x| + c_u|x|) = \tilde{\sigma}_f(|\theta|)|x| \end{aligned}$$

for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$, where $\tilde{\sigma}_f := \sigma_f(1 + c_u) \in \mathcal{K}_\infty$. \square

Proposition 7.52. *Suppose Assumptions 7.1 to 7.3, 7.5, and 7.8 hold and assume \hat{f} is Lipschitz continuous on bounded sets. Let $\rho > 0$, $\mathcal{S} := \text{lev}_\rho V_N^0$, and $\Theta \subseteq \mathbb{R}^{n_\theta}$ be compact. There exist $c_{b,1}, c_{b,2} > 0$ such that, for each $x \in \mathcal{S}$ and $\theta \in \Theta$,*

$$|V_N(x^+, \tilde{\mathbf{u}}(x)) - V_N(\hat{x}^+, \tilde{\mathbf{u}}(x))| \leq 2c_{b,1}|x||x^+ - \hat{x}^+| + c_{b,2}|x^+ - \hat{x}^+|^2 \quad (7.49)$$

where $\hat{x}^+ := \hat{f}_c(x)$ and $x^+ := f_c(x, \theta)$.

Proof. First, we seek to prove following bound on the incurred terminal penalty $\tilde{V}_f(\cdot, \cdot) := V_f(\hat{\phi}(N; \cdot, \cdot))$: for each $x \in \mathcal{S}$ and $\theta \in \Theta$,

$$|\tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) - \tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x))| \leq c_{a,1}|x||x^+ - \hat{x}^+| + c_{a,2}|x^+ - \hat{x}^+|^2 \quad (7.50)$$

where $x^+ := f_c(x, \theta)$ and $\hat{x}^+ := \hat{f}_c(x)$.

Using the identity $|y|_M^2 - |\hat{y}|_M^2 = |y - \hat{y}|_M^2 + 2(y - \hat{y})^\top M \hat{y}$ for any positive definite M and y, \hat{y} of appropriate dimensions, we have, for each $x \in \mathcal{S}$ and $\theta \in \Theta$,

$$\begin{aligned} \tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) - \tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) &= |\hat{\phi}(N; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(N; \hat{x}^+, \tilde{\mathbf{u}}(x))|_{P_f}^2 \\ &\quad + 2(\hat{\phi}(N; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(N; \hat{x}^+, \tilde{\mathbf{u}}(x)))^\top P_f \times \hat{\phi}(N; \hat{x}^+, \tilde{\mathbf{u}}(x)). \end{aligned} \quad (7.51)$$

where $x^+ := f_c(x, \theta)$ and $\hat{x}^+ := \hat{f}_c(x)$. By Proposition 7.51, there exists $c_x > 0$ such that $|\hat{x}^0(k; x)| \leq c_x|x|$ and therefore

$$|\hat{\phi}(k; \hat{f}_c(x), \tilde{\mathbf{u}}(x))| = |\hat{x}^0(k+1; x)| \leq c_x|x| \quad (7.52)$$

for each $k \in \mathbb{I}_{0:N-1}$ and $x \in \mathcal{S}$. By Assumptions 7.3 and 7.5, we have, for each $x \in \mathbb{X}_f$,

$$\underline{\sigma}(P_f)|\hat{f}(x, \kappa_f(x))|^2 \leq V_f(\hat{f}(x, \kappa_f(x))) \leq V_f(x) - \underline{\sigma}(Q)|x|^2 \leq [\bar{\sigma}(P_f) - \underline{\sigma}(Q)]|x|^2$$

and therefore

$$|\hat{f}(x, \kappa_f(x))| \leq \gamma_f|x|$$

where $\gamma_f := \sqrt{[\bar{\sigma}(P_f) - \underline{\sigma}(Q)]/\underline{\sigma}(P_f)}$. Then, since $\hat{x}^0(N; x) \in \mathbb{X}_f$ and \mathbb{X}_f is positively

invariant for $x^+ = \hat{f}(x, \kappa_f(x))$, we have

$$|\hat{\phi}(N; \hat{f}_c(x), \tilde{\mathbf{u}}(x))| = |\hat{f}(\hat{x}^0(N; x), u^0(N; x))| \leq \gamma_f |\hat{x}^0(N; x)| \leq \gamma_f c_x |x| \quad (7.53)$$

for each $x \in \mathcal{S}$. Since $(\mathcal{S}, \mathbb{U}, \Theta)$ are each bounded and f is continuous, $\mathcal{S}_0 := f(\mathcal{S}, \mathbb{U}, \Theta)$ is bounded. But this means $\mathcal{S}_{k+1} := \hat{f}(\mathcal{S}_k, \mathbb{U})$ is bounded for each $k \in \mathbb{I}_{\geq 0}$ (by induction), so $\bar{\mathcal{S}} := \bigcup_{k=0}^N \mathcal{S}_k$ is also bounded. Since \hat{f} is Lipschitz continuous on bounded sets, there exists $L_f > 0$ such that $|\hat{f}(x, u) - \hat{f}(\tilde{x}, \tilde{u})| \leq L_f |x - \tilde{x}, u - \tilde{u}|$ for all $x, \tilde{x} \in \bar{\mathcal{S}}$ and $u, \tilde{u} \in \mathbb{U}$. Then, for each $\theta \in \Theta$, we have

$$\begin{aligned} & |\hat{\phi}(k+1; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(k+1; \hat{x}^+, \tilde{\mathbf{u}}(x))| \\ &= |\hat{f}(\hat{\phi}(k; x^+, \tilde{\mathbf{u}}(x)), u^0(k; x)) - \hat{f}(\hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x)), u^0(k; x))| \\ &\leq L_f |\hat{\phi}(k; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x))| \end{aligned}$$

for each $k \in \mathbb{I}_{0:N-1}$, and therefore

$$|\hat{\phi}(k; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x))| \leq L_f^k |x^+ - \hat{x}^+|, \quad (7.54)$$

for each $k \in \mathbb{I}_{0:N}$, where $\hat{x}^+ := \hat{f}_c(x)$ and $x^+ := f_c(x, \theta)$. Finally, combining (7.51), (7.53), and (7.54), we have (7.50) for all $x \in \mathcal{S}$ and $\theta \in \Theta$, where $c_{a,1} := 2L_f^N \gamma_f c_x \bar{\sigma}(P_f)$ and $c_{a,2} := L_f^{2N} \bar{\sigma}(P_f)$.

Moving on to the proof of (7.49), we have, for each $x \in \mathcal{S}$ and $\theta \in \Theta$,

$$\begin{aligned} V_N(x^+, \tilde{\mathbf{u}}(x)) - V_N(N; \hat{x}^+, \tilde{\mathbf{u}}(x)) &= \sum_{k=0}^{N-1} |\hat{\phi}(k; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x))|_Q^2 \\ &\quad + 2(\hat{\phi}(k; x^+, \tilde{\mathbf{u}}(x)) - \hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x)))^\top Q \hat{\phi}(k; \hat{x}^+, \tilde{\mathbf{u}}(x)) \\ &\quad + \tilde{V}_f(x^+, \tilde{\mathbf{u}}(x)) - \tilde{V}_f(\hat{x}^+, \tilde{\mathbf{u}}(x)) \end{aligned} \quad (7.55)$$

where $\hat{x}^+ := \hat{f}_c(x)$ and $x^+ := f_c(x, \theta)$, and combining (7.50), (7.52), (7.54), and (7.55), we have (7.49) with $c_{b,1} := c_{a,1} + 2\bar{\sigma}(Q) \sum_{k=0}^{N-1} L_f^k c_x$ and $c_{b,2} := c_{a,2} + \bar{\sigma}(Q) \sum_{k=0}^{N-1} L_f^{2k}$. \square

Proof of Proposition 7.43. By Proposition 7.39, there exists $\tilde{\sigma}_f \in \mathcal{K}_\infty$ such that (7.28) for all $x \in \mathcal{S}$. Moreover, by Proposition 7.52, there exist $c_{b,1}, c_{b,2} > 0$ such that (7.49) for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$, where $x^+ := f_c(x, \theta)$ and $\hat{f}_c(x)$. Finally, (7.28) and (7.49) imply (7.31) for all $x \in \mathcal{S}$ and $\theta \in \mathbb{R}^{n_\theta}$, where $\sigma_V(\cdot) := c_{b,1}\tilde{\sigma}_f(\cdot) + c_{b,2}[\tilde{\sigma}_f(\cdot)]^2 \in \mathcal{K}_\infty$. \square

General nonlinear MPC

Next, we move on to the general nonlinear MPC results (Propositions 7.26, 7.41, 7.42 and 7.44). Again, several preliminary results are required.

Proposition 7.53. *For each $\alpha \in \mathcal{K}$ and $\gamma \in \mathcal{K}^2$, let $\gamma_1(s, t) := \alpha(\gamma(s, t))$, $\gamma_2(s, t) := \gamma(\alpha(s), t)$, and $\gamma_3(s, t) := \gamma(s, \alpha(t))$ for each $s, t \geq 0$. Then $\gamma_1, \gamma_2, \gamma_3 \in \mathcal{K}^2$.*

Proof. This fact follows directly from the closure of \mathcal{K} under composition (Kellett, 2014). For example, for each $s \geq 0$, we have $\gamma_2(\cdot, s) = \gamma(\alpha(\cdot), s) \in \mathcal{K}$ by closure under composition, $\gamma_2(s, \cdot) = \gamma(\alpha(s), \cdot) \in \mathcal{K}$ trivially, and γ_2 is continuous as it is a composition of continuous functions. \square

Proof of Proposition 7.41. Without loss of generality, assume \mathcal{S} and Θ contain the origin. By assumption, $C := \mathcal{S} \times \mathbb{U} \times \Theta$ is compact, and by Proposition 7.49, there exists $\sigma_f \in \mathcal{K}_\infty$ such

that

$$|f(x, u, \theta) - f(\tilde{x}, \tilde{u}, \tilde{\theta})| \leq \sigma_f(|(x, u, \theta) - (\tilde{x}, \tilde{u}, \tilde{\theta})|) \quad (7.56)$$

for all $(x, u, \theta), (\tilde{x}, \tilde{u}, \tilde{\theta}) \in C$. Specializing (7.56) to $(\tilde{x}, \tilde{u}, \tilde{\theta}) = (x, u, 0) \in C$ gives

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq \sigma_f(|\theta|) \quad (7.57)$$

for all $(x, u, \theta) \in C$. On the other hand, specializing (7.56) to $(\tilde{x}, \tilde{u}, \tilde{\theta}) = (0, 0, \theta) \in C$ gives

$$|f(x, u, \theta)| = |f(x, u, \theta) - f(0, 0, \theta)| \leq \sigma_f(|(x, u)|)$$

and therefore

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq |f(x, u, \theta)| + |\hat{f}(x, u)| \leq 2\sigma_f(|(x, u)|) \quad (7.58)$$

for all $(x, u, \theta) \in C$. Combining (7.57) and (7.58) gives

$$|f(x, u, \theta) - \hat{f}(x, u)| \leq \min\{2\sigma_f(|(x, u)|), \sigma_f(|\theta|)\}$$

for all $(x, u, \theta) \in C$, which is an upper bound that is clearly continuous, nondecreasing in each $|x|$ and $|\theta|$, and zero if either $|x|$ or $|\theta|$ is zero. To make the upper bound strictly increasing, pick any $\sigma_1, \sigma_2 \in \mathcal{K}$ and let $\gamma_f(s, t) := \min\{2\sigma_f(s), \sigma_f(t)\} + \sigma_1(s)\sigma_2(t)$ for each $s, t \geq 0$. Then $\gamma_f \in \mathcal{K}^2$, and (7.29) holds for all $(x, u, \theta) \in C$. \square

Proof of Proposition 7.42. First, we have $\gamma_f \in \mathcal{K}^2$ satisfying (7.29) for all $x \in \mathcal{S}$, $u \in \mathbb{U}$, and $\theta \in \Theta$ by Proposition 7.41. Using the bounds (7.9) and (7.20a) with $u = \kappa_N(x)$, we have, for each $x \in \mathcal{X}_N$,

$$\alpha_1(|\kappa_N(x)|) \leq \ell(x, \kappa_N(x)) \leq V_N^0(x) \leq \alpha_2(|x|)$$

and thus $|\kappa_N(x)| \leq \alpha_\kappa(|x|)$, where $\alpha_\kappa := \alpha_1^{-1} \circ \alpha_2 \in \mathcal{K}_\infty$. Then, for each $x \in \mathcal{S}$ and $\theta \in \Theta$,

$$\begin{aligned} |f_c(x, \theta) - \hat{f}_c(x)| &\leq \gamma_f(|(x, \kappa_N(x))|, |\theta|) \\ &\leq \gamma_f(|x| + |\kappa_N(x)|, |\theta|) \\ &\leq \gamma_f(|x| + \alpha_\kappa(|x|), |\theta|) = \tilde{\gamma}_f(|x|, |\theta|). \end{aligned}$$

where $\tilde{\gamma}_f(s, t) := \gamma_f(s + \alpha_\kappa(s), t)$ for each $s, t \geq 0$. Then $(\cdot) + \alpha_\kappa(\cdot) \in \mathcal{K}_\infty$, and $\tilde{\gamma}_f \in \mathcal{K}^2$ by Proposition 7.53. Finally, (7.30) holds for all $x \in \mathcal{S}$ and $\theta \in \Theta$. \square

Proof of Proposition 7.44. By Proposition 7.49, there exists $\alpha_b \in \mathcal{K}_\infty$ such that

$$V_N(x_1, \mathbf{u}_1) - V_N(x_2, \mathbf{u}_2) \leq \alpha_b(|(x_1 - x_2, \mathbf{u}_2 - \mathbf{u}_2)|) \quad (7.59)$$

for all $(x, \mathbf{u}), (\tilde{x}, \tilde{\mathbf{u}}) \in f(\mathcal{S}, \mathbb{U}, \Theta) \times \mathbb{U}^N$. Specializing (7.59) to $x_1 = x^+ := f_c(x, \theta)$, $x_2 = \hat{x}^+ := f_c(x)$, and $\mathbf{u}_1 = \mathbf{u}_2 = \tilde{\mathbf{u}}(x)$ gives

$$|V_N(x^+, \tilde{\mathbf{u}}(x)) - V_N(\hat{x}^+, \tilde{\mathbf{u}}(x))| \leq \alpha_b(|x^+ - \hat{x}^+|) \quad (7.60)$$

for each $x \in \mathcal{S}$ and $\theta \in \Theta$. By Proposition 7.42 there exists $\tilde{\gamma}_f \in \mathcal{K}^2$ satisfying (7.30) for all $x \in \mathcal{S}$ and $\theta \in \Theta$. Finally, combining (7.30) and (7.60) gives (7.32) with $\gamma_V(s, t) := \alpha_b(\tilde{\gamma}_f(s, t))$ for all $s, t \geq 0$, where $\gamma_V \in \mathcal{K}^2$ by Proposition 7.53. \square

Proof of Proposition 7.26. Let

$$\tilde{\gamma}(s, t) := \sup_{\tilde{s} \in (0, s)} \frac{\gamma(\tilde{s}, t)}{\alpha(\tilde{s})}$$

for each $s, t > 0$, so that

$$L := \limsup_{s \rightarrow 0^+} \frac{\gamma(s, \tau)}{\alpha(s)} = \lim_{s \rightarrow 0^+} \tilde{\gamma}(s, \tau).$$

Suppose $L < 1$. Then there exists $\delta_0 > 0$ such that $|\tilde{\gamma}(s, \tau) - L| < 1 - L$ for all $s \in (0, \delta_0]$. But $\tilde{\gamma}(s, t) \geq 0$ and $L \geq 0$ for all $s, t > 0$, so $\tilde{\gamma}(s, \tau) < 1$ for all $s \in (0, \delta_0]$ by the reverse triangle inequality. Therefore

$$\frac{\gamma(s, t)}{\alpha(s)} \leq \frac{\gamma(s, \tau)}{\alpha(s)} \leq \tilde{\gamma}(s, \tau) < 1$$

and $\gamma(s, t) < \alpha(s)$ for all $s \in (0, \delta_0]$ and $t \in [0, \tau]$.

If $\delta_0 \geq \rho$, the proof is complete with $\delta := \tau$. Otherwise, we must enlarge the interval in s by shrinking the interval in t . For each $t \in (0, \tau]$, let

$$\gamma_0(t) := \inf \{ s > 0 \mid \gamma(s, t) \geq \alpha(s) \}.$$

Since $\gamma(s, t) \leq \gamma(s, \tau) < \alpha(s)$ for each $s \in (0, \delta_0]$ and $t \in [0, \tau]$, we have $\gamma_0(t) > 0$. Then, by continuity of α and γ , $\gamma_0(t)$ must be equal to the first nonzero point of intersection if it exists. Otherwise $\gamma_0(t)$ is infinite. Note that γ_0 is a strictly decreasing function since, for any $t \in (0, \tau]$, we have $\gamma(\gamma_0(t), t') < \gamma(\gamma_0(t), t) = \alpha(\gamma_0(t))$ for all $t' \in (0, t)$. Moreover, $\lim_{t \rightarrow 0^+} \gamma_0(t) = \infty$ since, if γ_0 was upper bounded by some $\bar{\gamma} > 0$, we could take $\gamma(\bar{\gamma}, t) \geq \alpha(\bar{\gamma}) > 0$ for all $t \in (0, \tau]$, a contradiction of the fact that $\gamma(s, \cdot) \in \mathcal{K}$ for all $s > 0$. Then there must exist $\delta > 0$ such that $\gamma_0(\delta) > \rho$ and therefore $\gamma(s, t) < \alpha(s)$ for all $s \in (0, \rho]$ and $t \in [0, \delta]$. \square

7.B Additional examples details

7.B.1 Strong asymptotic stability counterexample

Consider the plant (7.33) and MPC defined in Section 7.5.1. We aim to show the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta$ is RES with $\delta = 3$, but not inherently strongly

stabilizing for any $\delta > 0$. By Lipschitz continuity of x^2 on bounded sets and 1/2-Hölder continuity of $\sqrt{|x|}$,

$$|x^2 - y^2| \leq 4|x - y|, \quad \forall x, y \in [-2, 2], \quad (7.61)$$

$$|\sigma(x) - \sigma(y)| \leq 2\sqrt{|x - y|}, \quad \forall x, y \in \mathbb{R}. \quad (7.62)$$

To show (7.61), note that, for each $\delta > 0$, we have

$$|x^2 - y^2| = |x + y||x - y| \leq 2\delta|x - y|$$

for all $x, y \in [-\delta, \delta]$, and take $\delta = 2$ to give (7.61). For (7.62), we first show $\sqrt{(\cdot)}$ is 1/2-Hölder continuous on $\mathbb{R}_{\geq 0}$:

$$|\sqrt{x} - \sqrt{y}| = \frac{|x - y|}{\sqrt{x} + \sqrt{y}} \leq \frac{|x - y|}{\sqrt{x} + \sqrt{y}} = \sqrt{|x - y|} \frac{\sqrt{|x - y|}}{\sqrt{x} + \sqrt{y}} \leq \sqrt{|x - y|}$$

for all $x, y \geq 0$, where the last inequality follows by the triangle inequality. Then we automatically get $|\sigma(x) - \sigma(y)| \leq \sqrt{|x - y|}$ if $x, y \geq 0$. On the other hand, if $x \geq 0$ and $y \leq 0$, we have

$$|\sigma(x) - \sigma(y)| = |\sqrt{x} + \sqrt{y}| \leq \sqrt{x} + \sqrt{-y} \leq 2\sqrt{x - y}.$$

Finally, flipping the signs of the prior arguments gives (7.62).

First, we derive the control law. The terminal set can be reached in a single move if and only if $|x| \leq 2$, so we have the steerable set $\mathcal{X}_1 = [-2, 2]$. Consider the problem *without* the terminal constraint. The objective is

$$V_1(x, u) = x^2 + u^2 + 4|x + u|$$

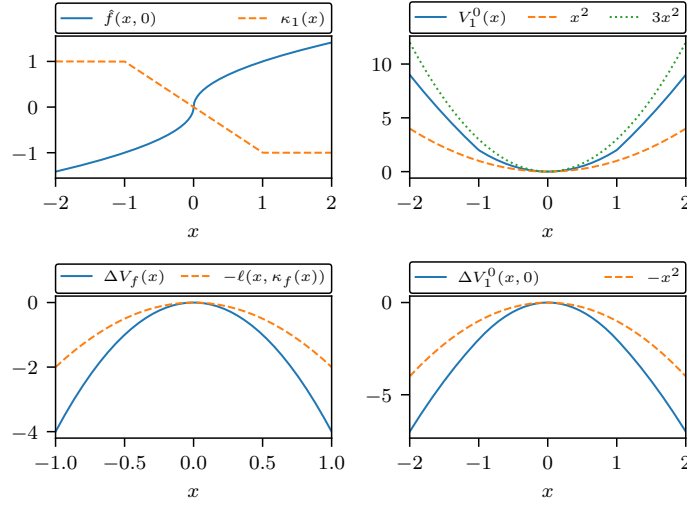


Figure 7.8: For the MPC of (7.33), plots of (top left) the open-loop dynamics and control law, (bottom left) the terminal cost difference, (top right) the optimal value function, and (bottom right) the cost difference, each with the relevant (nominal) bounds from Assumption 7.3 and (7.22).

which is increasing in u if $x > 1$ and $|u| \leq 1$, and decreasing in u if $x < -1$ and $|u| \leq 1$. Thus $V_1(x, \cdot)$ is minimized (over $|u| \leq 1$) by $\mathbf{u}^0(x) = -\text{sgn}(x)$ for all $x \notin [-1, 1]$. On the other hand, if $|x| \leq 1$, then $V_1(x, \cdot)$ is decreasing on $[-1, -x)$ and increasing on $(-x, 1]$. Thus $V_1(x, \cdot)$ is minimized (over $|u| \leq 1$) by $\mathbf{u}^0(x) = -x$ so long as $|x| \leq 1$. In summary, we have the control law $\kappa_1(x) := -\text{sat}(x)$. But

$$|\hat{f}(x, \kappa_1(x))| = \begin{cases} 0, & |x| \leq 1 \\ |x - \text{sgn}(x)| = |x| - 1, & 0 < |x| \leq 2 \end{cases}$$

so $u = \kappa_1(x)$ drives each state in $\mathcal{X}_1 = [-2, 2]$ to the terminal constraint $\mathbb{X}_f = [-1, 1]$. Therefore κ_1 is also the control law of the problem *with* the terminal constraint. The control law κ_1 is plotted, along with the unforced dynamics $\hat{f}(\cdot, 0)$, against $x \in \mathcal{X}_1$ in Figure 7.8 (top left).

Assumptions 7.1 and 7.4 are satisfied by definition, Assumption 7.2 is satisfied with $c_f :=$

8, and Assumption 7.3 is satisfied with $\kappa_f(x) := -x$ since $\hat{f}(x, \kappa_f(x)) = 0$ and

$$\Delta V_f(x) := V_f(\hat{f}(x, \kappa_f(x))) - V_f(x) = -4x^2 \leq -2x^2 = -\ell(x, \kappa_f(x))$$

for all $x \in \mathbb{X}_f$. See Figure 7.8 (bottom left) for plots of ΔV_f and $-\ell(\cdot, \kappa_f(\cdot))$. Therefore, by Theorem 7.20, the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta$ is RAS on $\mathcal{X}_1 = [-2, 2]$ with ISS Lyapunov function V_1^0 for some $\delta > 0$. Our next goal is to find such a $\delta > 0$.

First, however, let us establish that V_1^0 is a Lyapunov function for the modeled closed-loop $x^+ = \hat{f}(x, \kappa_1(x))$ in $\mathcal{X}_1 = [-2, 2]$. We already have $V_1^0(x) \geq x^2$ for all $|x| \leq 2$. For the upper bound, we have

$$V_1^0(x) = V_1(x, \kappa_1(x)) = \begin{cases} 2x^2, & |x| \leq 1, \\ x^2 + 4|x| - 3, & 1 < |x| \leq 2 \end{cases}$$

for each $|x| \leq 2$. But the polynomials $-2x^2 \pm 4x - 3$ have no real roots, so $4|x| - 3 < 2x^2$, and the above inequality gives $V_1^0(x) \leq 3x^2$ for all $|x| \leq 2$. Moreover, by (7.40), we have $\Delta V_1^0(x, \theta) \leq -x^2$, so $x^+ = \hat{f}(x, \kappa_1(x))$ is in fact *exponentially* stable on $\mathcal{X}_1 = [-2, 2]$. We plot V_1^0 and $\Delta V_1^0(\cdot, 0) := V_1^0(\hat{f}(\cdot, \kappa_1(\cdot))) - V_1^0(\cdot)$, along with their exponential Lyapunov bounds, in Figure 7.8 (right).

For robust positive invariance, let $|x| \leq 2$, $\theta \in \mathbb{R}$, $x^+ := f(x, \kappa_1(x), \theta)$, $\hat{x}^+ := \hat{f}(x, \kappa_1(x))$ and note that

$$x^+ = \sigma(\sigma^{-1}(\hat{x}^+) - \theta \text{sat}(x))$$

where $\sigma^{-1}(x) = \text{sgn}(x)|x|^2$, and therefore

$$|x^+| \leq \sqrt{|\hat{x}^+|^2 + |\theta| |\text{sat}(x)|} \leq \sqrt{1 + |\theta|}.$$

Then $|x^+| \leq 2$ so long as $|\delta| \leq 3$, so $\mathcal{X}_1 = [-2, 2]$ is RPI for $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 3$.

By continuity of f , V_1^0 , and κ_1 and Proposition 7.49, there exists $\sigma \in \mathcal{K}_\infty$ such that $|V_1^0(x^+) - V_1^0(\hat{x}^+)| \leq \sigma(|\theta|)$ and therefore $V_1^0(x^+) \leq V_1^0(\hat{x}^+) + |V_1^0(x^+) - V_1^0(\hat{x}^+)| \leq V_1^0(x) - x^2 + \sigma(|\theta|)$ for all $|x| \leq 2$ and $|\theta| \leq 3$, where $x^+ := f(x, \kappa_1(x), \theta)$ and $\hat{x}^+ := \hat{f}(x, \kappa_1(x))$. Therefore $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 3$ is not only RAS, but RES on \mathcal{X}_1 by Theorem 7.14.

We now aim to show strong stability is *not* achieved. For simplicity, we consider $\mathcal{S} := \text{lev}_2 V_1^0 = [-1, 1] = \mathbb{X}_f$ as the candidate basin of attraction. Let $|x| \leq 1$, $|\theta| \leq 3$, $x^+ := f(x, \kappa_1(x), \theta)$, and $\hat{x}^+ := \hat{f}(x, \kappa_1(x))$. Moreover, $\ell(x, \kappa_1(x)) \geq 2|x|^2 =: \alpha_3(|x|)$. Next, we have $\kappa_1(x) = -x$, $x^+ = \sigma(x\theta)$, and $\hat{x}^+ = 0$. Therefore

$$\begin{aligned} |V_1(x^+, \tilde{\mathbf{u}}(x)) - V_1(\hat{x}^+, \tilde{\mathbf{u}}(x))| &= |(x^+)^2 + 4|x^+|| \leq |x^+|^2 + 4|x^+| \\ &\leq |x||\theta| + 4\sqrt{|x||\theta|} =: \gamma_V(|x|, |\theta|) \end{aligned}$$

where $\gamma_V \in \mathcal{K}^2$. For each $t > 0$, we have $\frac{\gamma_V(s, t)}{\alpha_3(s)} = (st + 4\sqrt{st})/(2s^2) = t/(2s) + 2\sqrt{t}/s^{3/2}$, so $\lim_{s \rightarrow 0^+} \frac{\gamma_V(s, t)}{\alpha_3(s)} = \infty$ for all $t > 0$, and (7.24) is not satisfied.

As mentioned in the main text, (7.24) is sufficient but not necessary. But the cost difference curve is positive definite, as

$$\Delta V_1^0(x, \theta) = 2[\sigma(\theta x)]^2 - 2x^2 = 2(|\theta| - |x|)|x| > 0$$

for any $0 < |x| < |\theta| \leq 1$. In other words, θ can be arbitrarily small but nonzero, and the cost difference curve will remain positive definite near the origin.

7.B.2 Nonlinearizable yet inherently strongly stabilizing

Consider the plant (7.34) and MPC defined in Section 7.5.2. We aim to show the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta$ is RES in \mathcal{X}_1 with $\delta = 1$, and SES with $\delta = 1/2$.

To derive the control law, we first consider the problem *without* the terminal constraint (i.e., $\mathbb{X}_f = \mathbb{R}$). We have the objective

$$V_1(x, u) = x^2 + u^2 + 4(x + (1/2)\gamma(x) + u)^2.$$

Taking the partial derivative in u ,

$$\frac{\partial V_1}{\partial u}(x, u) = 8x + 4\gamma(x) + 10u$$

and setting that to zero gives the optimal input

$$\mathbf{u}^0(x) = -g(x) := -(4/5)x - (2/5)\gamma(x)$$

whenever $|g(x)| \leq 1$. Otherwise the solution saturates at $\mathbf{u}^0(x) = -\text{sgn}(g(x))$, so we have $\mathbf{u}^0(x) = \kappa_1(x) := -\text{sat}(g(x))$ for all $|x| \leq 2$.

To see where the control law $\kappa_1(x)$ saturates, first note

$$\frac{d^2 g}{dx^2}(x) = \frac{2}{5} \frac{d^2 \gamma}{dx^2}(x) = -\frac{8\pi^2 \sin(2\pi/x)}{5|x|^3}$$

for all $x \neq 0$, so $g(x)$ is strictly concave on $x \in [1/(n - 1/2), 1/n]$ and strictly convex on $x \in [1/n, 1/(n + 1/2)]$ for each $n \in \mathbb{I}$. Therefore $g(x)$ achieves a local maximum on each $x \in [1/(n - 1/2), 1/n]$, and the maximum is strictly decreasing with n . The last, and greatest, of these local maxima on $|x| \leq 2$ is achieved on $2/3 \leq x \leq 1$. Through numerical optimization, we find $\max_{0 \leq x \leq 1} g(x) = \max_{2/3 \leq x \leq 1} g(x) \approx 0.9849$. By strict convexity of $g(x)$ on $x \in [1, 2]$, $g(1) = 4/5$, and $g(2) = 8/5$, we have $\max_{1 \leq x \leq 2} g(x) = g(2) = 8/5$. Therefore $g(x)$ intersects the horizontal line at $u = 1$ exactly once over $x \in [-2, 2]$, and it does so at some $x^* \in [1, 2]$, which we can numerically verify is $x^* \approx 1.6989$. By symmetry,

$g(x)$ intersects $u = -1$ at $-x^*$. Finally, because $g(x)$ is strictly convex (concave) on $[1, 2]$ ($[-2, -1]$), it saturates on $(x^*, 2]$ (and $[-2, -x^*)$) and we have

$$\kappa_1(x) = \begin{cases} -(4/5)x - (2/5)\gamma(x), & |x| \leq x^*, \\ -\operatorname{sgn}(x), & x^* < |x| \leq 2. \end{cases}$$

For the problem *with* the terminal constraint, we have

$$|\hat{f}(x, \kappa_1(x))| = |(1/5)x + (3/5)\gamma(x)| \leq (1/5)x + (3/5)|\gamma(x)| \leq 4/5$$

for each $x \in [0, 1]$,

$$|\hat{f}(x, \kappa_1(x))| = |(1/5)x + (3/5)\gamma(x)| \leq |(1/5)x - (3/5)|\gamma(x)|| \leq (2/5)|x| \leq 4/5$$

for each $x \in [1, x^*]$, and

$$|\hat{f}(x, \kappa_1(x))| = |x + (1/2)\gamma(x) - 1| \leq |x| - 1 - (1/2)|\gamma(x)| \leq |x| - 1 \leq 1$$

for each $x \in [x^*, 2]$, where we have used the fact that $\gamma(x) \leq 0$ for all $x \in [1, 2]$. Therefore $|\hat{f}(x, \kappa_1(x))| \leq 1$ for all $x \in [0, 2]$, and the same holds for all $x \in [-2, 0]$ by symmetry. Therefore the terminal constraint $\mathbb{X}_f = [-1, 1]$ is automatically satisfied by the unconstrained control law, so $\kappa_1(x)$ is also the control law for the MPC *with* the terminal constraint. In Figure 7.9 (top left), we plot κ_1 and $\hat{f}(\cdot, 0)$ on \mathcal{X}_1 .

Assumptions 7.1 and 7.5 are satisfied by definition, and Assumption 7.2 is satisfied with $c_f := 4$. Let $\kappa_f(x) := -(1/2)(x + \gamma(x))$ for all $|x| \leq 1$. Then

$$|\kappa_f(x)| \leq (1/2)(|x| + |\gamma(x)|) \leq |x| \leq 1$$




Figure 7.9: For the MPC of (7.34), we plot as a function of x (top left) the open-loop dynamics and control law, (bottom left) the terminal cost difference, (top right) the optimal value function, and (bottom right) the cost difference, each with the relevant (nominal) bounds from Assumption 7.3 and (7.22).

for all $|x| \leq 1$, so $u = \kappa_f(x)$ is feasible in the terminal constraint. Moreover, $\hat{f}(x, \kappa_f(x)) = (1/2)x$, so

$$\Delta V_f(x) := V_f(\hat{f}(x, \kappa_f(x))) - V_f(x) + \ell(x, \kappa_f(x)) = -2x^2 + |\kappa_f(x)|^2 \leq -x^2 \leq 0$$

and Assumption 7.3 is satisfied. See Figure 7.9 (bottom left) for plots of ΔV_f and $-\ell(\cdot, \kappa_f(\cdot))$. By Theorem 7.22, the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta$ is RES on \mathcal{X}_1 with the ISS Lyapunov function V_1^0 for some $\delta > 0$. Our next aim is to find such a $\delta > 0$.

Let $|x| \leq 2$, $\theta \in \mathbb{R}$, $x^+ := f(x, \kappa_1(x), \theta)$, and $\hat{x}^+ := \hat{f}(x, \kappa_1(x))$. Then $x^+ = \hat{x}^+ + \theta \kappa_1(x)$, and we have

$$|x^+| \leq |\hat{x}^+| + |\theta| |\kappa_1(x)| \leq 1 + |\theta|$$

for all $\theta \in \mathbb{R}$. But this means $|x^+| \leq 2$ for all $|\theta| \leq 1$, so \mathcal{X}_1 is RPI for $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 1$. Continuity of f , ℓ , V_f , and κ_1 implies continuity of $V_1^0(f_c(\cdot, \cdot))$, at least for all $|x| \leq 2$ and $|\theta| \leq 1$ on which the function is well-defined. Then, by Proposition 7.49, there exists

$\sigma \in \mathcal{K}_\infty$ such that, if $|\theta| \leq 1$, we have $|V_1^0(x^+) - V_1^0(\hat{x}^+)| \leq \sigma(|\theta|)$, and therefore $V_1^0(x^+) \leq V_1^0(\hat{x}^+) + |V_1^0(x^+) - V_1^0(\hat{x}^+)| \leq V_1^0(x) - x^2 + \sigma(|\theta|)$. Finally, $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq 1$ is RES in $\mathcal{X}_1 = [-2, 2]$ by Theorem 7.14.

Next, we aim to show the MPC is inherently strongly stabilizing via Assumption 7.33 and Theorem 7.34. Consider the candidate Lyapunov function $V(x) := x^2$ for all $|x| \leq 2$ and $V(x) := \infty$ otherwise, and let $\rho \geq 4$, $\mathcal{S} := \text{lev}_\rho V = [-2, 2] = \mathcal{X}_1$, and $\delta_0 := 1$. If we can show Assumption 7.33(a,b) hold with these ingredients, then Assumption 7.33 will hold for all $\rho > 0$. Assumption 7.33(a) and (7.25a) are already satisfied with $|\theta| \leq \delta_0 = 1$, and \mathcal{S} is RPI, but it remains to construct the bound (7.25b). Throughout this derivation, let $x^+ := f(x, \kappa_1(x), \theta)$ and $\hat{x}^+ := \hat{f}(x, \kappa_1(x))$.

First, suppose $|x| \leq x^*$ and $|\theta| \leq 1$. Then the controller does not saturate, i.e., $\kappa_1(x) = -0.8x - 0.4\gamma(x)$, and we have in the nominal case $\hat{x}^+ = 0.2x + 0.1\gamma(x)$, $|\hat{x}^+| \leq 0.3|x|$, and

$$V(\hat{x}^+) - V(x) = |\hat{x}^+|^2 - |x|^2 \leq -0.91|x|^2. \quad (7.63)$$

Next, consider the identity

$$y^2 - z^2 = 2z(y - z) + (y - z)^2 \quad (7.64)$$

for all $y, z \in \mathbb{R}$. We have $x^+ = (0.2 - 0.8\theta)x + (0.1 - 0.4\theta)\gamma(x)$, so $|x^+ - \hat{x}^+| = |0.8\theta x + 0.4\theta\gamma(x)| \leq 1.2|\theta||x|$, and (7.64) implies

$$|V(x^+) - V(\hat{x}^+)| \leq 0.72|\theta||x|^2 + 1.44|\theta|^2|x|^2. \quad (7.65)$$

Next, suppose $x^* < x \leq 2$ and $|\theta| \leq 1$. Then the controller always saturates, i.e., $\kappa_1(x) = -1$. Since $\gamma(\tilde{x}) \leq 0$ for all $1 \leq \tilde{x} \leq 2$, we have $0 \leq 0.5x + 0.5\gamma(x) \leq 0.5x \leq 1$ and $\hat{x}^+ = x + 0.5\gamma(x) - 1 \leq 0.5x$. Moreover, $x - 1 > x^* - 1 > 0$, so $\hat{x}^+ = x + 0.5\gamma(x) - 1 >$

$0.5\gamma(x) \geq -0.5x$. Then we have $|\hat{x}^+| \leq 0.5|x|$ and

$$V(\hat{x}^+) - V(x) = |\hat{x}^+|^2 - |x|^2 \leq -0.75|x|^2. \quad (7.66)$$

Moreover, $|x^+ - \hat{x}^+| = |\theta|$ and (7.64) implies

$$|V(x^+) - V(\hat{x}^+)| \leq (1/x^*)|\theta||x|^2 + (1/x^*)^2|\theta|^2|x|^2 \quad (7.67)$$

where we have used the fact that $|x|/x^* > 1$. By symmetry, (7.66) and (7.67) also hold for $-2 \leq x < -x^*$.

Combining (7.63), (7.66), (7.67), and (7.68), we have

$$V(x^+) \leq V(x) - a_3|x|^2 + \sigma_V(|\theta|)|x|^2 \quad (7.68)$$

for all $x \in \mathcal{X}_N$, where $a_3 := 0.75$ and $\sigma_V(t) := \max \{ 0.72t + 1.44t^2, (2/x^*)t + (1/x^*)^2t^2 \}$ and Assumption 7.33 is satisfied. Finally, by Theorem 7.34 (and its proof), the closed-loop system $x^+ = f(x, \kappa_1(x), \theta)$, $|\theta| \leq \delta$ is SES in $\mathcal{X}_N = [-2, 2]$ for any $\delta \in (0, \sigma_V^{-1}(a_3))$. Thus, it suffices to take $|\theta| \leq \delta = 0.5$ since

$$\sigma_V(0.5) = \max \{ 0.72, 0.3809 \dots \} = 0.72 < 0.75 = a_3.$$

7.B.3 Upright pendulum

Consider the plant (7.36) and MPC defined in Section 7.5.3. It is noted in the main text that Assumptions 7.1, 7.2, 7.5, 7.8, and 7.9 are automatically satisfied. To design P_f and show

Assumption 7.3 holds, consider the linearization

$$x^+ = \underbrace{\begin{bmatrix} 1 & 0.1 \\ 0.1 & 1 \end{bmatrix}}_{=:A} x + \underbrace{\begin{bmatrix} 0 \\ 5 \end{bmatrix}}_{=:B} u \quad (7.69)$$

and the feedback gain $K := \begin{bmatrix} 2 & 2 \end{bmatrix}$, which stabilizes (7.69) because $A_K := A - BK = \begin{bmatrix} -1 & 0.1 \\ -0.9 & 0 \end{bmatrix}$ has eigenvalues of 0.9 and 0.1. Numerically solving the Lyapunov equation

$$A_K^\top P_f A_K - P_f = -2Q_K$$

where $Q_K := Q + K^\top R K = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix}$, we have a unique positive definite solution $P_f := \begin{bmatrix} 31.133\dots & 10.196\dots \\ 10.196\dots & 10.311\dots \end{bmatrix}$. Using the inequality $|\sin x_1 - x_1| \leq (1/6)|x_1|^3$ for all $x_1 \in \mathbb{R}$, we have

$$\begin{aligned} & |V_f(\hat{f}(x, -Kx)) - V_f(A_K x)| \\ &= 2x^\top A_K^\top P_f \begin{bmatrix} 0 \\ \Delta(\sin x_1 - x_1) \end{bmatrix} + [P_f]_{22} \Delta^2 (\sin x_1 - x_1)^2 \\ &\leq b|x|^4 + a|x|^6 \end{aligned}$$

for all $x \in \mathbb{R}^2$, where $a := \frac{[P_f]_{22} \Delta^2}{36} = 2.8643\dots \times 10^{-3}$ and $b := \frac{\Delta |A_K^\top P_f \begin{bmatrix} 0 \\ 1 \end{bmatrix}|}{3} = 0.045675\dots$

Moreover, $\underline{\sigma}(Q_K) = 1$, so

$$\begin{aligned} & V_f(\hat{f}(x, -Kx)) - V_f(x) + \ell(x, -Kx) \\ &= -|x|_{Q_K}^2 + V_f(\hat{f}(x, -Kx)) - V_f(A_K x) \\ &\leq -[1 - b|x|^2 - a|x|^4]|x|^2 \end{aligned}$$

for all $x \in \mathbb{R}^2$. The polynomial inside the brackets has roots at $x_* = -1.0231\dots$ and $x^* = 0.9774\dots$ and is positive in between. Recall $c_f := \underline{\sigma}(P_f)/8$. Then $\underline{\sigma}(P_f)|x|^2 \leq V_f(x) \leq c_f = \underline{\sigma}(P_f)/8$ implies $|x| \leq \frac{1}{2\sqrt{2}} < x^*$ and $|u| = |Kx| = 2(|x_1| + |x_2|) \leq 2\sqrt{2}|x| \leq 1$, so

Assumption 7.3 is satisfied with $\kappa_f(x) := -Kx = -2x_1 - 2x_2$, and P_f and \mathbb{X}_f as defined.

Chapter 8

Stability of offset-free MPC despite plant-model mismatch

Despite over twenty years of applied use and active research, there are no general results on the stability of offset-free MPC with respect to tracking errors. Sufficient conditions for which linear offset-free MPC stability implies offset-free performance were first established by Muske and Badgwell (2002); Pannocchia and Rawlings (2003). While Muske and Badgwell (2002); Pannocchia and Rawlings (2003) do not explicitly mention control of nonlinear plants, the results are widely applicable to both linear and nonlinear plants with asymptotically constant disturbances, as controller stability is assumed rather than explicitly demonstrated. In fact, Pannocchia and Rawlings (2003) demonstrate offset-free control on a highly nonlinear, non-isothermal reactor model.

Offset-free MPC designs with *nonlinear models* and tracking costs were first considered by Morari and Maeder (2012). For the special case of state feedback, Pannocchia et al. (2015) give a disturbance model and estimator design for which the offset-free MPC is provably asymptotically stable and offset-free. In Pannocchia et al. (2015), the state-feedback observer design is generalized to economic cost functions, and convergence to the optimal steady state is demonstrated. A general, output-feedback offset-free *economic* MPC was first proposed by Vaccari and Pannocchia (2017), who use gradient correction strategies to ensure the eco-

conomic MPC, if it converges, achieves the optimal steady-state performance. For further developments of offset-free economic MPC, we refer the reader to Pannocchia (2018); Faulwasser and Pannocchia (2019); Vaccari et al. (2021).

The results discussed thus far have assumed closed-loop stability rather than proven it. Some authors have proposed provably stable and output-tracking nonlinear MPC designs (Falugi, 2015; Limon et al., 2018; Köhler et al., 2020; Berberich et al., 2022b; Galuppini et al., 2023; Soloperto et al., 2023), but they all assume access to the plant dynamic equations, and none consider process and measurement disturbances.

In this chapter, we propose a nonlinear offset-free MPC design that has offset-free performance and asymptotic stability subject to plant-model mismatch, persistent disturbances, and changing references. Based on the results in Chapter 7, we use positive definite quadratic costs and assume differentiability of the plant and model equations to ensure the plant-model mismatch does not prevent stability with respect to the steady-state targets. To ensure the controller is robustly feasible, we soften any output constraints in the regulator using an exact penalty method, and to guarantee nominal regulator stability, we apply constraint backoffs to the steady-state target problem. Lipschitz continuity of the steady-state target problem solutions is required to guarantee robustness to estimate errors and setpoint and disturbance changes.

We outline the chapter as follows. In Section 8.1, the offset-free MPC design is presented. In Section 8.3, we establish asymptotic stability of the nominal system. In Section 8.4, we establish robust performance with respect to estimate errors, setpoint changes, and disturbance changes. In Section 8.5, we extend these results to the mismatched system using the approach from Chapter 7. Finally, in Section 8.7, we conclude the chapter with a discussion of limitations and future work.

8.1 Problem statement

8.1.1 System of interest

Consider the following discrete-time plant:

$$x_{\mathbb{P}}^+ = f_{\mathbb{P}}(x_{\mathbb{P}}, u, w_{\mathbb{P}}) \quad (8.1a)$$

$$y = h_{\mathbb{P}}(x_{\mathbb{P}}, u, w_{\mathbb{P}}) \quad (8.1b)$$

where $x_{\mathbb{P}} \in \mathbb{X} \subseteq \mathbb{R}^n$ is the *plant* state, $u \in \mathbb{U} \subseteq \mathbb{R}^{n_u}$ is the input, $y \in \mathbb{Y} \subseteq \mathbb{R}^{n_y}$ is the output, and $w_{\mathbb{P}} \in \mathbb{W} \subseteq \mathbb{R}^{n_w}$ is the *plant* disturbance. The functions $f_{\mathbb{P}}$ and $h_{\mathbb{P}}$ are not known. Instead, we assume access to a model of the plant,

$$x^+ = f(x, u, d) \quad (8.2a)$$

$$y = h(x, u, d) \quad (8.2b)$$

where $x \in \mathbb{X} \subseteq \mathbb{R}^n$ is the *model* state and $d \in \mathbb{D} \subseteq \mathbb{R}^{n_d}$ is the *model* disturbance. Without loss of generality, we assume the nominal plant and model functions are consistent, i.e.,

$$f(x, u, 0) = f_{\mathbb{P}}(x, u, 0), \quad h(x, u, 0) = h_{\mathbb{P}}(x, u, 0) \quad (8.3)$$

for all $(x, u) \in \mathbb{X} \times \mathbb{U}$. The plant disturbance $w_{\mathbb{P}}$ may include process and measurement noise, exogenous disturbances, parameter errors, discretization errors, and even unmodeled dynamics. The purpose of the model disturbance d is to align the plant and model outputs at steady state. The model disturbance d may include any of the plant disturbances and/or fictitious signals accounting for the effect of the plant disturbances on the steady-state output.

Example 8.1. Consider a single-state linear plant with parameter errors,

$$\begin{aligned} f_{\text{P}}(x_{\text{P}}, u, w_{\text{P}}) &= (\hat{a} + (w_{\text{P}})_1)x_{\text{P}} + (\hat{b} + (w_{\text{P}})_2)u \\ h_{\text{P}}(x_{\text{P}}, u, w_{\text{P}}) &= x_{\text{P}} + (w_{\text{P}})_3 \end{aligned}$$

and a single-state linear model with an input disturbance:

$$f(x, u, d) = \hat{a}x + \hat{b}(u + d), \quad h(x, u, d) = x.$$

For this example, the plant disturbance w_{P} includes both parameter errors and measurement noise, whereas the model disturbance only provides the means to shift the model steady states in response to plant disturbances.

The control objective is to drive the reference signal,

$$r = g(u, y) \tag{8.4}$$

to the setpoint r_{sp} using only knowledge of the model (8.2), past (u, y) data, and auxiliary setpoints $(u_{\text{sp}}, y_{\text{sp}})$ (to be defined). The setpoints $s_{\text{sp}} := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}})$ are possibly time-varying, but only the current value is available at a given time. The controller should be *offset-free* when the setpoint and plant disturbances are asymptotically constant, i.e.,

$$(\Delta s_{\text{sp}}(k), \Delta w_{\text{P}}(k)) \rightarrow 0 \quad \Rightarrow \quad r(k) - r_{\text{sp}}(k) \rightarrow 0$$

where $\Delta s_{\text{sp}}(k) := s_{\text{sp}}(k) - s_{\text{sp}}(k - 1)$ and $\Delta w_{\text{P}}(k) := w_{\text{P}}(k) - w_{\text{P}}(k - 1)$. Otherwise, the amount of offset should be robust to setpoint and disturbance *increments* $(\Delta s_{\text{sp}}, \Delta w_{\text{P}})$.

Remark 8.2. To achieve nominal consistency (8.3) and track the reference (8.4), we typically need the dimensional constraints $n_y \leq n_d$ and $n_r \leq n_u$, respectively. Otherwise their are

insufficient degrees of freedom to manipulate the output and reference at steady state with the disturbance and input, respectively.

Remark 8.3. We do not strictly require an asymptotically constant disturbance. For example, if $r_{\text{sp}}(k) = \sin(1/k)$ and $w_{\text{p}} \equiv 0$, then the setpoint increments go to zero $\Delta r_{\text{sp}}(k) = \sin(1/k) - \sin(1/(k-1)) = O(1/k^2)$. But the setpoint signal becomes approximately constant as $k \rightarrow \infty$, so we should expect the offset-free MPC to be approximately offset-free.

Throughout, we make the following assumptions on plant, model, and reference functions.

Assumption 8.4 (Continuity). The functions $g : \mathbb{U} \times \mathbb{Y} \rightarrow \mathbb{R}^{n_r}$, $(f_{\text{p}}, h_{\text{p}}) : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \rightarrow \mathbb{X} \times \mathbb{Y}$, and $(f, h) : \mathbb{X} \times \mathbb{U} \times \mathbb{D} \rightarrow \mathbb{X} \times \mathbb{Y}$ are continuous, and $f(0, 0, 0) = 0$, $h(0, 0, 0) = 0$, $g(0, 0) = 0$, and (8.3) holds for all $(x, u) \in \mathbb{X} \times \mathbb{U}$.

8.1.2 Constraints

The sets $(\mathbb{X}, \mathbb{Y}, \mathbb{D}, \mathbb{W})$ are physical constraints (e.g., nonnegativity of chemical concentrations, temperatures, pressures, etc.) that the systems (8.1), (8.2), and (8.3) automatically satisfy. These are hard constraints enforced only during state estimation. On the other hand, we enforce the hard constraint $u \in \mathbb{U}$ during both regulation and target selection. Additionally, we enforce soft joint input-output constraints of the form

$$\mathbb{Z}_y := \{ (u, y) \in \mathbb{U} \times \mathbb{Y} \mid c_i(u, y) \leq 0 \forall i \in \mathbb{I}_{1:n_c} \}$$

where $c : \mathbb{U} \times \mathbb{Y} \rightarrow \mathbb{R}^{n_c}$ is the soft constraint function. Having active constraints at steady state is problematic, so the constraints are sometimes tightened as follows:

$$\bar{\mathbb{Z}}_y := \{ (u, y) \mid c_i(u, y) + b_i \leq 0 \forall i \in \mathbb{I}_{1:n_c} \}$$

where $b \in \mathbb{R}_{>0}^{n_c}$ is the vector of back-off constants. No such constraint tightening is required for the input constraints. We assume the constraints and the back-off constant satisfy the following properties throughout.

Assumption 8.5 (Constraints). The sets (\mathbb{X}, \mathbb{Y}) are closed, $(\mathbb{U}, \mathbb{W}, \mathbb{D})$ are compact, and all contain the origin. The soft constraint function $c : \mathbb{U} \times \mathbb{Y} \rightarrow \mathbb{R}^{n_c}$ is continuous and

$$0 < b_i < -c_i(0, 0), \quad \forall i \in \mathbb{I}_{1:n_c}.$$

8.1.3 Offset-free model predictive control

Offset-free MPC consists of three parts or subroutines: target selection, regulation, and state estimation.

Steady-state target problem

Given a *model* disturbance $d \in \mathbb{D}$ and setpoint $r_{\text{sp}} \in \mathbb{R}^{n_r}$, we define the set of offset-free steady-state pairs by

$$\mathcal{Z}_O(r_{\text{sp}}, d) := \{ (x, u) \in \mathbb{X} \times \mathbb{U} \mid x = f(x, u, d), y = h(x, u, d), (u, y) \in \bar{\mathbb{Z}}_y, r_{\text{sp}} = g(u, y) \}. \quad (8.5)$$

To pick the best steady-state pair among members of $\mathcal{Z}_O(r_{\text{sp}}, d)$, it is customary to optimize the steady state with respect to some auxiliary setpoint pair $z_{\text{sp}} := (u_{\text{sp}}, y_{\text{sp}}) \in \bar{\mathbb{Z}}_y$ (typically chosen such that $r_{\text{sp}} = g(u_{\text{sp}}, y_{\text{sp}})$). For each $(r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathbb{R}^{n_r} \times \bar{\mathbb{Z}}_y \times \mathbb{D}$, we define the steady-state target problem (SSTP) by

$$V_s^0(\beta) := \min_{(x, u) \in \mathcal{Z}_O(r_{\text{sp}}, d)} \ell_s(u - u_{\text{sp}}, h(x, u, d) - y_{\text{sp}}) \quad (8.6)$$

where $\beta := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d)$ are the *SSTP parameters* and $\ell_s : \mathbb{R}^{n_u} \times \mathbb{R}^{n_y} \rightarrow \mathbb{R}_{\geq 0}$ is a steady-state cost function, typically a positive definite quadratic. We define the set of feasible SSTP parameters as

$$\mathcal{B} := \{ (r_{\text{sp}}, z_{\text{sp}}, d) \in \mathbb{R}^{n_r} \times \overline{\mathbb{Z}}_y \times \mathbb{D} \mid \mathcal{Z}_O(r_{\text{sp}}, d) \text{ is nonempty} \}. \quad (8.7)$$

To guarantee the existence of solutions to the SSTP (8.6), the following assumption is required.

Assumption 8.6. The function $\ell_s : \mathbb{R}^{n_u} \times \mathbb{R}^{n_y} \rightarrow \mathbb{R}_{\geq 0}$ is continuous and, for each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, at least one of the following properties holds:

- (i) $\mathcal{Z}_O(r_{\text{sp}}, d)$ is compact;
- (ii) with $V_s(x, u, \beta) := \ell_s(u - u_{\text{sp}}, h(x, u, d) - y_{\text{sp}})$, the function $V_s(\cdot, \cdot, \beta)$ is coercive in $\mathcal{Z}_O(r_{\text{sp}}, d)$, i.e., for any sequence $(x_k, u_k) \in \mathcal{Z}_O(r_{\text{sp}}, d)$ such that $\|(x_k, u_k)\| \rightarrow \infty$, we have $V_s(x_k, u_k, \beta) \rightarrow \infty$.

Under Assumptions 8.4 to 8.6, \mathcal{B} is nonempty and the SSTP (8.6) has solutions for all $\beta \in \mathcal{B}$. The solution to (8.6) may not be unique. Throughout, we assume some selection rule has been applied and denote the functions returning solutions to (8.6) by $z_s(\cdot) := (x_s(\cdot), u_s(\cdot)) : \mathcal{B} \rightarrow \mathbb{X} \times \mathbb{U}$.

Regulator

Given the SSTP parameters $\beta \in \mathcal{B}$, the regulator is defined as a finite horizon optimal control problem (FHOCP) with the steady-state targets $(x_s(\beta), u_s(\beta))$. We consider a FHOCP with a horizon length $N \in \mathbb{I}_{>0}$, stage cost $\ell : \mathbb{X} \times \mathbb{U} \times \mathcal{B} \rightarrow \mathbb{R}_{\geq 0}$, terminal cost $V_f : \mathbb{X} \times \mathcal{B} \rightarrow \mathbb{R}_{\geq 0}$, and terminal constraint $\mathbb{X}_f(\beta) \subseteq \mathbb{X}$ (to be defined). For each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, we define the terminal constraint (8.8), feasible initial state and input sequence pairs (8.9), feasible input sequences at $x \in \mathbb{X}$ (8.10), feasible initial states (8.11),

and feasible state-parameter pairs (8.12) by the sets

$$\mathbb{X}_f(\beta) := \text{lev}_{c_f} V_f(\cdot, \beta) \quad (8.8)$$

$$\mathcal{Z}_N(\beta) := \{ (x, \mathbf{u}) \in \mathbb{X} \times \mathbb{U}^N \mid \phi(N; x, \mathbf{u}, d) \in \mathbb{X}_f(\beta) \} \quad (8.9)$$

$$\mathcal{U}_N(x, \beta) := \{ \mathbf{u} \in \mathbb{U}^N \mid (x, \mathbf{u}) \in \mathcal{Z}_N(\beta) \} \quad (8.10)$$

$$\mathcal{X}_N(\beta) := \{ x \in \mathbb{X} \mid \mathcal{U}_N(x, \beta) \text{ is nonmpety} \} \quad (8.11)$$

$$\mathcal{S}_N := \{ (x, \beta) \in \mathbb{X} \times \mathcal{B} \mid \mathcal{U}_N(x, \beta) \text{ is nonmpety} \} \quad (8.12)$$

where $c_f > 0$ and $\phi(k; x, \mathbf{u}, d)$ denotes the solution to (8.2a) at time k given an initial state x , constant disturbance d , and sufficiently long input sequence \mathbf{u} . For each $(x, \mathbf{u}, \beta) \in \mathbb{X} \times \mathbb{U}^N \times \mathcal{B}$, we define the FHOCP objective by

$$V_N(x, \mathbf{u}, \beta) := V_f(\phi(N; x, \mathbf{u}, d), \beta) + \sum_{k=0}^{N-1} \ell(\phi(k; x, \mathbf{u}, d), u(k), \beta). \quad (8.13)$$

For each $(x, \beta) \in \mathcal{S}_N$, we define the FHOCP by

$$V_N^0(x, \beta) := \min_{\mathbf{u} \in \mathcal{U}_N(x, \beta)} V_N(x, \mathbf{u}, \beta). \quad (8.14)$$

As in Chapter 7, we take $V_N^0(x, \beta) := \infty$ for all infeasible pairs $(x, \beta) \notin \mathcal{S}_N$, according to the convention of Rockafellar and Wets (1998).

To guarantee closed-loop stability and robustness, we consider the following assumptions.

Assumption 8.7 (Terminal control law). There exists a function $\kappa_f : \mathbb{X} \times \mathcal{B} \rightarrow \mathbb{U}$ such that

$$V_f(f(x, \kappa_f(x, \beta), d), \beta) - V_f(x, \beta) \leq -\ell(x, \kappa_f(x, \beta), \beta)$$

for all $x \in \mathbb{X}_f(\beta)$ and $\beta := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$.

Assumption 8.8 (Quadratic costs). The stage and terminal costs take the form

$$\ell(x, u, \beta) = |x - x_s(\beta)|_Q^2 + |u - u_s(\beta)|_R^2 + \sum_{i=1}^{n_c} w_i \max \{ 0, c_i(u, h(x, u, d)) \}$$

$$V_f(x, \beta) = |x - x_s(\beta)|_{P_f(\beta)}^2$$

for each $(x, u) \in \mathbb{X} \times \mathbb{U}$ and $\beta := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, where Q , R , and $P_f(\beta)$ are positive definite matrices for each $\beta \in \mathcal{B}$, the function P_f is continuous, and $w_i > 0$ for each $i \in \mathbb{I}_{1:n_c}$.

Remark 8.9. With $\beta = (s_{\text{sp}}, d) \in \mathcal{B}$, Assumption 8.7 and the terminal set definition (8.8) imply $V_f(f(x, \kappa_f(x, \beta), d), \beta) \leq V_f(x, \beta) \leq c_f$ for all $x \in \mathbb{X}_f(\beta)$ and therefore $\mathbb{X}_f(\beta)$ is positive invariant for $x^+ = f(x, \kappa_f(x, \beta), d)$.

Assumptions 8.4 to 8.6 and 8.8 guarantee the existence of solutions to (8.14) for all $(x, \beta) \in \mathcal{S}_N$ (Rawlings et al., 2020, Prop. 2.4). We denote any such solution by $\mathbf{u}^0(x, \beta) = (u^0(0; x, \beta), \dots, u^0(N-1; x, \beta))$, and define the corresponding optimal state sequence by $\mathbf{x}^0(x, \beta) := (x^0(0; x, \beta), \dots, x^0(N; x, \beta))$ where $x^0(k; x, \beta) := \phi(k; x, \mathbf{u}^0(x, \beta), d)$. We define the FHOCP control law by $\kappa_N(x, \beta) := u^0(0; x, \beta)$.

Remark 8.10. Given Assumptions 8.4 to 8.6 and 8.8, it may be impossible to satisfy Assumption 8.7 without constraint back-offs, i.e., $b = 0$. This is because the terminal cost difference $V_f(f(x, \kappa_f(x, \beta), d)) - V_f(x)$ is, at best, negative definite with quadratic scaling (regardless of the target value), whereas the stage cost $\ell(x, \kappa_f(x, \beta), \beta)$ has quadratic scaling when the soft constraint is satisfied but linear scaling when the soft constraint is violated. Thus, if the constraints are active at the targets, the stage cost will always exceed the decrease in terminal cost if the state violates the constraints and is sufficiently small.

Example 8.11. Consider the scalar linear system $x^+ = x + u + d$, $y = x$, and $r = y$ with stage costs of the form Assumption 8.8 and the soft constraint function $c(u, y) = y - 1$. Let

$b = 0$ and $\beta = (1, 0, 1, 0)$. Clearly the target is reachable, and we can take the SSTP (8.6) solution $(x_s(\beta), u_s(\beta)) = (1, 0)$. Then we have stage costs of the form $\ell(x, u, \beta) = q(x - 1)^2 + ru^2 + w \max\{0, x - 1\}$ and $V_f(x, \beta) = p_f x^2$, where $q, r, w, p_f > 0$. Assumption 8.7 is not satisfied if there exists $x \in \mathbb{R}$ such that

$$\mathcal{F}(x, u) := p_f(x + u - 1)^2 - p_f(x - 1)^2 + q(x - 1)^2 + ru^2 + w \max\{0, x - 1\} > 0$$

for all $u \in \mathbb{R}$. Completing the squares gives

$$\begin{aligned} \mathcal{F}(x, u) &= (\tilde{a}u + \tilde{b}(x - 1))^2 + \tilde{c}(x - 1)^2 + w \max\{0, x - 1\} \\ &\geq \tilde{c}(x - 1)^2 + w \max\{0, x - 1\} \end{aligned}$$

for all $x \in \mathbb{R}$ and $u \in \mathbb{R}$, where $\tilde{a} := \sqrt{r + p_f}$, $\tilde{b} := \frac{p_f}{2\tilde{a}}$, and $\tilde{c} := q - \tilde{b}^2$. Ideally, we would have chosen (q, r, p_f) so that $\tilde{c} < 0$. But this means we can still take $0 < x - 1 < \sqrt{\frac{w}{\tilde{c}}}$ to give

$$\mathcal{F}(x, u) \geq \tilde{c}(x - 1)^2 + w(x - 1) > 0$$

for all $u \in \mathbb{R}$, no matter the chosen $w > 0$.

On the other hand, let $b = 1$ and $\beta = (0, 0, 0, 0)$. Again, the target is reachable and we can take the SSTP solution $(x_s(0), u_s(0)) = (0, 0)$. Notice that for both problems the backed-off constraint $c(u, y) + b$ is active at the solution. This time, however, we have

$$\begin{aligned} \mathcal{F}(x, u) &:= p_f(x + u)^2 - p_f x^2 + qx^2 + ru^2 + w \max\{0, x - 1\} \\ &= (\tilde{a}u + \tilde{b}x)^2 + \tilde{c}x^2 + w \max\{0, x - 1\} \end{aligned}$$

and with $\kappa_f(x, 0) := -\frac{\tilde{b}}{a}x$, we have

$$\mathcal{F}(x, \kappa_f(x, 0)) = \tilde{c}x^2 + w \max\{0, x - 1\}$$

for all $x \in \mathbb{R}$. Let $c_f = p_f$ and suppose $\tilde{c} < 0$. Then, for each $x \in \mathbb{X}_f(0)$, we have $|x| \leq 1$ and therefore

$$\mathcal{F}(x, \kappa_f(x, 0)) = \tilde{c}x^2 \leq 0.$$

State estimation

In practice, the SSTP and FHOCP are implemented with state and disturbance estimates rather than the true values. To this end, we consider any estimator that estimates both plant and disturbance states.

Definition 8.12. A joint state and disturbance estimator is a sequence of functions $\Phi_k : \mathbb{X} \times \mathbb{D} \times \mathbb{U}^k \times \mathbb{Y}^k \rightarrow \mathbb{X} \times \mathbb{D}$ defined for each $k \in \mathbb{I}_{\geq 0}$. For each $k \in \mathbb{I}_{\geq 0}$, we define the *state and disturbance estimates* by

$$(\hat{x}(k), \hat{d}(k)) := \Phi_k(\bar{x}, \bar{d}, \mathbf{u}_{0:k-1}, \mathbf{y}_{0:k-1}) \quad (8.15)$$

where $(\bar{x}, \bar{d}) \in \mathbb{X} \times \mathbb{D}$ is the initial guess at time $k = 0$, $\mathbf{u} \in \mathbb{U}^\infty$ is the input data, and $\mathbf{y} \in \mathbb{Y}^\infty$ is the output data.

Remark 8.13. Since the regulator requires a state estimate to compute, and the input directly affects the output, the current state and disturbance estimates $(\hat{x}(k), \hat{d}(k))$ must be functions of past data, not including the current measurement $y(k)$. Therefore, at time $k = 0$, there is no data available to update the prior guess, and most estimator designs will take Φ_0 as the identity map, i.e.,

$$(\hat{x}(0), \hat{d}(0)) := \Phi_0(\bar{x}, \bar{d}) = (\bar{x}, \bar{d}).$$

However, we can also consider models without direct feedthrough effects (i.e., $y = h(x, d)$) in which case Definition 8.12 can be modified so the estimator functions also take $y(k)$ as an argument.

The estimator (8.15) is designed according to the model (8.2) and thus has no knowledge of the plant state x_p or plant disturbance w_p . To analyze its performance and state the assumptions needed to establish offset-free performance, we consider the following noise model:

$$x^+ = f(x, u, d) + w \quad (8.16a)$$

$$d^+ = d + w_d \quad (8.16b)$$

$$y = h(x, u, d) + v \quad (8.16c)$$

where $\tilde{w} := (w, w_d, v) \in \tilde{\mathbb{W}}(x, u, d) \subseteq \mathbb{R}^{n_{\tilde{w}}}$ are the process, disturbance, and measurement noises, $n_{\tilde{w}} := n + n_d + n_y$, and

$$\tilde{\mathbb{W}}(x, u, d) := \{ (w, w_d, v) \mid (x^+, d^+, y) \in \mathbb{X} \times \mathbb{D} \times \mathbb{Y}, (8.16) \}$$

is a constraint set that ensures all quantities remain physical. We define the set of feasible trajectories by

$$\tilde{\mathbb{Z}}_e := \{ (\mathbf{x}, \mathbf{u}, \mathbf{d}, \mathbf{y}, \tilde{\mathbf{w}}) \in \mathbb{X}^\infty \times \mathbb{U}^\infty \times \mathbb{D}^\infty \times \mathbb{Y}^\infty \times (\mathbb{R}^{n_{\tilde{w}}})^\infty \mid (8.16) \text{ and } \tilde{w} = (w, w_d, v) \in \tilde{\mathbb{W}}(x, u, d) \}.$$

Finally, denoting the state, disturbance, and errors by

$$e_x(k) := x(k) - \hat{x}(k), \quad e_d(k) := d(k) - \hat{d}(k), \quad (8.17a)$$

$$e(k) := \begin{bmatrix} e_x(k) \\ e_d(k) \end{bmatrix}, \quad \bar{e} := \begin{bmatrix} x(0) - \bar{x} \\ d(0) - \bar{d} \end{bmatrix}, \quad (8.17b)$$

we define robust stability of the estimator (8.15) as follows.

Definition 8.14. The estimator (8.15) is *robustly globally exponentially stable* (RGES) for the system (8.16) if there exist constants $c_{e,1}, c_{e,2} > 0$ and $\lambda_e \in (0, 1)$ such that

$$|e(k)| \leq c_{e,1} \lambda_e^k |\bar{e}| + c_{e,2} \sum_{j=1}^k \lambda_e^{j-1} |\tilde{w}(k-j)|$$

for each $k \in \mathbb{I}_{\geq 0}$, prior guess $(\bar{x}, \bar{d}) \in \mathbb{X} \times \mathbb{D}$, and trajectories $(\mathbf{x}, \mathbf{u}, \mathbf{d}, \mathbf{y}, \tilde{\mathbf{w}}) \in \tilde{\mathbb{Z}}_e$, given definitions (8.15) and (8.17).

For the case with plant-model mismatch, the estimator (8.15) is not only assumed to be RGES for the system (8.16), but is also assumed to admit a robust global Lyapunov function.

Assumption 8.15. The initial estimator Φ_0 is the identity map. There exists a function $V_e : \mathbb{X} \times \mathbb{D} \times \mathbb{X} \times \mathbb{D} \rightarrow \mathbb{R}_{\geq 0}$ and constants $c_1, c_2, c_3, c_4, \delta_w > 0$ such that

$$c_1 |e(k)|^2 \leq V_e(k) \leq c_2 |e(k)|^2 \quad (8.18a)$$

$$V_e(k+1) \leq V_e(k) - c_3 |e(k)|^2 + c_4 |\tilde{w}(k)|^2 \quad (8.18b)$$

for all $(\bar{x}, \bar{d}) \in \mathbb{X} \times \mathbb{D}$, $(\mathbf{x}, \mathbf{u}, \mathbf{d}, \mathbf{y}, \tilde{\mathbf{w}}) \in \tilde{\mathbb{Z}}_e$, and $k \in \mathbb{I}_{\geq 0}$, where (8.15), (8.17), and $V_e(k) := V_e(x(k), d(k), \hat{x}(k), \hat{d}(k))$.

The following theorem establishes that Assumption 8.15 implies RGES of the estimator (8.15) for the system (8.16) (see Appendix 8.A.1 for proof).

Theorem 8.16. *Suppose the estimator (8.15) for the system (8.16) satisfies Assumption 8.15. Then the estimator is RGES under Definition 8.14.*

Remark 8.17. In Assumption 8.15, we assume Φ_0 is the identity map, and therefore $e(0) = \bar{e}$. However, as mentioned in Remark 8.13, if we consider models without direct input-output effects (i.e., $y = \hat{h}(x, d)$), then the estimator functions Φ_k may become a function of the current output $y(k)$ and it is no longer reasonable to assume Φ_0 is the identity map. Then $e(0) \neq \bar{e}$ in general. However, we can modify Definition 8.12 to include robustness to the current noise $\tilde{n}(k)$, and we can modify Assumption 8.15 to include a linear bound of the form $|e(0)| \leq \bar{a}_1|\bar{e}| + \bar{a}_2|\tilde{w}(0)|$, for some $\bar{a}_1, \bar{a}_2 > 0$, to again imply RGES of the estimator.

While Assumption 8.15 is satisfied for stable full-order observers of (8.16),¹ we know of no nonlinear results that guarantee a Lyapunov function characterization of stability (i.e., Assumption 8.15) for the full information estimation (FIE) or moving horizon estimation (MHE) algorithms. FIE and MHE were shown to be RGES for exponentially detectable and stabilizable systems by Allan and Rawlings (2021), but they use a Q -function to demonstrate stability. To the best of our knowledge, the closest construction is the N -step Lyapunov function of Schiller et al. (2023). If we treat the disturbance as a parameter, rather than an uncontrollable integrator, there are FIE and MHE algorithms for combined state and parameter estimation that could also be used to estimate the states and disturbances (Muntwiler et al., 2023; Schiller and Müller, 2023).²

¹A full-order state observer of (8.16) is a dynamical system, evolving in the same state space as (8.16), stabilized with respect to x by output feedback.

²The estimation algorithms of Muntwiler et al. (2023) produce RGES state estimates, but it is not shown the parameter estimates are RGES. The estimation algorithm of Schiller and Müller (2023) produces RGES state and parameter estimates, but only under a persistence of excitation condition.

8.2 Robust stability for tracking and estimation

In this section, we consider stabilization of the system,³

$$\xi^+ = F(\xi, u, \omega), \quad \omega \in \Omega(\xi, u). \quad (8.19)$$

The system (8.19) represents the evolution of an *extended* plant state $\xi \in \Xi \subseteq \mathbb{R}^{n_\xi}$ subject to the input $u \in \mathbb{U}$ and *extended* disturbance $\omega \in \Omega(\xi, u) \subseteq \mathbb{R}^{n_\omega}$ (to be defined). Greek letters are used for the extended state and disturbance (ξ, ω) to avoid confusion with the states and disturbances of (8.1), (8.2), and (8.16). Throughout, we assume Ξ is closed and $0 \in \Omega(\xi, u)$ and $F(\xi, u, \omega) \in \Xi$ for all $(\xi, u) \in \Xi \times \mathbb{U}$ and $\omega \in \Omega(\xi, u)$.

8.2.1 Robust stability with respect to two outputs

We first consider stabilization of (8.19) under state feedback,

$$\xi^+ = F_c(\xi, \omega), \quad \omega \in \Omega_c(\xi) \quad (8.20)$$

where $\kappa : \Xi \rightarrow \mathbb{U}$ is the control law, $F_c(\xi, \omega) := F(\xi, \kappa(\xi), \omega)$, and $\Omega_c(\xi) := \Omega(\xi, \kappa(\xi))$. We define robust positive invariance for the system (8.20) as follows.

Definition 8.18 (Robust positive invariance). A closed set $X \subseteq \Xi$ is *robustly positive invariant* (RPI) for the system (8.20) if $\xi \in X$ and $\omega \in \Omega_c(\xi)$ imply $F_c(\xi, \omega) \in X$.

Robust target- and setpoint-tracking stability are defined under the umbrella of input-to-state stability (ISS) with respect to two measurement functions (Tran et al., 2015). We slightly modify their definition by considering measurement functions of (ξ, ω) (rather than just ξ)

³To ensure unphysical states are not produced by additive disturbances, we let the disturbance set be a function of the state and input. However, we can convert (8.19) to a standard form by taking $\xi^+ = \tilde{F}(\xi, u, \omega)$, $\omega \in \Omega$ where $\tilde{F}(\xi, u, \omega) = F(\xi, \text{proj}_{\Omega(\xi, u)}(\omega))$, $\Omega := \bigcup_{(\xi, u) \in \Xi \times \mathbb{U}} \Omega(\xi, u)$, and $\text{proj}_{\Omega(\xi, u)}(\omega) = \text{argmin}_{\omega' \in \Omega(\xi, u)} |\omega - \omega'|$.

and structuring the measurement functions as norms of the outputs $\zeta_1 \in \mathbb{R}^{n_{\zeta_1}}$ and $\zeta_2 \in \mathbb{R}^{n_{\zeta_2}}$, where

$$\zeta_1 = G_1(\xi, \omega), \quad \zeta_2 = G_2(\xi, \omega). \quad (8.21)$$

The definition of Tran et al. (2015) can be reconstructed by taking G_1 and G_2 as scalar-valued, positive semidefinite functions of ξ .

Definition 8.19 (Robust stability w.r.t. two outputs). We say the system (8.20) (with outputs (8.21)) is *robustly asymptotically stable* (RAS) (on a RPI set $X \subseteq \Xi$) with respect to (ζ_1, ζ_2) if there exist $\beta_\zeta \in \mathcal{KL}$ and $\gamma_\zeta \in \mathcal{K}$ such that

$$|\zeta_1(k)| \leq \beta_\zeta(|\zeta_2(0)|, k) + \gamma_\zeta(\|\omega\|_{0:k}) \quad (8.22)$$

for each $k \in \mathbb{I}_{\geq 0}$ and trajectories $(\xi, \omega, \zeta_1, \zeta_2)$ satisfying (8.20), (8.21), and $\xi(0) \in X$. We say (8.20) is *robustly exponentially stable* (RES) w.r.t. (ζ_1, ζ_2) if it is RAS w.r.t. (ζ_1, ζ_2) with $\beta_\zeta(s, k) := c_\zeta \lambda_\zeta^k s$ for some $c_\zeta > 0$ and $\lambda_\zeta \in (0, 1)$.

For the nominal case (i.e., $\Omega(\xi, u) \equiv \{0\}$), we drop the word *robust* from Definitions 8.18 and 8.19 and simply write *positive invariant*, *asymptotically stable* (AS), and *exponentially stable* (ES). Moreover, if (8.20) is RAS (RES) w.r.t. (ζ, ζ) , where $\zeta = G(\xi, \omega)$, we simply say it is RAS (RES) w.r.t. ζ .

In Sections 8.3 and 8.4, we use Definition 8.19 demonstrate nominal stability and robustness to estimate error, noise, and SSTP parameter changes. The following cases of the system (8.19), control law $u = \kappa(\xi)$, and outputs (8.21) are considered.

1. *Nominal stability*: Let $\xi := x$, $u = \kappa(\xi) := \kappa_N(x, \beta)$, $\omega := 0$, $\zeta_1 := g(u, h(x, u, d)) - r_{\text{sp}}$, and $\zeta_2 := x - x_s(\beta)$. Then, for each *fixed* $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, the closed-loop

system has dynamics (8.20) and outputs (8.21) with

$$\begin{aligned} F(\xi, \omega) &:= f(x, \kappa_N(x, \beta), \beta) \\ G_1(\xi) &:= g(x, h(x, \kappa_N(x, \beta), d)) - r_{\text{sp}} \\ G_2(\xi) &:= x - x_s(\beta) \end{aligned}$$

for each $\xi \in \mathcal{X}_N^\rho := \text{lev}_\rho V_N^0$ and $\omega = 0$. AS (ES) w.r.t. ζ_2 corresponds to (exponential) target-tracking stability, and AS (ES) w.r.t. (ζ_1, ζ_2) corresponds to (exponential) setpoint-tracking stability.

2. *Robust stability (w.r.t. estimate error, noise, SSTP parameter changes):* Let $\xi := (\hat{x}, \hat{\beta})$, $\kappa(\xi) := \kappa_N(\xi)$, $\omega := (e, e^+, \Delta s_{\text{sp}}, \tilde{w})$, $\zeta_1 := r - r_{\text{sp}}$, $\zeta_2 := \hat{x} - x_s(\hat{\beta})$, where $r := g(u, h(\hat{x} + e_x, u, \hat{d} + e_d) + v)$ and $\hat{\beta} := (s_{\text{sp}}, \hat{d})$. Then the closed-loop system has dynamics (8.20) and outputs (8.21) with

$$\begin{aligned} F(\xi, \omega) &:= \begin{bmatrix} f(\hat{x} + e_x, \kappa_N(\hat{x}, \hat{\beta}), \hat{d} + e_d) + w - e_x^+ \\ s_{\text{sp}} + \Delta s_{\text{sp}} \\ \hat{d} + e_d + w_d - e_d^+ \end{bmatrix} \\ G_1(\xi) &:= g(x, h(\hat{x} + e_x, \kappa_N(\hat{x}, \hat{\beta}), \hat{d} + e_d) + v) - r_{\text{sp}}, \\ G_2(\xi) &:= \hat{x} - x_s(\hat{\beta}) \end{aligned}$$

for each $\xi = (\hat{x}, \hat{\beta})$ in a to-be-defined RPI set $\hat{\mathcal{S}}_N^\rho$ and $\omega \in \Omega_c(\xi)$ (to be defined). RAS (RES) of (8.20) w.r.t. ζ_2 alone corresponds to robust (exponential) target-tracking stability, and RAS (RES) w.r.t. (ζ_1, ζ_2) corresponds to robust (exponential) setpoint-tracking stability.

Remark 8.20. If (8.20) is RAS on $X \subseteq \Xi$ w.r.t. (ζ_1, ζ_2) , then $\omega(k) \rightarrow 0$ implies $\zeta_1(k) \rightarrow 0$ so long as $\xi(0) \in X$.

Remark 8.21. Definition 8.19 generalizes many ISS and input-to-output stability (IOS) defi-

nitions originally posed for continuous-time systems by Sontag and Wang (1995, 1999, 2000). However, only Definition 8.19 is suitable for analyzing both target- and setpoint-tracking performance of the offset-free MPC. ISS is not appropriate as the SSTP parameters β are often part of the extended state ξ . IOS and robust output stability allow the tracking performance to degrade with the magnitude of the SSTP parameters. While state-independent IOS (SI IOS) coincides with the special case of $\zeta = G_1(\xi) \equiv G_2(\xi)$ (e.g., for target-tracking), we find the setpoint-tracking error is more tightly bounded by the initial target-tracking error.

Next, we define an (exponential) ISS Lyapunov function with respect to the noise-free outputs

$$\zeta_1 = G_1(\xi), \quad \zeta_2 = G_2(\xi) \quad (8.23)$$

and show its existence implies RAS (RES) of (8.20) with respect to (ζ_1, ζ_2) (see Appendix 8.A.2 for proof).

Definition 8.22. Consider the system (8.20) with outputs (8.23). We call $V : \Xi \rightarrow \mathbb{R}_{\geq 0}$ an ISS Lyapunov function (on a RPI set $X \subseteq \Xi$) with respect to (ζ_1, ζ_2) if there exist $\alpha_i \in \mathcal{K}_{\infty}$, $i \in \mathbb{I}_{1:3}$ and $\sigma \in \mathcal{K}$ such that, for each $\xi \in X$ and $\omega \in \Omega_c(\xi)$,

$$\alpha_1(|G_1(\xi)|) \leq V(\xi) \leq \alpha_2(|G_2(\xi)|) \quad (8.24a)$$

$$V(F_c(\xi, \omega)) \leq V(\xi) - \alpha_3(V(\xi)) + \sigma(|\omega|). \quad (8.24b)$$

We say V is an *exponential ISS Lyapunov function* with respect to (ζ_1, ζ_2) if it is an ISS Lyapunov function with respect to (ζ_1, ζ_2) with $\alpha_i(\cdot) = a_i(\cdot)^b$ for some $a_i, b > 0$, $i \in \mathbb{I}_{1:3}$.

Theorem 8.23. If the system (8.20) with outputs (8.23) admits an (exponential) ISS Lyapunov function $V : \Xi \rightarrow \mathbb{R}_{\geq 0}$ on an RPI set $X \subseteq \Xi$ with respect to (ζ_1, ζ_2) , then it is RAS (RES) on X with respect to (ζ_1, ζ_2) .

Similarly to Definitions 8.18 and 8.19, we call V a *Lyapunov function* or *exponential Lyapunov function* w.r.t. (ζ_1, ζ_2) if it satisfies Definition 8.22 in the nominal case (i.e., $\Omega(\xi, u) \equiv \{0\}$). Moreover, we note that the proof of Theorem 8.23 trivially extends to the nominal case by setting $\omega = 0$ throughout.

Remark 8.24. If $\zeta = G_1(\xi) \equiv G_2(\xi)$, then it suffices to replace (8.24b) with $V(F_c(\xi, \omega)) \leq V(\xi) - \tilde{\alpha}_3(|G_1(\xi)|) + \sigma(|\omega|)$ to establish ISS with respect to ζ , where $\tilde{\alpha}_3 \in \mathcal{K}_\infty$. Then (8.24b) holds with $\alpha_3 := \tilde{\alpha}_3 \circ \alpha_2^{-1}$.

8.2.2 Combined controller-estimator robust stability

In applications without plant-model mismatch, it suffices to consider RES of each of the controller and estimator subsystems to establish RES of the combined system. This is because the controller and estimator error systems are connected *sequentially*, with the target- and setpoint-tracking errors having no influence on the estimation errors. However, as we show in Section 8.5, plant-model mismatch makes this a *feedback interconnection*, with the tracking errors influencing the state estimate errors and vice versa. Therefore it is necessary to analyze stability of the combined system.

We define the *extended* sensor output $v \in \Upsilon \subseteq \mathbb{R}^{n_v}$ by

$$v = H(\xi, u, \omega). \quad (8.25)$$

Assume Υ is closed and $H(\xi, u, \omega) \in \Upsilon$ for all $(\xi, u) \in \Xi \times \mathbb{U}$ and $\omega \in \Omega(\xi, u)$. We consider the extended state estimator

$$\hat{\xi}(k) := \Phi_k^\xi(\bar{\xi}, \mathbf{u}_{0:k-1}, \mathbf{v}_{0:k-1}) \quad (8.26)$$

where $\bar{\xi} \in \hat{\Xi} \subseteq \mathbb{R}^{n_\xi}$ is the prior guess and $\Phi_k^\xi : \hat{\Xi} \times \mathbb{U}^k \times \Upsilon^k \rightarrow \hat{\Xi}, k \in \mathbb{I}_{\geq 0}$. The set $\hat{\Xi}$ is

closed but is not necessarily the same, let alone of the same dimension, as Ξ . We consider stabilization via state estimate feedback,

$$u = \hat{\kappa}(\hat{\xi}) \quad (8.27)$$

where $\hat{\kappa} : \hat{\Xi} \rightarrow \mathbb{U}$. Finally, we define a RPI set as follows.

Definition 8.25. A closed set $S \subseteq \Xi \times \hat{\Xi}$ is RPI for the system (8.19) and (8.25)–(8.27) if $(\xi(k), \hat{\xi}(k)) \in S$ for all $k \in \mathbb{I}_{\geq 0}$ and $(\boldsymbol{\xi}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{v})$ satisfying (8.19), (8.25), (8.26), and (8.27), and $(\xi(0), \bar{\xi}) \in S$.

With plant-model mismatch, the *extended* plant and model states to evolve on different spaces. Thus, we define the estimator error $\varepsilon \in \mathbb{R}^{n_\xi}$ as the deviation of the estimate $\hat{\xi}$ from an arbitrary function $G_\varepsilon : \Xi \rightarrow \hat{\Xi}$ of the state ξ ,

$$\varepsilon(k) = G_\varepsilon(\xi(k)) - \hat{\xi}(k), \quad \bar{\varepsilon} := G_\varepsilon(\xi(0)) - \bar{\xi}. \quad (8.28)$$

Finally, we define robust stability with respect to the outputs

$$\zeta_1 = G_1(\xi, \hat{\xi}, u, \omega), \quad \zeta_2 = G_2(\xi, \hat{\xi}, u, \omega) \quad (8.29)$$

similarly to Definition 8.19.

Definition 8.26. The system (8.19) and (8.25)–(8.27) (with outputs (8.29)) is RAS in a RPI set $S \subseteq \mathbb{X} \times \hat{\mathbb{X}}$ with respect to (ζ_1, ζ_2) if there exist functions $\beta_\zeta, \gamma_\zeta \in \mathcal{KL}$ such that

$$|(\zeta_1(k), \varepsilon(k))| \leq \beta_\zeta(|(\zeta_2(0), \bar{\varepsilon})|, k) + \sum_{i=0}^k \gamma_\zeta(|\omega(k-i)|, i) \quad (8.30)$$

for all $k \in \mathbb{I}_{\geq 0}$ and all trajectories $(\boldsymbol{\xi}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{v}, \boldsymbol{\varepsilon}, \boldsymbol{\zeta}_1, \boldsymbol{\zeta}_2)$ satisfying (8.19), (8.25)–(8.28) and

(8.29), and $(\xi(0), \bar{\xi}) \in \mathcal{S}$. We say (8.19) and (8.25)–(8.27) is RES w.r.t. (ζ_1, ζ_2) if it is RAS w.r.t. (ζ_1, ζ_2) with $\beta_\zeta(s, k) := c_\zeta \lambda_\zeta^k s$ and $\gamma_\zeta(s, k) := \lambda_\zeta^k \sigma_\zeta(s)$ for some $c_\zeta > 0$, $\lambda_\zeta \in (0, 1)$, and $\sigma_\zeta \in \mathcal{K}$.

As in Section 8.2.1, we say (8.19) and (8.25)–(8.27) is RAS (RES) w.r.t. $\zeta = G(\xi, \omega)$ if it is RAS (RES) w.r.t. (ζ, ζ) .

In Section 8.5, we establish robustness of offset-free MPC with plant-model mismatch in terms of Definition 8.26, using the following definition of the system (8.19) and (8.25)–(8.27), estimate errors (8.28), and outputs (8.29):

3. *With mismatch:* Let $\xi := (x_P, \alpha)$, $\hat{\xi} := (\hat{x}, \hat{\beta})$, $u := \kappa_N(\hat{\xi})$, $\omega := (\Delta s_{\text{sp}}, \Delta w_P)$, $v := (y, \Delta s_{\text{sp}})$, $\varepsilon := (x_P + \Delta x_s(\alpha), s_{\text{sp}}, d_s(\alpha)) - \hat{\xi}$, $\zeta_1 := r - r_{\text{sp}}$, $\zeta_2 := \hat{x} - x_s(\hat{\beta})$, where $r := g(u, h_P(x, u, w_P))$, $\alpha := (s_{\text{sp}}, w_P)$, $\hat{\beta} := (s_{\text{sp}}, \hat{d})$, and $(\Delta x_s(\alpha), d_s(\alpha))$ are to be defined. Then the closed-loop system has dynamics (8.19) and (8.25)–(8.27), errors (8.28), and outputs (8.29) with

$$F(\xi, u, \omega) := \begin{bmatrix} f_P(x_P, u, w_P) \\ s_{\text{sp}} + \Delta s_{\text{sp}} \\ w_P + \Delta w_P \end{bmatrix}, \quad H(\xi, u, \omega) := \begin{bmatrix} h_P(\xi, u, w_P) \\ \Delta s_{\text{sp}} \end{bmatrix},$$

$$\Phi_k^\xi(\bar{\xi}, \mathbf{u}_{0:k-1}, \mathbf{v}_{0:k-1}) := (\hat{x}(k), s_{\text{sp}}(k), \hat{d}(k)), \quad G_\varepsilon(\xi) := \begin{bmatrix} x_P + \Delta x_s(\alpha) \\ d_s(\alpha) \end{bmatrix}$$

$$G_1(\xi, u, \omega) := g(u, h_P(x_P, u, w_P)) - r_{\text{sp}}, \quad G_2(\hat{\xi}) := \hat{x} - x_s(\hat{\beta})$$

for each $(\xi, \hat{\xi}) = (x, \beta, \hat{x}, \hat{\beta})$ in a to-be-defined RPI set $\mathcal{S}_N^{\rho, \tau}$ and $\omega \in \Omega_c(\xi)$ (to be defined), where $(\hat{x}(k), \hat{d}(k)) := \Phi_k(\bar{x}, \bar{d}, \mathbf{u}_{0:k-1}, \mathbf{y}_{0:k-1})$ as in Definition 8.12.

As in Section 8.2.1, RAS (RES) w.r.t. ζ_2 corresponds to robust (exponential) target-tracking stability, and RAS (ES) w.r.t. (ζ_1, ζ_2) corresponds to robust (exponential) setpoint-tracking stability.

Remark 8.27. If (8.19) and (8.25)–(8.27) is RAS on a RPI set $\mathcal{S} \subseteq \Xi \times \hat{\Xi}$ w.r.t. (ζ_1, ζ_2) , then

$\omega(k) \rightarrow 0$ implies $(\zeta_1(k), \varepsilon(k)) \rightarrow 0$ so long as $(\xi(0), \bar{\xi}) \in \mathcal{S}$ (cf. (Allan and Rawlings, 2021, Prop. 3.11)).

To analyze stability of the system (8.19), (8.25), (8.26), and (8.27), we use the following theorem (see Appendix 8.A.3 for proof).

Theorem 8.28. *Consider the system (8.19) and (8.25)–(8.27) with errors (8.28) and output $\zeta = G(\hat{\xi})$. Suppose $\Phi_0^{\hat{\xi}}$ is the identity map and there exist constants $a_i, b_i > 0, i \in \mathbb{I}_{1:4}$, a RPI set $\mathcal{S} \subseteq \mathbb{X} \times \hat{\mathbb{X}}$, and functions $V : \hat{\Xi} \rightarrow \mathbb{R}_{\geq 0}$, $V_\varepsilon : \Xi \times \hat{\Xi} \rightarrow \mathbb{R}_{\geq 0}$, and $\sigma, \sigma_\varepsilon \in \mathcal{K}$ such that $\frac{a_4 c_4}{a_3 c_1} < 1$, $\frac{a_4 c_4}{a_3 c_3} < \frac{c_1}{c_1 + c_2}$, and, for all trajectories $(\xi, \hat{\xi}, \mathbf{u}, \omega, \mathbf{v}, \varepsilon, \zeta)$ satisfying (8.19) and (8.25)–(8.28), $\zeta = G(\hat{\xi})$, and $(\xi(0), \bar{\xi}) \in \mathcal{S}$, we also satisfy*

$$a_1 |\zeta|^2 \leq V(\hat{\xi}) \leq a_2 |\zeta|^2 \quad (8.31a)$$

$$V(\hat{\xi}^+) \leq V(\hat{\xi}) - a_3 |\zeta|^2 + a_4 |(\varepsilon, \varepsilon^+)|^2 + \sigma(|\omega|) \quad (8.31b)$$

$$c_1 |\varepsilon|^2 \leq V_\varepsilon(\xi, \hat{\xi}) \leq c_2 |\varepsilon|^2 \quad (8.31c)$$

$$V_\varepsilon(\xi^+, \hat{\xi}^+) \leq V_\varepsilon(\xi, \hat{\xi}) - c_3 |\varepsilon|^2 + c_4 |\zeta|^2 + \sigma_\varepsilon(|\omega|). \quad (8.31d)$$

Then the system (8.19) and (8.25)–(8.27) is RES in \mathcal{S} w.r.t. ζ .

8.3 Nominal offset-free performance

In this section, we consider the application of offset-free MPC to the model (8.2) in the *nominal* case (i.e., without estimate errors or setpoint and disturbance changes). Consider the

following *modeled* closed-loop system:

$$x^+ = f_c(x, \beta) := f(x, \kappa_N(x, \beta), d) \quad (8.32a)$$

$$y = h_c(x, \beta) := h(x, \kappa_N(x, \beta), d) \quad (8.32b)$$

$$r = g_c(x, \beta) := g(\kappa_N(x, \beta), h_c(x, \beta)) \quad (8.32c)$$

where $(x, \beta) := (x, r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{S}_N$. For each $\rho > 0$ and $\beta \in \mathcal{B}$, we define the candidate domain of stability

$$\mathcal{X}_N^\rho(\beta) := \text{lev}_\rho V_N^0(\cdot, \beta). \quad (8.33)$$

In the following theorem, we establish *nominal* stability and offset-free performance of the modeled closed-loop system (8.32), under Assumptions 8.4 to 8.8 and with constant, known setpoints $s_{\text{sp}} = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}})$ and disturbance d .

Theorem 8.29. *Suppose Assumptions 8.4 to 8.8 hold. Let $\rho > 0$.*

(a) *For each compact $\mathcal{B}_c \subseteq \mathcal{B}$, there exist constants $a_1, a_2, a_3 > 0$ such that*

$$a_1|x - x_s(\beta)|^2 \leq V_N^0(x, \beta) \leq a_2|x - x_s(\beta)|^2 \quad (8.34a)$$

$$V_N^0(f_c(x, \beta), \beta) \leq V_N^0(x, \beta) - a_3|x - x_s(\beta)|^2 \quad (8.34b)$$

for all $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta \in \mathcal{B}_c$.

(b) *For each $\beta \in \mathcal{B}$, the system (8.32a) is ES on $\mathcal{X}_N^\rho(\beta)$ with respect to the target-tracking error*

$$\delta x := x - x_s(\beta).$$

(c) *For each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, the system (8.32a) is AS on $\mathcal{X}_N^\rho(\beta)$ with respect to*

$$(\delta r, \delta x), \text{ where } \delta r := g_c(x, \beta) - r_{\text{sp}} \text{ is the setpoint-tracking error.}$$

(d) *If g and h are Lipschitz continuous on bounded sets, then part (c) can be upgraded to ES.*

We include a proof of Theorem 8.29 in Appendix 8.B.1. Two details of the proof are re-

quired for the subsequent results. First, from (Rawlings et al., 2020, Prop. 2.4), we have

$$V_N(f_c(x, \beta), \tilde{\mathbf{u}}(x, \beta), \beta) \leq V_N^0(x, \beta) - \ell(x, \kappa_N(x, \beta), \beta) \quad (8.35)$$

for all $(x, \beta) \in \mathcal{S}_N$, where

$$\tilde{\mathbf{u}}(x, \beta) := (u^0(1; x, \beta), \dots, u^0(N-1; x, \beta), \kappa_f(x^0(N; x, \beta), \beta)) \quad (8.36)$$

is a suboptimal sequence for $x^+ := f_c(x, \beta)$. Second, for each $(x, \beta) \in \mathcal{S}_N$, the suboptimal sequence $\tilde{\mathbf{u}}(x, \beta)$ steers the system from $f_c(x, \beta)$ to the terminal constraint $\mathbb{X}_f(\beta)$ in $N-1$ moves and keeps it there (by Assumption 8.7). Therefore $\tilde{\mathbf{u}}(x, \beta) \in \mathcal{U}_N(f_c(x, \beta), \beta)$ and $f_c(x, \beta) \in \mathcal{X}_N(\beta)$.

Remark 8.30. Theorem 8.29(a) provides Lyapunov bounds that are *uniform* in the SSTP parameters β on compact subsets $\mathcal{B}_c \subseteq \mathcal{B}$. This implies a guaranteed decay rate $\lambda \in (0, 1)$ for the deviation of the state from its target $x - x_s(\beta)$, although this guaranteed rate may become arbitrarily close to 1 as we expand the size of the compact set \mathcal{B}_c .

8.4 Offset-free performance without mismatch

In this section, we show offset-free MPC (without plant-model mismatch) is robust to estimate errors and setpoint and disturbance changes. We assume the actual plant evolves according to the noisy model equations (8.16). We assume the setpoints evolve according to

$$s_{\text{sp}}^+ = s_{\text{sp}} + \Delta s_{\text{sp}} \quad (8.37)$$

where $s_{\text{sp}} := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}})$ and $\Delta s_{\text{sp}} := (\Delta r_{\text{sp}}, \Delta u_{\text{sp}}, \Delta y_{\text{sp}})$. At each time, we define $\beta := (s_{\text{sp}}, d)$ and $\Delta \beta := (\Delta s_{\text{sp}}, w_d)$ and sometimes write $\beta^+ = \beta + \Delta \beta$. Taking the approach

of (Rawlings et al., 2020, Sec. 4.6), the estimate error system evolves as

$$\hat{x}^+ = f(\hat{x} + e_x, u, \hat{d} + e_d) + w - e_x^+ \quad (8.38a)$$

$$\hat{d}^+ = \hat{d} + e_d + w_d - e_d^+ \quad (8.38b)$$

$$y = h(\hat{x} + e_x, u, \hat{d} + e_d) + v. \quad (8.38c)$$

We lump the perturbation terms from (8.37) and (8.38) into a single disturbance variable, defined as $\tilde{d} := (e, e^+, \Delta s_{\text{sp}}, \tilde{w})$. To ensure the noise does not result in unphysical states, disturbances, or measurements, we define the set of admissible perturbations as

$$\begin{aligned} \tilde{\mathbb{D}}(\hat{x}, u, \hat{d}) := \{ \tilde{d} = (e_x, e_d, e_x^+, e_d^+, \Delta s_{\text{sp}}, \tilde{w}) \mid (8.38), \\ (\hat{x}^+, \hat{d}^+) \in \mathbb{X} \times \mathbb{D}, \tilde{w} \in \tilde{\mathbb{W}}(\hat{x} + e_x, u, \hat{d} + e_d) \} \end{aligned}$$

for each $(\hat{x}, u, \hat{d}) \in \mathbb{X} \times \mathbb{U} \times \mathbb{D}$. The closed-loop estimate error system, defined by (8.6), (8.14), (8.15), (8.37), and (8.38), evolves as

$$\hat{x}^+ = \hat{f}_c(\hat{x}, \hat{\beta}, \tilde{d}) := f(\hat{x} + e_x, \kappa_N(\hat{x}, \hat{\beta}), \hat{d} + e_d) + w - e_x^+ \quad (8.39a)$$

$$\hat{\beta}^+ = \hat{f}_{\beta,c}(\hat{\beta}, \tilde{d}) := \begin{bmatrix} s_{\text{sp}} + \Delta s_{\text{sp}} \\ \hat{d} + e_d + w_d - e_d^+ \end{bmatrix} \quad (8.39b)$$

$$y = \hat{h}_c(\hat{x}, \hat{\beta}, \tilde{d}) := h(\hat{x} + e_x, \kappa_N(\hat{x}, \hat{\beta}), \hat{d} + e_d) + v$$

$$r = \hat{g}_c(\hat{x}, \hat{\beta}, \tilde{d}) := g(\kappa_N(\hat{x}, \hat{\beta}), h_c(\hat{x}, \hat{\beta}, \tilde{d}))$$

where $\hat{\beta} := (s_{\text{sp}}, \hat{d})$.

8.4.1 Steady-state target problem assumptions

Even with bounds on the estimate errors and setpoint and disturbance changes, there are no guarantees the SSTP (8.6) is feasible at all times. Moreover, there is no guarantee the SSTP solutions themselves are robust to disturbance estimate errors. To guarantee robust feasibility of the SSTP (8.6) and robustness of the targets themselves, we make the following assumption.

Assumption 8.31. There exists a compact set $\mathcal{B}_c \subseteq \mathcal{B}$ and constant $\delta_0 > 0$ such that

- (i) $\hat{\mathcal{B}}_c := \{ (s, \hat{d}) \mid (s, d) \in \mathcal{B}_c, |e_d| \leq \delta_0, \hat{d} := d - e_d \in \mathbb{D} \} \subseteq \mathcal{B}$; and
- (ii) z_s is continuous on $\hat{\mathcal{B}}_c$.

Assumption 8.31(i) guarantees robust feasibility of the SSTP so long as $\beta \in \mathcal{B}_c^\infty$ and $\|\mathbf{e}_d\| \leq \delta_0$. Whenever Assumption 8.31(i) is satisfied, it is convenient to define

$$\tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) := \{ \tilde{d} \in \tilde{\mathbb{D}}(\hat{x}, \kappa_N(\hat{x}, \hat{\beta}), \hat{\beta}) \mid \hat{f}_{\beta,c}(\hat{\beta}, \tilde{d}) \in \hat{\mathcal{B}}_c \}$$

for each $(\hat{x}, \hat{\beta}) \in \mathcal{S}_N$. As long as the disturbance always lies in $\tilde{\mathbb{D}}(\hat{x}, \hat{\beta})$, the SSTP is feasible at all times.

In the following lemma, we show Assumption 8.31 holds for some $\mathcal{B}_c = \delta \mathbb{B}^{n_\beta}$ when a rank condition is satisfied by the system linearized at the origin (see Appendix 8.C for proof).

Lemma 8.32. *Suppose Assumptions 8.4 and 8.5 hold, each of the sets $\mathbb{X}, \mathbb{U}, \mathbb{D}$ contain neighborhoods of the origin, the functions f, g, h, ℓ_s are twice continuously differentiable, $\ell_s(0, 0) = 0$, $\partial_{(u,y)} \ell_s(0, 0) = 0$, $\partial_{(u,y)}^2 \ell_s(0, 0)$ is positive definite,*

$$M_1 := \begin{bmatrix} A - I & B \\ H_y C & H_y D + H_u \end{bmatrix} \quad (8.40a)$$

is full row rank, and (A, C) is detectable, where

$$A := \partial_x f(0, 0, 0), \quad B := \partial_u f(0, 0, 0), \quad (8.40b)$$

$$C := \partial_x h(0, 0, 0), \quad D := \partial_u h(0, 0, 0), \quad (8.40c)$$

$$H_y := \partial_y g(0, 0), \quad H_u := \partial_u g(0, 0). \quad (8.40d)$$

Then there exists a compact set $\mathcal{B}_c \subseteq \mathcal{B}$ and a function $z_s : \mathcal{B} \rightarrow \mathbb{X} \times \mathbb{U}$ satisfying all parts of Assumption 8.31. Moreover, $z_s(\beta)$ uniquely solves (8.6) for all $\beta \in \hat{\mathcal{B}}_c$.

8.4.2 Robust stability

In Proposition 8.33, we establish recursive feasibility of the FHOCP given feasibility of the SSTP at each time for sufficiently small $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta})$. For brevity, we defer the proof to Appendix 8.B.2. However, we sketch the proof as follows. First, we show the suboptimal input sequence $\tilde{\mathbf{u}}(x, \hat{\beta})$ is recursively feasible. Second, we establish a cost decrease of the form

$$V_N(\hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}^+) \leq V_N^0(\hat{x}, \hat{\beta}) - a_3 |\delta \hat{x}|^2 + \sigma_r (|\tilde{d}|) \quad (8.41)$$

where $a_3 > 0$, $\sigma_r \in \mathcal{K}_\infty$, and $\delta \hat{x} := \hat{x} - x_s(\hat{\beta})$ is the target-tracking error. Third, we use this cost decrease to show the FHOCP is recursively feasible.

Proposition 8.33. *Suppose Assumptions 8.4 to 8.8 and 8.31 hold and let $\rho > 0$. There exists $\sigma_r \in \mathcal{K}_\infty$ and $a_3, \delta > 0$ such that*

(a) $\tilde{\mathbf{u}}(\hat{x}, \hat{\beta}) \in \mathcal{U}_N(\hat{x}^+, \hat{\beta}^+)$,

(b) (8.41) holds, and

(c) $\hat{x}^+ \in \mathcal{X}_N^\rho(\hat{\beta}^+)$,

for all $\hat{\beta} \in \hat{\mathcal{B}}_c$, $\hat{x} \in \mathcal{X}_N^\rho(\hat{\beta})$ and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta \mathbb{B}^{n_a}$, where $(\hat{x}^+, \hat{\beta}^+)$ are defined as in (8.39).

Finally, we present the main result of this section.

Theorem 8.34. *Suppose Assumptions 8.4 to 8.8 and 8.31 hold and let $\rho > 0$. There exists $\delta > 0$ such that*

(a) *the following set is RPI for the closed-loop system (8.39) with disturbance $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$:*

$$\hat{\mathcal{S}}_N^\rho := \{(\hat{x}, \hat{\beta}) \in \mathcal{S}_N \mid \hat{x} \in \mathcal{X}_N^\rho(\hat{\beta}), \hat{\beta} \in \hat{\mathcal{B}}_c\}; \quad (8.42)$$

(b) *there exist $a_i > 0, i \in \mathbb{I}_{1:3}$ and $\sigma_r \in \mathcal{K}_\infty$ such that*

$$a_1|\delta\hat{x}|^2 \leq V_N^0(\hat{x}, \hat{\beta}) \leq a_2|\delta\hat{x}|^2 \quad (8.43a)$$

$$V_N^0(\hat{x}^+, \hat{\beta}^+) \leq V_N^0(\hat{x}, \hat{\beta}) - a_3|\delta\hat{x}|^2 + \sigma_r(|\tilde{d}|) \quad (8.43b)$$

for all $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$ and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$, given (8.39) and $\delta\hat{x} := \hat{x} - x_s(\hat{\beta})$;

(c) *the closed-loop system (8.39) with disturbance $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$ is RES on $\hat{\mathcal{S}}_N^\rho$ with respect to the target-tracking error $\delta\hat{x} := \hat{x} - x_s(\hat{\beta})$;*

(d) *the closed-loop system (8.39) with disturbance $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$ is RAS on $\hat{\mathcal{S}}_N^\rho$ with respect to $(\delta r, \delta\hat{x})$, where $\delta r := \hat{g}_c(\hat{x}, \hat{\beta}, \tilde{d}) - r_{\text{sp}}$ is the setpoint-tracking error and $\hat{\beta} = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, \tilde{d})$; and*

(e) *if g and h are Lipschitz continuous on bounded sets, then part (d) can be upgraded to RES.*

To prove Theorem 8.34(d,e), we require the following proposition (see Appendix 8.B.3 for proof).

Proposition 8.35. *Let Assumptions 8.4 to 8.8 hold, $\rho, \delta > 0$, and $\mathcal{B}_c \subseteq \mathcal{B}$ be compact. There exist $\sigma_r, \sigma_g \in \mathcal{K}_\infty$ such that*

$$|g_c(\hat{x}, \hat{\beta}) - r_{\text{sp}}| \leq \sigma_r(|\hat{x} - x_s(\hat{\beta})|) \quad (8.44a)$$

$$|\hat{g}_c(\hat{x}, \hat{\beta}, \tilde{d}) - r_{\text{sp}}| \leq |g_c(\hat{x}, \hat{\beta}) - r_{\text{sp}}| + \sigma_g(|\tilde{d}|) \quad (8.44b)$$

for all $\hat{x} \in \mathcal{X}_N^\rho(\beta)$, $\hat{\beta} = (r_{\text{sp}}, z_{\text{sp}}, d) \in \mathcal{B}_c$, and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$. If g and h are Lipschitz on bounded sets, then we can take $\sigma_r(\cdot) := c_r(\cdot)$ and $\sigma_g(\cdot) := c_g(\cdot)$ for some $c_r, c_g > 0$.

Proof of Theorem 8.34. (a)—If $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N$ and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta})$, then $\hat{\beta}^+ := \hat{f}_{\beta,c}(\hat{\beta}, \tilde{d}) \in \hat{\mathcal{B}}_c$ by construction of $\tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta})$, and by Proposition 8.33(c), there exists $\delta > 0$ such that $\hat{x}^+ := \hat{f}_c(\hat{x}, \hat{\beta}, \tilde{d}) \in \mathcal{X}_N^\rho(\hat{\beta})$ so long as $|\tilde{d}| \leq \delta$.

(b)—Theorem 8.29 gives (8.43a), and Proposition 8.33(a,b) and the principle of optimality give (8.43b).

(c)—This follows from part (b) due to Theorem 8.23.

(d)—Let $(\hat{\mathbf{x}}, \hat{\beta}, \tilde{\mathbf{d}}, \mathbf{r})$ satisfy (8.39), $(\hat{x}(0), \hat{\beta}(0)) \in \hat{\mathcal{S}}_N^\rho$, $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$, and $r = \hat{g}_c(\hat{x}, \hat{\beta}, \tilde{d})$. Define $\delta r := r - r_{\text{sp}}$ and $\delta \hat{r} = g_c(\hat{x}, \hat{\beta}) - r_{\text{sp}}$ where $\hat{\beta} = (r_{\text{sp}}, z_{\text{sp}}, \hat{d})$. Then

$$\alpha_1(|\delta \hat{r}|) := a_1[\sigma_r^{-1}(|\delta \hat{r}|)]^2 \leq a_1|\delta \hat{x}|^2 \leq V_N^0(\hat{x}, \hat{\beta})$$

by Proposition 8.35 and part (b). Moreover, V_N^0 is an ISS Lyapunov function on $\hat{\mathcal{S}}_N^\rho$ with respect to $(\delta \hat{r}, \delta \hat{x})$, and RAS on $\hat{\mathcal{S}}_N^\rho$ with respect to $(\delta \hat{r}, \delta \hat{x})$ follows by Theorem 8.23. Then RAS w.r.t. $(\delta \hat{r}, \delta \hat{x})$ and (Rawlings and Ji, 2012, Eq. (1)) gives

$$\begin{aligned} |\delta r(k)| &\leq \sigma_r(|\delta \hat{r}(k)|) + \sigma_g(|\tilde{d}(k)|) \\ &\leq \sigma_r(c\lambda^k|\delta \hat{x}(0)| + \gamma(\|\tilde{\mathbf{d}}\|_{0:k-1})) + \sigma_g(|\tilde{d}(k)|) \\ &\leq \sigma_r(2c\lambda^k|\delta \hat{x}(0)|) + \sigma_r(2\gamma(\|\tilde{\mathbf{d}}\|_{0:k-1})) + \sigma_g(|\tilde{d}(k)|) \\ &\leq \sigma_r(2c\lambda^k|\delta \hat{x}(0)|) + (\sigma_r \circ 2\gamma + \sigma_g)(\|\tilde{\mathbf{d}}\|_{0:k}) \\ &=: \beta_r(|\delta \hat{x}(0)|, k) + \gamma_r(\|\tilde{\mathbf{d}}\|_{0:k}) \end{aligned} \tag{8.45}$$

for all $k \in \mathbb{I}_{\geq 0}$ and some $c > 0$, $\lambda \in (0, 1)$, and $\gamma \in \mathcal{K}$.

(e)—If g and h are Lipschitz continuous on bounded sets, then by Proposition 8.35, we can repeat part (d) with $\sigma_r(\cdot) := c_r(\cdot)$ and some $c_r > 0$. \square

8.5 Offset-free MPC under mismatch

In this section, we show offset-free MPC, *despite (sufficiently small) plant-model mismatch*, is robust to setpoint and disturbance changes. We consider the plant (8.1), setpoint dynamics (8.37), and plant disturbance dynamics

$$w_p^+ = w_p + \Delta w_p. \quad (8.46)$$

With $\alpha := (s_{\text{sp}}, w_p)$ and $\Delta\alpha := (\Delta s_{\text{sp}}, \Delta w_p)$, we have the relationship $\alpha^+ = \alpha + \Delta\alpha$. The SSTP and regulator are designed with the model (8.2), and the estimator is designed with the noisy model (8.16).

8.5.1 Target selection under mismatch

With plant-model mismatch, the connection between the steady-state targets and plant steady states becomes more complicated. To guarantee there is a plant steady state providing offset-free performance and that we can align the plant and model steady states using the disturbance estimate, we make the following assumptions about the SSTP.

Assumption 8.36. There exist compact sets $\mathcal{A}_c \subseteq \mathbb{R}^{n_r} \times \overline{\mathbb{Z}}_y \times \mathbb{W}$ and $\mathcal{B}_c \subseteq \mathcal{B}$ containing the origin, continuous functions $(x_{p,s}, d_s) : \mathcal{A}_c \rightarrow \mathbb{X} \times \mathbb{D}$, and a constant $\delta_0 > 0$ for which

- (a) $\hat{\mathcal{B}}_c$ (as defined in Assumption 8.31) is contained in \mathcal{B} ;
- (b) z_s is Lipschitz continuous on $\hat{\mathcal{B}}_c$;
- (c) for each $\alpha = (s_{\text{sp}}, w_p) \in \mathcal{A}_c$, the pair $(x_{p,s}, d_s) = (x_{p,s}(\alpha), d_s(\alpha))$ is the unique solution to

$$x_{p,s} = f_p(x_{p,s}, u_s(s_{\text{sp}}, d_s), w_p) \quad (8.47a)$$

$$y_s(s_{\text{sp}}, d_s) := h_p(x_{p,s}, u_s(s_{\text{sp}}, d_s), w_p) \quad (8.47b)$$

where $y_s(s_{\text{sp}}, d_s) := h(x_s(s_{\text{sp}}, d_s), u_s(s_{\text{sp}}, d_s), d_s)$;

(d) $(s_{\text{sp}}, d_s(s_{\text{sp}}, w_{\text{p}})) \in \mathcal{B}_c$ for all $(s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$; and

(e) $(s_{\text{sp}}, 0) \in \mathcal{A}_c$ for all $(s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$.

For each $\alpha = (s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$, Assumption 8.36 guarantees there is a unique *model* disturbance $d_s(\alpha)$ to estimate and the SSTP (8.6) is robustly feasible at $\beta = (s_{\text{sp}}, d_s(\alpha))$. Of course, the system cannot be stabilized for unbounded plant-model mismatch. Given Assumption 8.36, we define

$$\mathcal{A}_c(\delta_w) := \{ (s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c \mid |w_{\text{p}}| \leq \delta_w \}$$

$$\mathbb{A}_c(\alpha, \delta_w) := \{ \Delta\alpha \in \mathbb{R}^{n_\alpha} \mid \alpha + \Delta\alpha \in \mathcal{A}_c(\delta_w) \}.$$

Then $\mathcal{A}_c(\delta_w)$ is RPI for the system $\alpha^+ = \alpha + \Delta\alpha$, $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w)$, and if $\|e_d\| \leq \delta_0$, then $\hat{\beta} = (s_{\text{sp}}, d_s(\alpha) - e_d) \in \hat{\mathcal{B}}_c$ and the SSTP is feasible at all times.

Assumption 8.36 can be verified through a linearization analysis that is similar to the standard linear offset-free conditions (Muske and Badgwell, 2002; Pannocchia and Rawlings, 2003) (see Appendix 8.C for proof).

Lemma 8.37. *Suppose the conditions of Lemma 8.32 hold, $f_{\text{p}}, h_{\text{p}}$ are twice continuously differentiable, and*

$$M_2 := \begin{bmatrix} A - I & B_d \\ C & C_d \end{bmatrix} \tag{8.48}$$

is invertible, given the definitions (8.40), $B_d := \partial_d f(0, 0, 0)$, and $C_d := \partial_d h(0, 0, 0)$. Then there exist compact sets $\mathcal{A}_c \subseteq \mathbb{R}^{n_r} \times \overline{\mathbb{Z}}_y \times \mathbb{W}$ and $\mathcal{B}_c \subseteq \mathcal{B}$ containing neighborhoods of the origin and functions $z_s : \mathcal{B} \rightarrow \mathbb{X} \times \mathbb{U}$ and $(x_{\text{p},s}, d_s) : \mathcal{A}_c \rightarrow \mathbb{X} \times \mathbb{D}$ satisfying all parts of Assumption 8.36. Moreover, $z_s(\beta)$ and $(x_{\text{p},s}(\alpha), d_s(\alpha))$ are the unique solutions to (8.6) and (8.47) for all $\alpha = (s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$ and $\beta := (s_{\text{sp}}, d_s(\alpha))$.

8.5.2 State estimation and regulation under mismatch

Given Assumption 8.36, we can define a “true” model state as $x := x_p - \Delta x_s(\alpha)$ where $\Delta x_s := x_{p,s}(\alpha) - x_s(s_{sp}, d_s(\alpha))$ and $\alpha = (s_{sp}, w_p)$. Then the plant (8.1) can be rewritten in terms of the model state x as

$$x^+ = f_p(x + \Delta x_s(\alpha), u, w_p) - \Delta x_s(\alpha^+) \quad (8.49a)$$

$$y = h_p(x + \Delta x_s(\alpha), u, w_p). \quad (8.49b)$$

Alternatively, the plant can be written as (8.16), where

$$w := f_p(x + \Delta x_s(\alpha), u, w_p) - f(x, u, d_s(\alpha)) - \Delta x_s(\alpha^+) \quad (8.50a)$$

$$w_d := d_s(\alpha^+) - d_s(\alpha) \quad (8.50b)$$

$$v := h_p(x + \Delta x_s(\alpha), u, w_p) - h(x, u, d_s(\alpha)). \quad (8.50c)$$

Clearly $\tilde{w} := (w, w_d, v) \in \mathbb{W}(x, u, d)$ by construction. Under Assumption 8.15, the state and disturbance estimator (8.15) is RGES for the constructed model state x and noise vector \tilde{w} .

The noise vector \tilde{w} is still a function of the model state x , input u , and steady-state parameters α . Therefore, we bound it by more manageable variables, i.e., the tracking error $z - z_s(\beta)$, estimate errors e , plant disturbance w_p , and changes to the plant steady-state parameters $\Delta\alpha$. To this end, the following differentiability assumption is required.

Assumption 8.38. The derivatives $\partial_{(x,u)} f_p$ and $\partial_{(x,u)} h_p$ exist and are continuous on $\mathbb{X} \times \mathbb{U} \times \mathbb{W}$. The functions f, h and g are continuously differentiable on $\mathbb{X} \times \mathbb{U} \times \mathbb{D}$ and $\mathbb{U} \times \mathbb{Y}$.

Remark 8.39. Assumption 8.38 implies f, h are Lipschitz continuous on bounded sets.

Consider the closed-loop system

$$x^+ = f_P(x + \Delta x_s(\alpha), \kappa_N(\hat{x}, \hat{\beta}), w_P) - \Delta x_s(\alpha^+) \quad (8.51a)$$

$$\alpha^+ = \alpha + \Delta\alpha \quad (8.51b)$$

$$y = h_P(x + \Delta x_s(\alpha), \kappa_N(\hat{x}, \hat{\beta}), w_P). \quad (8.51c)$$

In the following propositions, we establish cost decreases for estimator and regulator Lyapunov functions for (8.51) (see Appendices 8.B.4 and 8.B.5 for proofs).

Proposition 8.40. *Suppose Assumptions 8.4 to 8.8, 8.15, 8.36, and 8.38 hold. Let $\rho > 0$. There exist $\hat{c}_3, \delta_w > 0$ and $\hat{\sigma}_w, \hat{\sigma}_\alpha \in \mathcal{K}_\infty$ such that*

$$V_e^+ \leq V_e - \hat{c}_3|e|^2 + \hat{\sigma}_w(|w_P|)|\delta\hat{x}|^2 + \hat{\sigma}_\alpha(|\Delta\alpha|) \quad (8.52)$$

so long as $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$, $x \in \mathbb{X}$, $\alpha = (s_{\text{sp}}, w_P) \in \mathcal{A}_c(\delta_w)$, $\Delta\alpha = (\Delta s_{\text{sp}}, \Delta w_P) \in \mathbb{A}_c(\alpha, \delta_w)$, and $|e|, |e^+| \leq \delta_0$, where $V_e(k) := V_e(x(k), d_s(\alpha(k)), \hat{x}(k), \hat{d}(k))$, (8.17), (8.50), and (8.51).

Proposition 8.41. *Let Assumptions 8.4 to 8.8, 8.36, and 8.38 hold and $\rho > 0$. There exist $\tilde{a}_3, \tilde{a}_4, \delta, \delta_w > 0$ and $\tilde{\sigma}_\alpha \in \mathcal{K}_\infty$ such that*

$$V_N^0(\hat{x}^+, \hat{\beta}^+) \leq V_N^0(\hat{x}, \hat{\beta}) - \tilde{a}_3|\delta\hat{x}|^2 + \tilde{a}_4|(e, e^+)|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|) \quad (8.53)$$

so long as $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$, $x \in \mathbb{X}$, $\alpha = (s_{\text{sp}}, w_P) \in \mathcal{A}_c(\delta_w)$, $\Delta\alpha = (\Delta s_{\text{sp}}, \Delta w_P) \in \mathbb{A}_c(\alpha, \delta_w)$, and $|\tilde{d}| \leq \delta$, where $\tilde{d} := (e, e^+, \Delta s_{\text{sp}}, \tilde{w})$, (8.17), (8.50), and (8.51).

8.5.3 Main result

Finally, we state the main result of this section.

Theorem 8.42. *Suppose Assumptions 8.4 to 8.8, 8.15, 8.36, and 8.38 hold and let $\rho > 0$. There exists $\tau, \delta_w, \delta_\alpha > 0$ such that, with*

$$\mathcal{S}_N^{\rho;\tau} := \{ (x, \alpha, \hat{x}, \hat{\beta}) \in \mathbb{X} \times \mathcal{A}_c \times \hat{\mathcal{S}}_N^\rho \mid V_e(x, d_s(\alpha), \hat{x}, \hat{d}) \leq \tau, \alpha = (s_{\text{sp}}, w_{\text{P}}), \hat{\beta} = (s_{\text{sp}}, \hat{d}) \}$$

the following statements hold:

- (a) *the set $\mathcal{S}_N^{\rho;\tau}$ is RPI for the closed-loop system (8.15) and (8.51) with the disturbance $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w) \cap \delta_\alpha \mathbb{B}^{n_\alpha}$;*
- (b) *the closed-loop system (8.15) and (8.51) with the disturbance $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w) \cap \delta_\alpha \mathbb{B}^{n_\alpha}$ is RES on $\mathcal{S}_N^{\rho;\tau}$ with respect to the target-tracking error $\delta\hat{x} := \hat{x} - x_s(\hat{\beta})$; and*
- (c) *the closed-loop system (8.15) and (8.51) with the disturbance $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w) \cap \delta_\alpha \mathbb{B}^{n_\alpha}$ is RES on $\mathcal{S}_N^{\rho;\tau}$ with respect to $(\delta r, \delta\hat{x})$, where $\delta r := r - r_{\text{sp}}$ is the setpoint-tracking error, $\alpha = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, w_{\text{P}})$, $r = g(\kappa_N(\hat{x}, \hat{\beta}), y)$, and (8.51c).*

Proof. (a)—We already have that $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$ and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta \mathbb{B}^{n_d}$ implies $(\hat{x}^+, \hat{\beta}^+) \in \hat{\mathcal{S}}_N^\rho$ for some $\delta > 0$. To keep the trajectory of $(x, \alpha, \hat{x}, \hat{\beta})$ in $\mathcal{S}_N^{\rho;\tau}$ at all times, it suffices to show there exist $\tau, \delta_w, \delta_\alpha > 0$ such that $\alpha \in \mathcal{A}_c(\delta_w)$, $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w) \cap \delta_\alpha \mathbb{B}^{n_\alpha}$, and $V_e := V_e(x, d_s(\alpha), \hat{x}, \hat{d}) \leq \tau$ implies $V_e^+ := V_e(x^+, \hat{x}^+) \leq \tau$ and $|(e, e^+, w)| \leq \delta$.

By Propositions 8.40 and 8.49 (in Appendix 8.B.5), there exist constants $\hat{c}_3, \tilde{c}_e, \delta_w > 0$ and functions $\hat{\sigma}_w, \hat{\sigma}_\alpha, \tilde{\sigma}_w, \tilde{\sigma}_\alpha \in \mathcal{K}_\infty$ satisfying (8.52) and

$$|\tilde{d}|^2 \leq \tilde{c}_e |(e, e^+)|^2 + \tilde{\sigma}_w(|w_{\text{P}}|) |\delta\hat{x}|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|) \quad (8.54)$$

so long as $\alpha = (s_{\text{sp}}, w_{\text{P}}) \in \mathcal{A}_c(\delta_w)$ and $\Delta\alpha \in \mathbb{A}_c(\alpha, \delta_w)$.

Assume, without loss of generality, that $\delta_w < (\frac{4c_2\tilde{c}_3}{a_1c_1\hat{c}_3} \hat{\sigma}_w + \tilde{\sigma}_w)^{-1} (\frac{a_1\delta^2}{\rho})$, which implies

$\frac{2c_2\hat{\sigma}_w(\delta_w)\rho}{a_1\hat{c}_3} < \left(\delta^2 - \frac{\tilde{\sigma}_w(\delta_w)\rho}{a_1}\right) \frac{c_1}{2\tilde{c}_e}$ and $\frac{\tilde{\sigma}_w(\delta_w)\rho}{a_1} < \delta^2$. Then we can take

$$\tau \in \left(\frac{2c_2\hat{\sigma}_w(\delta_w)\rho}{a_1\hat{c}_3}, \left(\delta^2 - \frac{\tilde{\sigma}_w(\delta_w)\rho}{a_1} \right) \frac{c_1}{2\tilde{c}_e} \right)$$

which implies $\frac{\tau\hat{c}_3}{2c_2} > \frac{\hat{\sigma}_w(\delta_w)\rho}{a_1}$ and $\delta^2 > \frac{2\tilde{c}_e\tau}{c_1} + \frac{\tilde{\sigma}_w(\delta_w)\rho}{a_1}$.

From (8.52), we have

$$V_e^+ \leq \begin{cases} \frac{\tau}{2} + \frac{\hat{\sigma}_w(\delta_w)\rho}{a_1} + \hat{\sigma}_\alpha(|\Delta\alpha|), & V_e \leq \frac{\tau}{2} \\ \tau - \frac{\tau\hat{c}_3}{2c_2} + \frac{\hat{\sigma}_w(\delta_w)\rho}{a_1} + \hat{\sigma}_\alpha(|\Delta\alpha|), & \frac{\tau}{2} < V_e \leq \tau. \end{cases}$$

But $\hat{c}_3 \leq c_2$ (otherwise we could show $V_e < 0$ with $w_p = 0$, $\Delta\alpha = 0$, and $e \neq 0$) so $\frac{\tau}{2} \geq \frac{\tau\hat{c}_3}{2c_2} > \frac{\hat{\sigma}_w(\delta_w)\rho}{a_1}$ and we have $V_e^+ \leq \tau$ so long as $|\Delta\alpha| \leq \delta_{\alpha,1} := \hat{\sigma}_\alpha^{-1}\left(\frac{\tau\hat{c}_3}{2c_2} - \frac{\hat{\sigma}_w(\delta_w)\rho}{a_1}\right)$, which is positive by construction. Moreover, $V_e, V_e^+ \leq \tau$ implies $|(e, e^+)|^2 = |e|^2 + |e^+|^2 \leq \frac{2\tau}{c_1}$ and by (8.54),

$$\begin{aligned} |\tilde{d}|^2 &\leq \tilde{c}_e|(e, e^+)|^2 + \tilde{\sigma}_w(|w_p|)|\hat{x} - x_s(\hat{\beta})|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|) \\ &\leq \frac{2\tilde{c}_e\tau}{c_1} + \tilde{\sigma}_w(\delta_w)\rho^2 + \tilde{\sigma}_\alpha(\delta_\alpha) \\ &\leq \delta^2 \end{aligned}$$

so long as $|\Delta\alpha| \leq \delta_{\alpha,2} := \tilde{\sigma}_\alpha^{-1}\left(\delta^2 - \frac{2\tilde{c}_e\tau}{c_1} - \frac{\tilde{\sigma}_w(\delta_w)\rho}{a_1}\right)$, which exists and is positive by construction.

Finally, we can take $\delta_\alpha := \min\{\delta_{\alpha,1}, \delta_{\alpha,2}\}$ to achieve $(x, \alpha, \hat{x}, \hat{\beta}) \in \mathcal{S}_N^{\rho;\tau}$ at all times.

(b)—From part (a), we already have $\tau, \delta_w, \delta_\alpha > 0$ such that $\mathcal{S}_N^{\rho;\tau}$ is RPI. By Assumption 8.15 and Theorem 8.34 we have (8.18a) and (8.43a) at all times for some $a_1, a_2, c_1, c_2 > 0$. By Propositions 8.40 and 8.41, there exist $\hat{c}_3, \tilde{a}_3, \tilde{a}_4 > 0$ and $\hat{\sigma}_w, \hat{\sigma}_\alpha, \sigma_\alpha \in \mathcal{K}_\infty$ such that (8.52) and (8.53) at all times. Assume, without loss of generality, that $\delta_w < \hat{\sigma}_w^{-1}\left(\min\left\{\frac{c_1\tilde{a}_3}{\tilde{a}_4}, \frac{a_3\hat{c}_3}{a_4} - \frac{c_1}{c_1+c_2}\right\}\right)$. By Theorem 8.28, the system is RES on $\mathcal{S}_N^{\rho;\tau}$ w.r.t. $\delta\hat{x}$.

(c)—By Proposition 8.35, there exist $c_r, c_g > 0$ such that $|\delta r| \leq c_r |\delta \hat{x}| + c_g |\tilde{d}|$ where $\tilde{d} := (e, e^+, \Delta_{s_{sp}}, \tilde{w})$. Combining this inequality with (8.18a), (8.52), and (8.54) gives

$$|\delta r| \leq c_{r,x} |\delta \hat{x}| + c_{r,e} |e| + \tilde{\gamma}_r (|\Delta \alpha|)$$

where $c_{r,x} := c_r + c_g (\sqrt{\tilde{\sigma}_\alpha(\delta_w)} + \sqrt{\tilde{c}_e \hat{\sigma}_\alpha(\delta_w)})$, $c_{r,e} := c_g \sqrt{\tilde{c}_e} (1 + \sqrt{c_2 - \hat{c}_3})$, and $\tilde{\gamma}_r := c_g (\sqrt{\tilde{\sigma}_\alpha} + \sqrt{\tilde{c}_e \hat{\sigma}_\alpha})$. Then

$$|(\delta r, e)| \leq \tilde{c}_r |(\delta \hat{x}, e)| + \tilde{\gamma}_r (|\Delta \alpha|)$$

where $\tilde{c}_r := c_{r,x} + c_{r,e} + 1$. Finally, RES w.r.t. $\delta \hat{x}$ gives

$$|(\delta \hat{x}(k), e(k))| \leq \tilde{c} \lambda^k |(\delta \hat{x}(0), \bar{e})| + \sum_{j=0}^k \lambda^j \tilde{\gamma} (|\Delta \alpha(k-j)|)$$

for some $\tilde{c} > 0$, $\lambda \in (0, 1)$, and $\tilde{\gamma} \in \mathcal{K}$, and therefore

$$\begin{aligned} |(\delta r(k), e(k))| &\leq \tilde{c}_r |(\delta \hat{x}(k), e(k))| + \tilde{\gamma}_r (|\Delta \alpha(k)|) \\ &\leq c \lambda^k |(\delta \hat{x}(0), \bar{e})| + \sum_{j=0}^k \lambda^j \gamma (|\Delta \alpha(k-j)|) \end{aligned}$$

where $c := \tilde{c}_r \tilde{c} > 0$ and $\gamma := \tilde{c}_r \tilde{\gamma} + \tilde{\gamma}_r \in \mathcal{K}_\infty$. □

8.6 Examples

In this section, we illustrate the main results using the example systems depicted in Figure 8.1. We compare two MPCs in our experiments.




Figure 8.1: Example systems.

Offset-free MPC The offset-free MPC (OFMPC) uses (8.6) and (8.14) and the following MHE problem:

$$\min_{(\mathbf{x}, \mathbf{d}) \in \mathbb{X}^{T_k+1} \times \mathbb{D}^{T_k+1}} V_T^{\text{MHE}}(\mathbf{x}, \mathbf{d}, \mathbf{u}, \mathbf{y}) \quad (8.55)$$

where $T_k := \min \{ k, T \}$, $T \in \mathbb{I}_{>0}$, and

$$\begin{aligned} V_T^{\text{MHE}}(\mathbf{x}, \mathbf{d}, \mathbf{u}, \mathbf{y}) := & \sum_{j=0}^{T_k-1} |x_{j+1} - f(x_j, u(j), d_j)|_{Q_w^{-1}}^2 \\ & + |d_{j+1} - d_j|_{Q_d^{-1}}^2 + |y(j) - h(x_j, u(j), d_j)|_{R_v^{-1}}^2. \end{aligned} \quad (8.56)$$

For simplicity, a prior term is not used. Let $\hat{x}(j; \mathbf{u}, \mathbf{y})$ and $\hat{d}(j; \mathbf{u}, \mathbf{y})$ denote solutions to the above problem, and define the estimates by

$$\hat{x}(k) := \hat{x}(k; \mathbf{u}_{k-T_k:k-1}, \mathbf{y}_{k-T_k:k-1}), \quad \hat{d}(k) := \hat{d}(k; \mathbf{u}_{k-T_k:k-1}, \mathbf{y}_{k-T_k:k-1}).$$

Tracking MPC The nominal tracking MPC (TMPC) uses (8.6) and (8.14) and the following MHE problem:

$$\min_{\mathbf{x} \in \mathbb{X}^{T_k+1}} V_T^{\text{MHE}}(\mathbf{x}, 0, \mathbf{u}_{k-T_k:k-1}, \mathbf{y}_{k-T_k:k-1}) \quad (8.57)$$

With solutions denoted by $\hat{x}(j; \mathbf{u}, \mathbf{y})$, we define the estimates by

$$\hat{x}(k) := \hat{x}(k; \mathbf{u}_{k-T_k:k-1}, \mathbf{y}_{k-T_k:k-1}), \quad \hat{d}(k) := 0.$$

We also construct, in the proof of the following lemma, terminal ingredients satisfying Assumption 8.7.

Lemma 8.43. *Suppose Assumptions 8.4 to 8.6 and 8.31 hold with $\mathcal{B} = \hat{\mathcal{B}}_c$ and $n_c = 0$, let $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{n_u \times n_u}$ be positive definite, and $\partial_{(x,u)}^2 f_i, i \in \mathbb{I}_{1:n}$ exist and are bounded on $\mathbb{X} \times \mathbb{U} \times \mathbb{D}$. For each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}$, let*

$$A(\beta) := \partial_x f(z_s(\beta), d), \quad B(\beta) := \partial_u f(z_s(\beta), d).$$

If $(A(\beta), B(\beta))$ is stabilizable for each $\beta \in \mathcal{B}$, then there exist functions $\kappa_f : \mathbb{X} \times \mathcal{B}$ and $P_f : \mathcal{B} \rightarrow \mathbb{R}^{n \times n}$, and a constant $c_f > 0$ satisfying Assumptions 8.7 and 8.8.

Proof. Throughout this proof, we let $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, \beta) \in \mathcal{B}$. Since $(A(\beta), B(\beta))$ is stabilizable, there exists a positive definite $P(\beta)$ that uniquely solves the following discrete algebraic Riccati equation,

$$P = A^\top P A + Q - A^\top P B (B^\top P B + R)^{-1} B^\top P A$$

where dependence on β has been suppressed for brevity. The solution P is continuous at each (A, B, Q, R) such that (A, B) is stabilizable and (Q, R) are positive definite (Sun, 1998).⁴

⁴In fact, Sun (1998) only needed $(A, Q^{1/2})$ detectable to derive perturbation bounds. But Assumption 8.8

Moreover, since f is twice differentiable and (x_s, u_s) are continuous on \mathcal{B} , then (A, B) must be continuous on \mathcal{B} . Therefore P is continuous on \mathcal{B} and Assumption 8.8 holds for $P_f(\beta) := 2P(\beta)$.

Next, with $K := PB(B^\top PB + R)^{-1}$, $A_K := A - BK$, and $Q_K := Q + K^\top RK$, we have $A_K^\top P_f A_K - P_f = -2Q_K$, where dependence on β has been suppressed for brevity. Then

$$V_f(\bar{x}^+, \beta) - V_f(x, \beta) \leq -2|\delta x|_{Q_K(\beta)}^2 \quad (8.58)$$

where $\bar{x}^+ := A_K(\beta)\delta x + x_s(\beta)$ and $\delta x := x - x_s(\beta)$. Since the second derivatives $\partial_{(x,u)}^2 f_i, i \in \mathbb{I}_{1:n}$ are bounded, there exists $\bar{c} > 0$ (independent of β) such that $|x^+ - \bar{x}^+| \leq \bar{c}|\delta x|^2$ where $x^+ := f(x, \kappa_f(x, \beta), d)$ and $\kappa_f(x, \beta) := -K(\beta)\delta x + u_s(\beta)$.⁵ Therefore, with $a(\beta) := 2\bar{c}\bar{\sigma}([A_K(\beta)]^\top P_f(\beta))$ and $b(\beta) := \bar{c}^2\bar{\sigma}(P_f(\beta))$, we have

$$|V_f(x^+, \beta) - V_f(\bar{x}^+, \beta)| \leq a(\beta)|\delta x|^3 + b(\beta)|\delta x|^4 \quad (8.59)$$

and combining (8.58) with (8.59), we have

$$\begin{aligned} & V_f(x^+, \beta) - V_f(x, \beta) + \ell(x, \kappa_f(x, \beta), \beta) \\ & \leq -|\delta x|_{Q_K(\beta)}^2 + V_f(x^+, \beta) - V_f(\bar{x}^+, \beta) \\ & \leq -[c(\beta) - b(\beta)|\delta x| - a(\beta)|\delta x|^2]|\delta x|^2 \end{aligned} \quad (8.60)$$

where $c(\beta) := \underline{\sigma}(Q_K(\beta))$. The polynomial $p_\beta(s) = c(\beta) - b(\beta)s - a(\beta)s^2$ has roots at $s_\pm(\beta) = \frac{-b(\beta) \pm \sqrt{[b(\beta)]^2 + 4a(\beta)c(\beta)}}{2a(\beta)}$ and is positive in between. Moreover, s_\pm are continuous over \mathcal{B} because (a, b, c) are as well, and $s_\pm(\beta)$ are positive and negative, respectively. Define $c_f := \min_{\beta \in \mathcal{B}} \underline{\sigma}(P_f(\beta))[s_+(\beta)]^2$ which exists and is positive due to continuity and positive-definiteness of Q , so we get this automatically.

⁵This follows by applying Taylor's theorem to $e(x, \beta) := x^+ - \bar{x}^+$ at $(x_s(\beta), d)$ and noting the intercept and first derivative (in x) is zero.

ity of x_+ and $\underline{\sigma}(P_f(\cdot))$ and compactness of \mathcal{B} . Finally, we have that $V_f(x, \beta) \leq c_f$ implies $\underline{\sigma}(P_f(\beta))|\delta x|^2 \leq V_f(x, \beta) \leq c_f$ and therefore $|\delta x| \leq \sqrt{\frac{c_f}{\underline{\sigma}(P_f(\beta))}} \leq s_+(\beta)$ and (8.60) implies Assumption 8.7 with $P_f(\beta)$ and $c_f > 0$ as constructed. \square

8.6.1 Simple pendulum

Consider the following nondimensionalized pendulum system (Figure 8.1a):

$$\dot{x} = F_P(x, u, w_P) := \begin{bmatrix} x_2 \\ \sin x_1 - (w_P)_1^2 x_2 + (\hat{k} + (w_P)_2)u + (w_P)_3 \end{bmatrix} \quad (8.61a)$$

$$y = h_P(x, u, w_P) := x_1 + (w_P)_4 \quad (8.61b)$$

$$r = g(u, y) := y \quad (8.61c)$$

where $(x_1, x_2) \in \mathbb{X} := \mathbb{R}^2$ are the angle and angular velocity, $u \in \mathbb{U} := [-1, 1]$ is the (dimensionless) motor voltage, $\hat{k} = 5 \text{ rad/s}^2$ is the estimated motor gain, $(w_P)_1$ is an air resistance factor, $(w_P)_2$ is the error in the motor gain estimate, $(w_P)_3$ is an externally applied torque, and $(w_P)_4$ is the measurement noise. Let $\psi(t; x, u, w_P)$ denote the solution to (8.61) at time t given $x(0) = x$, $u(t) = u$, and $w_P(t) = w_P$. We model the discretization of (8.61) by

$$x^+ = f_P(x, u, w_P) := x + \Delta F_P(x, u, w_P) + (w_P)_5 r_d(x, u, w_P) \quad (8.62a)$$

where $(w_P)_5$ scales the discretization error, r_d is a residual function given by

$$r_d(x, u, w_P) := \int_0^\Delta [F_P(x(t), u, w_P) - F_P(x, u, w_P)] dt \quad (8.62b)$$

and $x(t) = \psi(t; x, u, w_P)$. Assuming a zero-order hold on the input u and disturbance w_P , the system (8.61) is discretized (exactly) as (8.62) with $(w_P)_5 \equiv 1$. We model the system with

$w_p = w(d) := (0, 0, d, 0, 0)$, i.e.,

$$x^+ = f(x, u, d) := f_p(x, u, w(d)) = x + \Delta \begin{bmatrix} x_2 \\ \sin x_1 + \hat{k}u + d \end{bmatrix} \quad (8.63a)$$

$$y = h(x, u, d) := h_p(x, u, w(d)) = x_1 \quad (8.63b)$$

and therefore we do not need access to the residual function r_d to design the offset-free MPC.

For the following simulations, assume $w_p \in \mathbb{W} := [-3, 3]^3 \times [-0.05, 0.05] \times \{0, 1\}$, and let the sample time be $\Delta = 0.1$ s. Regardless of objective ℓ_s , the SSTP (8.6) is uniquely solved by

$$x_s(\beta) := \begin{bmatrix} r_{\text{sp}} \\ 0 \end{bmatrix}, \quad u_s(\beta) := -\frac{1}{\hat{k}}(\sin r_{\text{sp}} + d)$$

for each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}_c$, where

$$\mathcal{B}_c := \{ (r, u, y, d) \in \mathbb{R}^4 \mid |\sin r + d|, |\sin y + d| \leq \hat{k}, |u| \leq 1 \}$$

and $\delta_0 > 0$. Likewise, the solution to (8.47) is

$$x_{p,s}(\alpha) := \begin{bmatrix} r_{\text{sp}} \\ 0 \end{bmatrix}, \quad d_s(\alpha) := \frac{\hat{k}(w_p)_3 - (w_p)_2 \sin r_{\text{sp}}}{\hat{k} + (w_p)_2} \quad (8.64)$$

for each $\alpha = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, w_p) \in \mathcal{A}_c$, where

$$\mathcal{A}_c := \{ (r, u, y, w) \in \mathbb{R}^3 \times \mathbb{W} \mid |\sin r + (w_p)_3|, |\sin y + (w_p)_3| \leq \hat{k} + (w_p)_2, |u| \leq 1 \}.$$

Notice that \mathcal{A}_c and \mathcal{B}_c are compact and satisfy Assumption 8.36. We define a regulator with $N := 20$, $\mathbb{U} := [-1, 1]$, $\ell_s(u, y) = |u|^2 + |y|^2$, $\ell(x, u, \Delta u, \beta) := |x - x_s(\beta)|^2 + 10^{-2}(u -$

$u_s(\beta))^2 + 10^2(\Delta u)^2$,⁶ $V_f(x, \beta) := |x - x_s(\beta)|_{P_f(\beta)}^2$, and $\mathbb{X}_f := \text{lev}_{c_f} V_f$, where $P_f(\beta)$ and $c_f \approx 0.4364$ are chosen according to the proof of Lemma 8.43 to satisfy Assumptions 8.7 and 8.8. Assumption 8.5 is satisfied trivially and Assumptions 8.4, 8.36, and 8.38 are satisfied since smoothness of F implies that ψ , r , and f are smooth (Hale, 1980, Thm. 3.3). Finally, we use MHE designs (8.55) and (8.57) for the offset-free MPC and tracking MPC, respectively, where $T = 5$, $Q_w := \begin{bmatrix} 10^{-3} & \\ & 10^{-6} \end{bmatrix}$, and $Q_d := R_v := 1$. While the estimators defined by (8.55) and (8.57) should be RGES (Allan and Rawlings, 2021), it is not known if they satisfy Assumption 8.15. If Assumption 8.15 is satisfied, then Theorem 8.42 gives robust stability with respect to the tracking errors.

We present the results of numerical experiments in Figure 8.2. To ensure numerical accuracy, the plant (8.61) is simulated by four 4th-order Runge-Kutta steps per sample time. Unless otherwise specified, we consider, in each simulation, unmodeled air resistance $(w_p)_1 \equiv 1$, motor gain error $(w_p)_2 \equiv 2$, an exogenous torque $(w_p)_3(k) = 3H(k - 240)$, the discretization parameter $(w_p)_4 \equiv 1$, no measurement noise $(w_p)_5 \equiv 0$, and a reference signal $r_{\text{sp}}(k) = \pi H(5 - k) + \frac{\pi}{2} H(k - 120)$, where H denotes the unit step function. The setpoint brings the pendulum from the resting state $x_1 = \pi$, to the upright position $x_1 = 0$, to the half-way position $x_1 = \frac{\pi}{2}$.

In the first experiment, we consider the case without plant-model mismatch, i.e., $(w_p)_1 \equiv 0$ and $(w_p)_2 \equiv 0$ (Figure 8.2a). Both offset-free and tracking MPC remove offset after the setpoint changes. However, only offset-free MPC removes offset after the disturbance is injected. Without a disturbance model, the tracking MPC cannot produce useful steady-state targets, and the pendulum drifts far from the setpoint. Moreover, the tracking MPC produces pathological state estimates, with nonzero velocity at steady state.

The second experiment considers the case with plant-model mismatch, i.e., $(w_p)_1 \equiv 1$

⁶The $\Delta u(k) := u(k) - u(k - 1)$ penalty is a standard generalization used by practitioners to “smooth” the closed-loop response in a tuneable fashion.

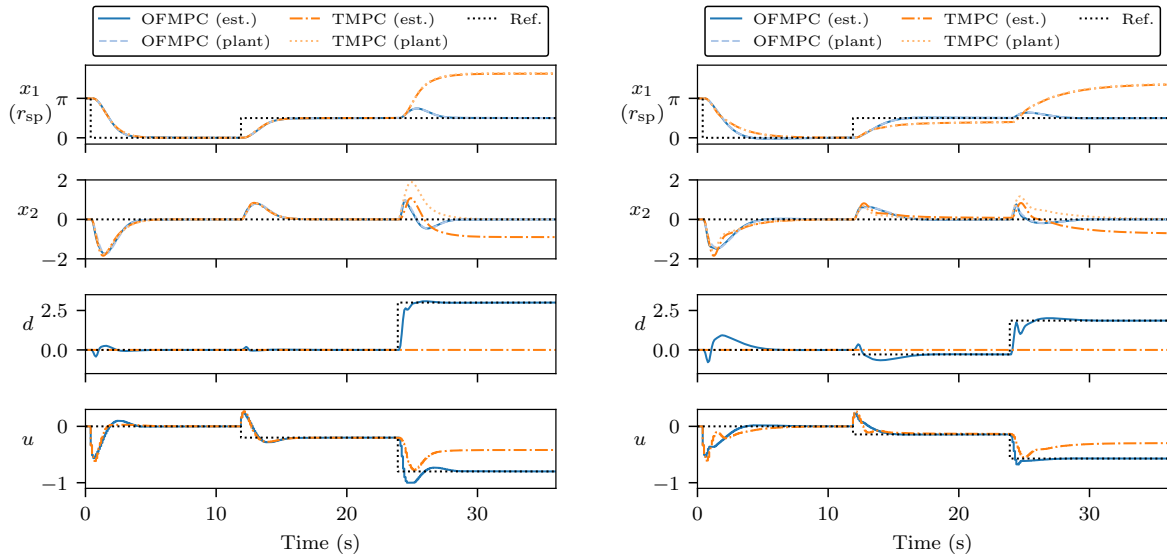
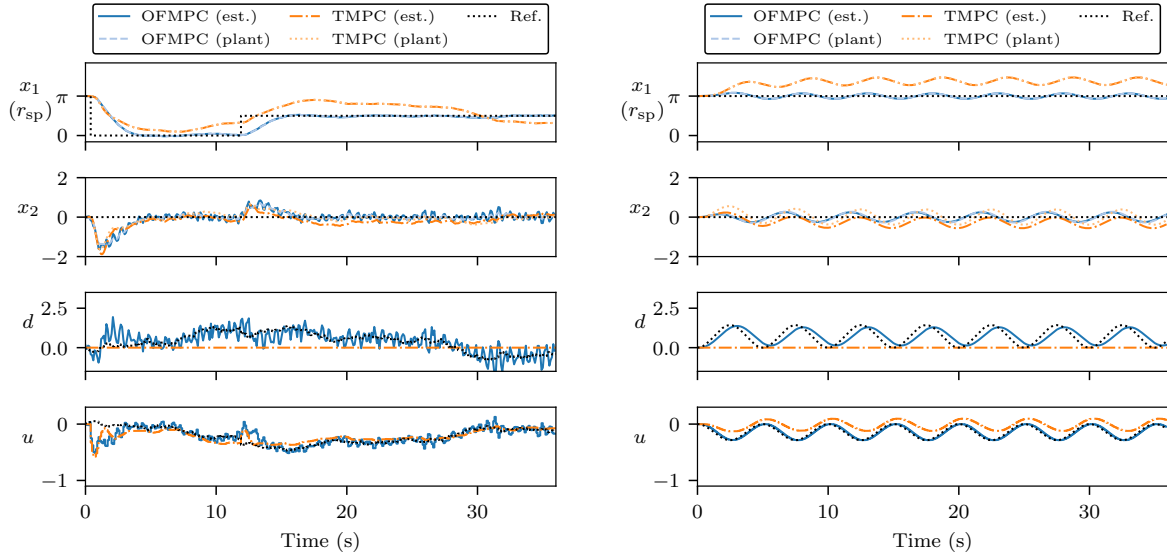
(a) No mismatch: $(w_p)_1 \equiv 0$ and $(w_p)_2 \equiv 0$.(b) Mismatch: $(w_p)_1 \equiv 1$ and $(w_p)_2 \equiv 2$.(c) Noise and mismatch: $(w_p)_3^+ = (w_p)_3 + (\Delta w_p)_3$, $(\Delta w_p)_3 \sim \mathcal{N}(0, 10^{-2})$, and $(w_p)_4 \sim \mathcal{N}(0, 10^{-4})$.(d) Oscillating disturbance and mismatch: $(w_p)_3(k) = 1 - \cos(\frac{2\pi k}{50})$ and $r_{sp}(k) \equiv \pi$.

Figure 8.2: Simulated closed-loop trajectories for the offset-free MPC and tracking MPC of (8.61). Solid blue and dot-dashed orange lines represent the closed-loop estimates and inputs (\hat{x}, \hat{d}, u) for the offset-free MPC and tracking MPC. Dashed blue and dotted orange lines represent the closed-loop plant states x_p for the offset-free MPC and tracking MPC. Dotted black lines represent the intended steady-state targets and disturbance values $(x_{p,s}, d_s, u_s)$ found by solving (8.6) and (8.47). We set $(w_p)_1 \equiv 1$, $(w_p)_2 \equiv 2$, $(w_p)_3(k) = 3H(k - 240)$, $(w_p)_4 \equiv 1$, $(w_p)_5 \equiv 0$, and $r_{sp}(k) = \pi H(5 - k) + \frac{\pi}{2} H(k - 120)$, unless otherwise specified.

and $(w_p)_2 \equiv 2$ (Figure 8.2b). As in the first experiment, both the tracking MPC and offset-free MPC bring the pendulum to the upright position $x_1 = 0$, without offset. However, only the offset-free MPC brings the pendulum to the half-way position $x_1 = \frac{\pi}{2}$. The tracking MPC, not accounting for motor gain errors, provides an insufficient force and does not remove offset. Note the intended disturbance estimate $d_s = \frac{13}{7}$ is a smaller value than the actual injected disturbance $(w_p)_3 = 3$, as underestimation of the motor gain necessitates a smaller disturbance value to be corrected. Again, the tracking MPC produces pathological state estimates.

The third experiment follows the second, except the exogenous torque is an integrating disturbance $(w_p)_3^+ = (w_p)_3 + (\Delta w_p)_3$ where $(w_p)_3 \sim N(0, 10^{-2})$ and we have measurement noise $(w_p)_5 \sim N(0, 10^{-4})$ (Figure 8.2c). In this experiment, we see the remarkable ability of offset-free MPC to track a reference subject to random disturbances. While the tracking MPC is robust to the disturbance $(w_p)_3$, it is not robust to the disturbance changes $(\Delta w_p)_3$ and wanders far from the setpoint as a result. On the other hand, offset-free MPC is robust to both and exhibits practically offset-free performance. We remark that, while the example is mechanical in nature, we are illustrating a behavior that is often desired in chemical process control, where process specifications must be met despite constantly, but slowly varying upstream conditions.

In the fourth and final experiment, the pendulum maintains the resting position $r_{sp}(k) \equiv \pi$ subject to an oscillating torque $(w_p)_3(k) = 1 - \cos(\frac{2\pi k}{50})$ (Figure 8.2d). Tracking MPC wanders away from the setpoint, whereas offset-free MPC oscillates around it with small amplitude. We note the disturbance estimate \hat{d} does not ever “catch” the intended value d_s as the disturbance model has no ability to match its *velocity* or *acceleration*. More general integrator schemes (e.g., double or triple integrators) could provide more dynamic tracking performance at the cost of a higher disturbance dimension (c.f., Maeder and Morari (2010) or (Zagrobelny, 2014, Ch. 5)).




Figure 8.3: Nominal steady states for the CSTR (8.65).

8.6.2 Continuous stirred-tank reactor

Consider the following nonisothermal continuous stirred-tank reactor (CSTR) (Hicks and Ray, 1971; Kameswaran and Biegler, 2006) (Figure 8.1b):

$$\begin{aligned} \dot{x} &= F_P(x, u, w_P) \\ &:= \begin{bmatrix} \theta^{-1}(1 + (w_P)_1 - x_1) - ke^{(w_P)_2 - M/x_2} x_1 \\ \theta^{-1}(x_f - x_2) + ke^{(w_P)_2 - M/x_2} x_1 - \gamma u(x_2 - x_c - (w_P)_3) \end{bmatrix} \end{aligned} \quad (8.65a)$$

$$y = h_P(x, u, w_P) := x_2 + (w_P)_4 \quad (8.65b)$$

$$r = g(u, y) := y \quad (8.65c)$$

where $(x_1, x_2) \in \mathbb{X} := \mathbb{R}_{\geq 0}^2$ are the dimensionless concentration and temperature, $u \in \mathbb{U} := [0, 2]$ is the dimensionless coolant flowrate, $\theta = 20$ s is the residence time, $k = 300$ s⁻¹ is the rate coefficient, $M = 5$ is the dimensionless activation energy, $x_f = 0.3947$ and $x_c = 0.3816$ are dimensionless feed and coolant temperatures, $\gamma = 0.117$ s⁻¹ is the heat transfer coefficient, $(w_P)_1$ is a kinetic modeling error, $(w_P)_2$ is a change to the coolant temperature,

and $(w_p)_4$ is the measurement noise. Again, we discretize the system (8.65) via the equations (8.62), where the continuous system is recovered with $(w_p)_5 = 1$ and zero-order holds on u and w_p . The system is modeled with $w_p = w(d) := (0, d, 0, 0, 0)$, i.e.,

$$x^+ = f(x, u, d) := x + \Delta \begin{bmatrix} \theta^{-1}(1 - x_1) - ke^{-M/x_2}x_1 \\ \theta^{-1}(x_f - x_2) + ke^{-M/x_2}x_1 - \gamma u(x_2 - x_c - d) \end{bmatrix} \quad (8.66a)$$

$$y = h(x, u, d) := x_2. \quad (8.66b)$$

The goal in the following simulations is to control the CSTR (8.65) from a nominal steady state $(x(0), u(-1)) \approx (0.9831, 0.3918, 0.8305)$ to a temperature setpoint $r_{\text{sp}} \in [0.6, 0.7]$. In this range the nominal steady states are unstable. Moreover, there is a Hopf bifurcation at $(x_{\text{Hopf}}, u_{\text{Hopf}}) \approx (0.1728, 0.7009, 0.6973)$. We plot the nominal steady states (i.e., $w_p = 0$) along with the initial steady state $x(0)$ and the Hopf bifurcation x_{Hopf} in Figure 8.3.

For the following simulations, assume disturbance set is $w_p \in \mathbb{W} := [-0.05, 0.05]^4 \times \{0, 1\}$, and let the sample time be $\Delta = 1$ s. Regardless of objective ℓ_s , the SSTP (8.6) is uniquely solved by

$$x_s(\beta) := \begin{bmatrix} \frac{1}{1 + \theta ke^{-M/r}} \\ r_{\text{sp}} \end{bmatrix}, \quad u_s(\beta) := \frac{x_f - r + 1 - (x_s(\beta))_1}{\theta\gamma(r - x_c - d)}$$

for each $\beta = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d) \in \mathcal{B}_c$, where

$$\mathcal{B}_c := [0.6, 0.7] \times \mathbb{U} \times [0.6, 0.7] \times [-0.1, 0.1]$$

and $\delta_0 > 0$. Likewise, the solution to (8.47) is

$$x_{P,s}(\alpha) := \begin{bmatrix} r_{\text{sp}} \\ 0 \end{bmatrix}, \quad d_s(\alpha) \quad (8.67)$$

for each $\alpha = (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$, where

$$\mathcal{A}_c := [0.6, 0.7] \times \mathbb{U} \times [0.6, 0.7] \times \mathbb{W}.$$

It is straightforward to verify \mathcal{A}_c and \mathcal{B}_c are compact and satisfy Assumption 8.36.

We define a regulator with $N := 150$, $\ell(x, u, \Delta u, \beta) := |x - x_s(\beta)|_Q^2 + 10^{-3}(u - u_s(\beta))^2 + (\Delta u)^2$, $Q = \begin{bmatrix} 10^{-3} & \\ & 1 \end{bmatrix}$, $V_f(x, \beta) := |x - x_s(\beta)|_{P_f(\beta)}^2$, and $\mathbb{X}_f := \text{lev}_{c_f} V_f$, where $P_f(\beta)$ and $c_f \approx 6.7031 \times 10^{-16}$ are chosen according to the proof of Lemma 8.43 to satisfy Assumption 8.7.⁷ Finally, we use MHE designs (8.55) and (8.57) for the offset-free MPC and tracking MPC, respectively, where $T := N$, $Q_w := 10^{-4}I$, $Q_d := 10^{-2}$, and $R_v := 1$. As in the simple pendulum example, if Assumption 8.15 is satisfied, then Theorem 8.42 implies the offset-free MPC can robustly track setpoints despite plant-model mismatch.

The results of the CSTR experiments are presented in Figure 8.4. Throughout these experiments, the plant (8.65) is simulated by ten 4th-order Runge-Kutta steps per sample time. Unless otherwise specified, each simulation is carried out with error in the feed concentration $(w_{\text{p}})_1 \equiv -0.05$, error in the activation energy $(w_{\text{p}})_2 \equiv -0.05$, a step in the coolant temperature $(w_{\text{p}})_3(k) = -0.05H(k - 300)$, no measurement noise $(w_{\text{p}})_4 \equiv 0$, the discretization parameter $(w_{\text{p}})_5 \equiv 1$, and a constant reference signal $r_{\text{sp}} \equiv 0.65$.

In the first experiment, we consider the case without plant-model mismatch, i.e., $(w_{\text{p}})_1 \equiv 0$ and $(w_{\text{p}})_2 \equiv 0$ (Figure 8.4a). As in the pendulum experiment, both offset-free and tracking MPC remove offset after the setpoint changes, but only offset-free MPC removes offset after the disturbance is injected. We also note that, after the disturbance is injected, the tracking MPC state estimates are slightly different than the plant states.

We consider plant-model mismatch $(w_{\text{p}})_1 \equiv -0.05$ and $(w_{\text{p}})_2 \equiv -0.05$ in the second

⁷While c_f was chosen near machine precision, the CSTR tends to evolve to the nearest stable steady state, and the horizon is chosen long enough to easily achieve this steady state to a high degree of precision. Thus, the system remains robust despite the tight terminal constraint.

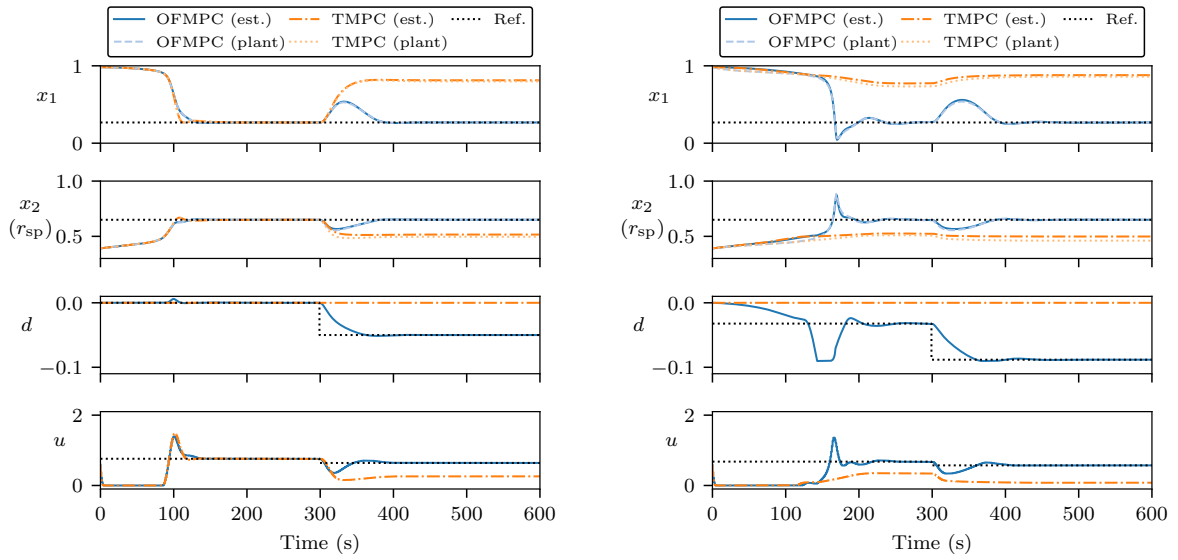
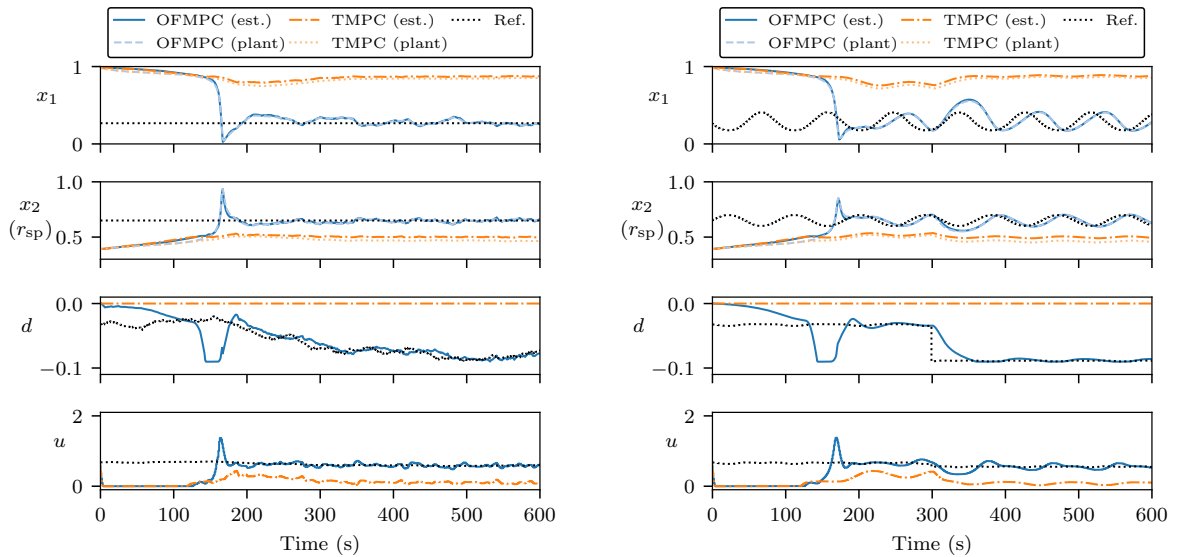
(a) No mismatch: $(w_p)_1 \equiv 0$ and $(w_p)_2 \equiv 0$.(b) Mismatch: $(w_p)_1 \equiv -0.05$ and $(w_p)_2 \equiv -0.05$.(c) Noise and mismatch: $(w_p)_3^+ = (w_p)_3 + (\Delta w_p)_3$, $(\Delta w_p)_3 \sim N(0, 10^{-6})$, and $(w_p)_4 \sim N(0, 10^{-4})$. (d) Oscillating setpoint: $r_{sp}(k) = 0.05 \sin(\frac{2\pi k}{90}) + 0.65$.

Figure 8.4: Simulated closed-loop trajectories for the offset-free MPC and tracking MPC of the CSTR (8.65). Solid blue and dot-dashed orange lines represent the closed-loop estimates and inputs (\hat{x}, \hat{d}, u) for the offset-free MPC and tracking MPC. Dashed blue and dotted orange lines represent the closed-loop plant states x_p for the offset-free MPC and tracking MPC. Dotted black lines represent the intended steady-state targets and disturbance values $(x_{p,s}, d_s, u_s)$ found by solving (8.6) and (8.47). We set $(w_p)_1 \equiv -0.05$, $(w_p)_2 \equiv -0.05$, $(w_p)_3(k) = -0.05H(k - 300)$, $(w_p)_4 \equiv 0$, $(w_p)_5 \equiv 1$, and $r_{sp} \equiv 0.65$ unless otherwise specified.

experiment (Figure 8.2b). The offset-free MPC is able to track the reference and reject the disturbance despite mismatch, this time at the cost of a significant temperature spike around $k = 170$. On the other hand, the tracking MPC fails to bring the temperature above $x_2 = 0.5$, far from the setpoint $r_{\text{sp}} = 0.65$.

In the third experiment, the coolant temperature is an integrating disturbance $(w_{\text{P}})_3^+ = (w_{\text{P}})_3 + (\Delta w_{\text{P}})_3$, $(\Delta w_{\text{P}})_3 \sim \text{N}(0, 10^{-6})$, and we have measurement noise $(w_{\text{P}})_4 \sim \text{N}(0, 10^{-4})$ (Figure 8.4c). As in the corresponding pendulum experiment, offset-free MPC tracks the reference despite the randomly drifting disturbance. Here we are illustrating a behavior that is often desired in chemical process control, where process specifications must be met despite constantly, but slowly varying upstream conditions. We remark that, while the pendulum example is mechanical in nature, it illustrated the same property. The tracking MPC, on the other hand, still cannot handle the plant-model mismatch and fails to bring the temperature up to the setpoint.

In the fourth and final experiment, the setpoint follows an oscillating pattern $r_{\text{sp}}(k) = 0.05 \sin(\frac{2\pi k}{90}) + 0.65$. Tracking MPC again fails bring the temperature up to the setpoint. Offset-free MPC closely follows the setpoint, substantially deviating from it only at the start-up phase and when the coolant temperature disturbance is injected. Again, we note that a precise tracking of this disturbance and reference signal could be accomplished by more general integrator schemes. (c.f., Maeder and Morari (2010) or (Zagrobelny, 2014, Sec. 5.3, 5.4)).

8.7 Conclusions

In this chapter, we presented a nonlinear offset-free MPC design that is robustly stable with respect to setpoint- and target-tracking errors, despite persistent disturbances and plant-model mismatch. Our results are significantly stronger than the standard offset-free sufficient conditions that can be found in the literature. Notably, we do not assume stability of the

closed-loop system to guarantee offset-free performance. The results are illustrated in numerical experiments.

These results form a foundation on which offset-free performance guarantees can be established on a wider class of MPC designs and applications. The results without mismatch (Theorem 8.34) should also extend to the control of plants with parameter drifts. A few limitations of this work, notably the requirement of a Lyapunov function for the estimator (Assumption 8.15), and the necessity of quadratic costs (Assumption 8.8), are also possible areas of future research.

Appendices

8.A Proofs of robust estimation and tracking stability

8.A.1 Proof of Theorem 8.16

Proof of Theorem 8.16. First, note that $c_3 \leq c_2$, as otherwise, this would imply $V_e(k+1) \leq 0$ whenever $\tilde{w}(k) = 0$. We combine the upper bound (8.18a) and bound on the difference (8.18b) to give

$$V_e(k+1) \leq \lambda V_e(k) + c_4 |\tilde{w}(k)|^2$$

where $\lambda := 1 - \frac{c_3}{c_2} \in (0, 1)$. Recursively applying the above inequality gives

$$\begin{aligned} V_e(k) &\leq \lambda^k V_e(0) + \sum_{j=1}^k c_4 \lambda^{j-1} |\tilde{w}(k-j)|^2 \\ &\leq c_2 \lambda^{k+1} |\bar{e}|^2 + \sum_{j=1}^k c_4 \lambda^{j-1} |\tilde{w}(k-j)|^2 \end{aligned}$$

noting that $e(0) = \bar{e}$ because Φ_0 is the identity map. Finally,

$$|e(k)| \leq \sqrt{\frac{V_e(k)}{c_1}} \leq c_{e,1} \lambda_e^k |\bar{e}| + c_{e,2} \sum_{j=1}^{k+1} \lambda_e^{j-1} |\tilde{w}(k-j)|$$

where $c_{e,1} := \sqrt{\frac{c_2}{c_1}}$, $c_{e,2} := \sqrt{\frac{c_4}{c_1}}$, and $\lambda_e := \sqrt{\lambda}$. □

8.A.2 Proof of Theorem 8.23

Proof of Theorem 8.23. Suppose $X \subseteq \Xi$ is RPI for (8.20). Let the functions $V : \Xi \rightarrow \mathbb{R}_{\geq 0}$ and $\alpha_i, \sigma \in \mathcal{K}_\infty, i \in \mathbb{I}_{1:3}$ satisfy (8.24) for all $\xi \in X$ and $\omega \in \Omega_c(\xi)$. Let $(\xi, \omega, \zeta_1, \zeta_2)$ satisfy (8.20) and $\xi(0) \in X$.

Asymptotic case. The proof of this part follows similarly to (Jiang and Wang, 2001, Lem. 3.5) and (Tran et al., 2015, Thm. 1). We start by noting (8.24b) can be rewritten

$$V(F_c(\xi, \omega)) \leq (\text{ID} - \alpha_4)(V(\xi)) + \sigma(\|\omega\|) \quad (8.68)$$

where $\alpha_4 := \alpha_3 \circ \alpha_2^{-1} \in \mathcal{K}_\infty$. Without loss of generality, we can assume $\text{ID} - \alpha_4 \in \mathcal{K}$ (Jiang and Wang, 2001, Lem. B.1). Let $\rho \in \mathcal{K}_\infty$ such that $\text{ID} - \rho \in \mathcal{K}_\infty$.

Let $b := \alpha_4^{-1}(\rho^{-1}(\sigma(\|\omega\|)))$ and $D := \{\xi \in \Xi \mid V(\xi) \leq b\}$. The following intermediate result is required.

Lemma 8.44. *If there exists $k_0 \in \mathbb{I}_{\geq 0}$ such that $\xi(k_0) \in D$, then $\xi(k) \in D$ for all $k \geq k_0$.*

Proof. Suppose $k \geq k_0$ and $\xi(k) \in D$. Then $V(\xi(k)) \leq b$ and by (8.68),

$$\begin{aligned} V(\xi(k+1)) &\leq (\text{ID} - \alpha_4)(V(\xi(k))) + \sigma(\|\omega\|) \\ &\leq (\text{ID} - \alpha_4)(b) + \sigma(\|\omega\|) \\ &= \underbrace{-(\text{ID} - \rho)(\alpha_4(b))}_{\leq 0} + \underbrace{b - \rho(\alpha_4(b)) + \sigma(\|\omega\|)}_{=0} \leq b. \end{aligned}$$

The result follows by induction. □

Next, let $j_0 := \min \{k \in \mathbb{I}_{\geq 0} \mid \xi(k) \in D\}$. The above lemma gives $V(\xi(k)) \leq \gamma(\|\omega\|)$ for all $k \geq j_0$, where $\gamma := \alpha_4^{-1} \circ \rho^{-1} \circ \sigma$. On the other hand, if $k < j_0$, then we have

$\rho(\alpha_4(V(\xi(k)))) > \sigma(\|\boldsymbol{\omega}\|)$ and therefore

$$\begin{aligned} V(\xi(k+1)) - V(\xi(k)) &\leq -\alpha_4(V(\xi(k))) + \sigma(\|\boldsymbol{\omega}\|) \\ &= -\alpha_4(V(\xi(k))) + \rho(\alpha_4(V(\xi(k)))) - \rho(\alpha_4(V(\xi(k)))) + \sigma(\|\boldsymbol{\omega}\|) \\ &\leq -\alpha_4(V(\xi(k))) + \rho(\alpha_4(V(\xi(k)))). \end{aligned}$$

By (Jiang and Wang, 2002, Lem. 4.3), there exists $\beta \in \mathcal{KL}$ such that

$$\alpha_1(|\zeta_1(k)|) \leq V(\xi(k)) \leq \beta(V(\xi(0)), k) \leq \beta(\alpha_2(|\zeta_2(0)|), k).$$

Combining the above inequalities gives

$$|\zeta_1(k)| \leq \max\{\beta_\zeta(|\zeta_2(0)|, k), \gamma_\zeta(\|\boldsymbol{\omega}\|)\} \leq \beta_\zeta(|\zeta_2(0)|, k) + \gamma_\zeta(\|\boldsymbol{\omega}\|)$$

where $\beta_\zeta(s, k) := \alpha_1^{-1}(\beta(\alpha_2(s), k))$ and $\gamma_\zeta := \alpha_1^{-1} \circ \gamma$. Finally, causality lets us drop future terms of ω from the signal norm in the above inequality and simply write (8.22).

Exponential case. Suppose, additionally, that $\alpha_i(\cdot) := a_i(\cdot)^b$, $i \in \mathbb{I}_{1:3}$. Without loss of generality, we can assume $\lambda := 1 - a_3 \in (0, 1)$. Recursively applying (8.24b) gives

$$\begin{aligned} V(\xi(k)) &\leq \lambda^k V(\xi(0)) + \sum_{i=1}^k \lambda^{i-1} \sigma(|\omega(k-i)|) \\ &\leq \lambda^k a_2 |\zeta_2(0)|^b + \frac{\sigma(\|\boldsymbol{\omega}\|_{0:k-1})}{1-\lambda}. \end{aligned}$$

Applying (8.24a), we have

$$|\zeta_1(k)| \leq \left(\frac{a_2}{a_1} \lambda^k |\zeta_2(0)|^b + \frac{\sigma(\|\boldsymbol{\omega}\|_{0:k-1})}{a_1(1-\lambda)} \right)^{1/b}.$$

If $b \geq 1$, the triangle inequality gives

$$|\zeta_1(k)| \leq c_\zeta \lambda_\zeta^k |\zeta_2(0)| + \gamma_\zeta(\|\boldsymbol{\omega}\|_{0:k-1}) \quad (8.69)$$

with $c_\zeta := \left(\frac{a_2}{a_1}\right)^{1/b}$, $\lambda_\zeta := \lambda^{1/b}$, and $\gamma_\zeta(\cdot) := \left(\frac{\sigma(\cdot)}{a_1(1-\lambda)}\right)^{1/b}$. Otherwise, if $b < 1$, then convexity gives (8.69) with $c_\zeta := \frac{1}{2} \left(\frac{2a_2}{a_1}\right)^{1/b}$, $\lambda_\zeta := \lambda^{1/b}$, and $\gamma_\zeta(\cdot) := \frac{1}{2} \left(\frac{2\sigma(\cdot)}{a_1(1-\lambda)}\right)^{1/b}$. \square

8.A.3 Proof of Theorem 8.28

Proof of Theorem 8.28. Throughout, we fix $k \in \mathbb{I}_{\geq 0}$ and drop dependence on k when it is understood from context. Let the trajectories $(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}, \mathbf{u}, \boldsymbol{\omega}, \mathbf{v}, \boldsymbol{\varepsilon}, \zeta)$ satisfy (8.19) and (8.25)–(8.28), $\zeta = G(\hat{\boldsymbol{\xi}})$, and $(\boldsymbol{\xi}(0), \bar{\boldsymbol{\xi}}) \in \mathcal{S}$, where \mathcal{S} is RPI. Suppose Φ_0^ξ is the identity map. Let $a_i, b_i > 0, i \in \mathbb{I}_{1:4}$, $V : \hat{\Xi} \rightarrow \mathbb{R}_{\geq 0}$, $V_\varepsilon : \Xi \times \hat{\Xi} \rightarrow \mathbb{R}_{\geq 0}$, and $\sigma, \sigma_\varepsilon \in \mathcal{K}$ satisfy $\frac{a_4 c_4}{a_3 c_1} < 1$, $\frac{a_4 c_4}{a_3 c_3} < \frac{c_1}{c_1 + c_2}$, and (8.31).

Joint Lyapunov function Combining the fact $|(\boldsymbol{\varepsilon}, \boldsymbol{\varepsilon}^+)|^2 = |\boldsymbol{\varepsilon}|^2 + |\boldsymbol{\varepsilon}^+|^2$ with the inequalities (8.31), we have

$$\begin{aligned} V(\hat{\boldsymbol{\xi}}^+) - V(\hat{\boldsymbol{\xi}}) &\stackrel{(8.31b)}{\leq} -a_3 |\zeta|^2 + a_4 |\boldsymbol{\varepsilon}|^2 + a_4 |\boldsymbol{\varepsilon}^+|^2 + \sigma(|\boldsymbol{\omega}|) \\ &\stackrel{(8.31c)}{\leq} -a_3 |\zeta|^2 + a_4 |\boldsymbol{\varepsilon}|^2 + \frac{a_4}{c_1} V_\varepsilon(\boldsymbol{\xi}^+, \hat{\boldsymbol{\xi}}^+) + \sigma(|\boldsymbol{\omega}|) \\ &\stackrel{(8.31d)}{\leq} -\tilde{a}_3 |\zeta|^2 + a_4 \left(1 - \frac{c_3}{c_1}\right) |\boldsymbol{\varepsilon}|^2 + \frac{a_4}{c_1} V_\varepsilon(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}) + \tilde{\sigma}(|\boldsymbol{\omega}|) \\ &\stackrel{(8.31c)}{\leq} -\tilde{a}_3 |\zeta|^2 + \tilde{a}_4 |\boldsymbol{\varepsilon}|^2 + \tilde{\sigma}(|\boldsymbol{\omega}|) \end{aligned}$$

where $\tilde{a}_3 := a_3 - \frac{a_4 c_4}{c_1}$, $\tilde{a}_4 := a_4 \left(1 + \frac{c_2 - c_3}{c_1}\right)$, and $\tilde{\sigma} := \frac{a_4}{c_1} \sigma_\varepsilon + \sigma \in \mathcal{K}$. Note that $\tilde{a}_3 = a_3 \left(1 - \frac{a_4 c_4}{a_3 c_1}\right) > 0$ by assumption, and $\tilde{a}_4 > 0$ since $c_2 > c_3$.

Let $W(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}) := V(\hat{\boldsymbol{\xi}}) + q V_\varepsilon(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}})$ where $q > 0$. With $b_1 := \min\{a_1, q c_1\}$, we have the

lower bound,

$$b_1|(\zeta, \varepsilon)|^2 = b_1|\zeta|^2 + b_1|\varepsilon|^2 \leq a_1|\zeta|^2 + qc_1|\varepsilon|^2 \leq V(\hat{\xi}) + qV_\varepsilon(\xi, \hat{\xi}) =: W(\xi, \hat{\xi}). \quad (8.70)$$

With $b_2 := \max \{ a_2, qc_2 \}$, we have the upper bound

$$W(\xi, \hat{\xi}) := V(\hat{\xi}) + qV_\varepsilon(\xi, \hat{\xi}) \leq a_2|\zeta|^2 + qc_2|\varepsilon|^2 \leq b_2|\zeta|^2 + b_2|\varepsilon|^2 = b_2|(\zeta, \varepsilon)|^2. \quad (8.71)$$

For the cost decrease, we first note that $\frac{a_4c_4}{a_3c_3} < \frac{c_1}{c_1+c_2}$ implies

$$\tilde{a}_4c_4 = a_4 \left(\frac{c_1 + c_2}{c_1} - \frac{c_3}{c_1} \right) c_4 < a_4 \left(\frac{a_3c_3}{a_4c_4} - \frac{c_3}{c_1} \right) c_4 = a_3c_3 - \frac{a_4c_3c_4}{c_1} = \tilde{a}_3c_3$$

and therefore $\frac{\tilde{a}_4}{c_3} < \frac{\tilde{a}_3}{c_4}$. With $q \in \left(\frac{\tilde{a}_4}{c_3}, \frac{\tilde{a}_3}{c_4} \right)$, we have

$$W(\xi^+, \hat{\xi}^+) \leq V(\hat{\xi}^+) + qV_\varepsilon(\xi^+, \hat{\xi}^+) \leq W(\xi, \hat{\xi}) - b_3|(\zeta, \varepsilon)|^2 + \sigma_W(|\omega|) \quad (8.72)$$

where $b_3 := \min \{ \tilde{a}_3 - qc_4, qc_3 - \tilde{a}_4 \} > 0$ and $\sigma_W := \tilde{\sigma} + q\sigma_\varepsilon \in \mathcal{K}$ by construction.

Robust exponential stability Substituting (8.71) into (8.72) gives

$$W(\xi^+, \hat{\xi}^+) \leq \lambda W(\xi, \hat{\xi}) - b_3|(\zeta, \varepsilon)|^2 + \sigma_W(|\omega|) \quad (8.73)$$

where $\lambda := 1 - \frac{b_3}{b_2}$ and we can assume $\lambda \in (0, 1)$ since

$$b_2 \geq qc_2 > qc_3 > qc_3 - \tilde{a}_4 \geq b_3.$$

Recursively applying (8.73) gives

$$\begin{aligned} W(\xi(k), \hat{\xi}(k)) &\leq \lambda^k W(\xi(0), \hat{\xi}(0)) + \sum_{i=1}^k \lambda^{i-1} \sigma(|\omega(k-i)|) \\ &\leq b_2 \lambda^k |(\zeta(0), \varepsilon(0))|^2 + \sum_{i=1}^k \lambda^{i-1} \sigma(|\omega(k-i)|) \end{aligned}$$

where the second inequality follows from (8.71). Finally, by (8.70) and the triangle inequality, we have

$$|(\zeta(k), e(k))| \leq c_\zeta \lambda_\zeta^k |(\zeta(0), \varepsilon(0))| + \sum_{i=1}^k \gamma_\zeta(|\omega(k-i)|, i)$$

where $c_\zeta := \sqrt{\frac{b_2}{b_1}}$, $\lambda_\zeta := \sqrt{\lambda}$, and $\gamma_\zeta(s, k) := \lambda_\zeta^{k-1} \sqrt{\frac{\sigma(s)}{b_1}}$. □

8.B Proofs of offset-free MPC stability

8.B.1 Proof of Theorem 8.29

We begin by proving Theorem 8.29(a,b).

Proof of Theorem 8.29(a,b). (a)—Suppose $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta \in \mathcal{B}_c$. From the main text, $\tilde{\mathbf{u}}(x, \beta)$ is feasible, so

$$V_N^0(f_c(x, \beta), \beta) \leq V_N(f_c(x, \beta), \tilde{\mathbf{u}}(x, \beta), \beta)$$

and, applying the inequality (8.35), we have

$$V_N^0(f_c(x, \beta), \beta) \leq V_N^0(x, \beta) - \ell(x, \kappa_N(x, \beta), \beta).$$

But

$$\underline{\sigma}(Q) |x - x_s(\beta)|^2 \leq \ell(x, \kappa_N(x, \beta), \beta) \leq V_N^0(x, \beta)$$

so the lower bound (8.34a) and the cost decrease (8.34b) both hold with $a_1 = a_3 = \underline{\sigma}(Q)$.

To establish the upper bound of (8.34a), we first note that since P_f is continuous and positive definite, and \mathcal{B}_c is compact, the maximum $\gamma := \max_{\beta \in \mathcal{B}_c} \bar{\sigma}(P_f(\beta)) > 0$ exists. Then $|x - x_s(\beta)| \leq \varepsilon := \sqrt{\frac{c_f}{\gamma}}$ implies

$$V_f(x, \beta) \leq \bar{\sigma}(P_f(\beta))|x - x_s(\beta)|^2 \leq \gamma|x - x_s(\beta)|^2 \leq c_f$$

and therefore $x \in \mathbb{X}_f(\beta)$. By monotonicity of the value function (Rawlings et al., 2020, Prop. 2.18) we have $V_N^0(x, \beta) \leq V_f(x, \beta)$ whenever $x \in \mathbb{X}_f(\beta)$, and therefore

$$V_N^0(x, \beta) \leq V_f(x, \beta) \leq \gamma|x - x_s(\beta)|^2$$

whenever $|x - x_s(\beta)| \leq \varepsilon$. On the other hand, if $|x - x_s(\beta)| > \varepsilon$, then

$$V_N^0(x, \beta) \leq \rho \leq \frac{\rho}{\varepsilon^2}|x - x_s(\beta)|^2.$$

Finally, we have the upper bound (8.34a) with $a_2 := \max\{\gamma, \frac{\rho}{\varepsilon^2}\}$.

(b)—Let $\beta \in \mathcal{B}$. We already have that $V_N^0(\cdot, \beta)$ is a Lyapunov function (for the system (8.32), on $\mathcal{X}_N^\rho(\beta)$) with respect to $x - x_s(\beta)$, and $f_c(x, \beta) \in \mathcal{X}_N(\beta)$ for all $x \in \mathcal{X}_N^\rho(\beta)$ by recursive feasibility. We can choose any compact set $\mathcal{B}_c \subseteq \mathcal{B}$ containing β to achieve the descent property (8.34b). Then, for each $x \in \mathcal{X}_N^\rho(\beta)$, we have

$$V_N^0(f_c(x, \beta), \beta) \leq V_N^0(x, \beta) - a_1|x - x_s(\beta)|^2 \leq \rho$$

and therefore $f_c(x, \beta) \in \mathcal{X}_N^\rho(\beta)$. In other words, $\mathcal{X}_N^\rho(\beta)$ is positive invariant for the system (8.32a). Finally, ES in $\mathcal{X}_N^\rho(\beta)$ w.r.t. $x - x_s(\beta)$ follows from Theorem 8.23. \square

To prove Theorem 8.29(c,d), we need a few preliminary results.

Proposition 8.45. *Suppose Assumptions 8.4 to 8.8 hold. Let $\rho > 0$ and $\mathcal{B}_c \subseteq \mathcal{B}$ be compact.*

There exist $c_x, c_u > 0$ such that

$$|x^0(j; x, \beta) - x_s(\beta)| \leq c_x |x - x_s(\beta)| \quad (8.74a)$$

$$|u^0(k; x, \beta) - u_s(\beta)| \leq c_u |x - x_s(\beta)| \quad (8.74b)$$

for each $x \in \mathcal{X}_N^\rho(\beta)$, $\beta \in \mathcal{B}_c$, $j \in \mathbb{I}_{1:N}$, and $k \in \mathbb{I}_{1:N-1}$.

Proof. Throughout, we fix $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta \in \mathcal{B}_c$. Unless otherwise specified, the constructed constants and functions are independent of (x, β) . By Theorem 8.29(a), there exists $a_2 > 0$ satisfying the upper bound (8.43a). Since P_f is continuous and positive definite and \mathcal{B}_c is compact, the minimum $\gamma := \min_{\beta \in \mathcal{B}_c} \underline{\sigma}(P_f(\beta))$ exists and is positive. Moreover, since Q, R are positive definite, we have $\underline{\sigma}(Q), \underline{\sigma}(R) > 0$. For each $k \in \mathbb{I}_{0:N-1}$,

$$\begin{aligned} \underline{\sigma}(Q) |x^0(k; x, \beta) - x_s(\beta)|^2 &\leq |x^0(k; x, \beta) - x_s(\beta)|_Q^2 \\ &\leq V_N^0(x, \beta) \leq a_2 |x - x_s(\beta)|^2 \\ \gamma |x^0(N; x, \beta) - x_s(\beta)|^2 &\leq |x^0(N; x, \beta) - x_s(\beta)|_{P_f(\beta)}^2 \\ &\leq V_N^0(x, \beta) \leq a_2 |x - x_s(\beta)|^2 \\ \underline{\sigma}(R) |u^0(k; x, \beta) - u_s(\beta)|^2 &\leq |u^0(k; x, \beta) - u_s(\beta)|_R^2 \\ &\leq V_N^0(x, \beta) \leq a_2 |x - x_s(\beta)|^2. \end{aligned}$$

Thus, (8.74) holds for all $j \in \mathbb{I}_{1:N}$ and $k \in \mathbb{I}_{1:N-1}$ with $c_x := \max \left\{ \sqrt{\frac{a_2}{\underline{\sigma}(Q)}}, \sqrt{\frac{a_2}{\gamma}} \right\}$ and $c_u := \sqrt{\frac{a_2}{\underline{\sigma}(R)}}$. \square

Proposition 8.46. *Suppose Assumptions 8.4 to 8.8 hold. Let $\rho > 0$, $\mathcal{B}_c \subseteq \mathcal{B}$ be compact. There*

exists $\sigma_r \in \mathcal{K}_\infty$ such that

$$|g_c(x, \beta) - r_{\text{sp}}| \leq \sigma_r(|x - x_s(\beta)|) \quad (8.75)$$

for each $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta = (r_{\text{sp}}, z_{\text{sp}}, d) \in \mathcal{B}_c$. Moreover, if g and h are Lipschitz continuous on bounded sets, then (8.75) holds on the same sets with $\sigma_r(\cdot) := c_r(\cdot)$ and some $c_r > 0$.

Proof. By Proposition 7.49, there exists $\tilde{\sigma}_r \in \mathcal{K}_\infty$ such that

$$|g(u, h(x, u, d)) - g(\tilde{u}, h(\tilde{x}, \tilde{u}, \tilde{d}))| \leq \tilde{\sigma}_r(|(x - \tilde{x}, u - \tilde{u}, \beta - \tilde{\beta})|)$$

for all $x, \tilde{x} \in \mathcal{X}_N^\rho$, $u, \tilde{u} \in \mathbb{U}$, and $\beta = (r, z, d), \tilde{\beta} = (\tilde{r}, \tilde{z}, \tilde{d}) \in \mathcal{B}_c$. Fix $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta \in \mathcal{B}_c$. The following constructions are independent of (x, β) unless otherwise specified.

By Proposition 8.45, there exists $c_u > 0$ such that

$$|\kappa_N(x, \beta) - u_s(\beta)| \leq c_u |x - x_s(\beta)|$$

Combining these inequalities gives

$$\begin{aligned} |g_c(x, \beta) - r_{\text{sp}}| &\leq \tilde{\sigma}_r(|(x - x_s(\beta), \kappa_N(x, \beta) - u_s(\beta))|) \\ &\leq \tilde{\sigma}_r((1 + c_u)|x - x_s(\beta)|) \\ &\leq \sigma_r(|x - x_s(\beta)|) \end{aligned}$$

where $\sigma_r(\cdot) := \tilde{\sigma}_r((1 + c_u)(\cdot)) \in \mathcal{K}_\infty$. If we also have that g and h are Lipschitz on bounded sets, then we can take $\sigma_r(\cdot) := c_r(\cdot)$ and $c_r := L_r(1 + c_u) > 0$, where $L_r > 0$ is the Lipschitz constant for $g(u, h(x, u, d))$ over $\mathcal{X}_N^\rho \times \mathbb{U} \times \mathcal{B}_c$. \square

Proof of Theorem 8.29(c,d). Fix $x \in \mathcal{X}_N^\rho(\beta)$ and $\beta \in \mathcal{B}$. Let $\mathcal{B}_c \subseteq \mathcal{B}$ be compact, containing β .

Define $\delta r := g_c(x, \beta) - r_{\text{sp}}$ and $\delta x := x - x_s(\beta)$.

(c)—By Proposition 8.46, there exists $\sigma_r \in \mathcal{K}_\infty$ such that (8.75) holds. Then

$$\alpha_1(|\delta r|) := a_1[\sigma_r^{-1}(|\delta r|)]^2 \leq a_1|\delta x|^2 \leq V_N^0(x, \beta)$$

so $V_N^0(\cdot, \beta)$ is a Lyapunov function on $\mathcal{X}_N^\rho(\beta)$ w.r.t. $(\delta r, \delta x)$, and AS on $\mathcal{X}_N^\rho(\beta)$ w.r.t. $(\delta r, \delta x)$ follows by Theorem 8.23.

(d)—If g and h are Lipschitz continuous on bounded sets, then by Proposition 8.46, we can repeat part (c) with $\alpha_1(\cdot) := a_1 c_r^{-2}(\cdot)^2$ and some $c_r > 0$. Then $V_N^0(\cdot, \beta)$ is an exponential Lyapunov function on $\mathcal{X}_N^\rho(\beta)$ w.r.t. $(\delta r, \delta x)$, and ES on $\mathcal{X}_N^\rho(\beta)$ w.r.t. $(\delta r, \delta x)$ follows by Theorem 8.23. \square

8.B.2 Proof of Proposition 8.33

To establish Proposition 8.33, we require the following result.

Proposition 8.47. *Suppose Assumptions 8.4 to 8.8 and 8.31 hold and let $\rho > 0$. The set*

$$\hat{\mathcal{X}}_N^\rho := \bigcup_{\hat{\beta} \in \hat{\mathcal{B}}_c} \mathcal{X}_N^\rho(\hat{\beta})$$

is compact, where $\hat{\mathcal{B}}_c$ is defined as in Assumption 8.31(i).

Proof. Consider the lifted set

$$\mathcal{F} := \{ (\hat{x}, \mathbf{u}, \hat{\beta}) \in \mathbb{X} \times \mathbb{U}^N \times \hat{\mathcal{B}}_c \mid V_f(\phi(N; \hat{x}, \mathbf{u}, \hat{\beta})) \leq c_f, V_N(\hat{x}, \mathbf{u}, \hat{\beta}) \leq \rho \}.$$

Notice $\hat{\mathcal{X}}_N^\rho$ is equivalent to the projection of \mathcal{F} onto the first coordinate, i.e., $\hat{\mathcal{X}}_N^\rho = P(\mathcal{F})$ where $P(\hat{x}, \mathbf{u}, \hat{\beta}) = \hat{x}$. Since P is continuous, the image $\hat{\mathcal{X}}_N^\rho = P(\mathcal{F})$ is compact whenever \mathcal{F} is compact. Thus, it suffices to show \mathcal{F} is compact.

The set \mathcal{F} is closed because $(\mathbb{X}, \mathbb{U}, \hat{\mathcal{B}}_c)$ are closed, and continuity of (f, x_s, u_s, ℓ, V_f) implies continuity of $V_f(\phi(N; \cdot, \cdot, \cdot))$ and $V_N(\cdot, \cdot, \cdot)$. Next, we show \mathcal{F} is bounded. Since x_s is continuous and $\hat{\mathcal{B}}_c$ is compact, the maximum $\rho_s := \max_{\hat{\beta} \in \hat{\mathcal{B}}_c} |x_s(\hat{\beta})|$ exists and is finite. For each $(\hat{x}, \mathbf{u}, \hat{\beta}) \in \mathcal{F}$, we have $V_N^0(\hat{x}, \hat{\beta}) \leq V_N(\hat{x}, \mathbf{u}, \hat{\beta}) \leq \rho$ by construction. But $V_N^0(\hat{x}, \hat{\beta}) \geq \underline{\sigma}(Q)|\hat{x} - x_s(\hat{\beta})|^2$, so this implies $|\hat{x} - x_s(\hat{\beta})| \leq \sqrt{\frac{\rho}{\underline{\sigma}(Q)}}$ and therefore $|\hat{x}| \leq \sqrt{\frac{\rho}{\underline{\sigma}(Q)}} + \rho_s$. But \mathbf{u} and $\hat{\beta}$ always lie in compact sets, so \mathcal{F} must be bounded. \square

Proof of Proposition 8.33. Let $\hat{\beta} \in \hat{\mathcal{B}}_c$, $\hat{x} \in \mathcal{X}_N^\rho(\hat{\beta})$, and $|\tilde{d}| \leq \delta_0$ such that $\hat{\beta}^+ := \hat{f}_{\beta, c}(\hat{\beta}, \tilde{d}) \in \hat{\mathcal{B}}_c$. For brevity, let

$$\begin{aligned} \bar{x}^+ &:= f_c(\hat{x}, \hat{\beta}), & \bar{x}^+(N) &:= \phi(N; \bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{d}), & \bar{x}(N) &:= x^0(N; \hat{x}, \hat{\beta}), \\ \hat{x}^+ &:= \hat{f}_c(\hat{x}, \hat{\beta}, \tilde{d}), & \hat{x}^+(N) &:= \phi(N; \hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{d}^+). \end{aligned}$$

Recall $\tilde{d} := (e, e^+, \Delta\beta, w, v)$, $e := (e_x, e_d)$, $e^+ := (e_x^+, e_d^+)$, and $\Delta\beta := (\Delta, w_d)$.

From Proposition 8.47, the set $\hat{\mathcal{X}}_N^\rho$ is compact. Since the functions (f, x_s, u_s, P_f) are continuous, so are (V_f, V_N) . By Proposition 7.49, there exist $\sigma_f, \sigma_{V_f}, \sigma_{V_N} \in \mathcal{K}_\infty$ such that

$$|f(x_1, \mathbf{u}_1, \hat{d}_1) - f(x_2, \mathbf{u}_2, \hat{d}_2)| \leq \sigma_f(|(x_1 - x_2, \mathbf{u}_1 - \mathbf{u}_2, \hat{d}_1 - \hat{d}_2)|) \quad (8.76)$$

$$|V_f(\phi(N; x_1, \mathbf{u}_1, \hat{d}_1), \hat{\beta}_1) - V_f(\phi(N; x_2, \mathbf{u}_2, \hat{d}_2), \hat{\beta}_2)| \leq \sigma_{V_f}(|(x_1 - x_2, \mathbf{u}_1 - \mathbf{u}_2, \hat{\beta}_1 - \hat{\beta}_2)|) \quad (8.77)$$

$$|V_N(x_1, \mathbf{u}_1, \hat{\beta}_1) - V_N(x_2, \mathbf{u}_2, \hat{\beta}_2)| \leq \sigma_{V_N}(|(x_1 - x_2, \mathbf{u}_1 - \mathbf{u}_2, \hat{\beta}_1 - \hat{\beta}_2)|) \quad (8.78)$$

for all $x_1 \in \mathbb{X}$, $x_2 \in \hat{\mathcal{X}}_N^\rho$, $u_1, u_2 \in \mathbb{U}$, $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{U}^N$, and $\hat{\beta}_1 = (s_1, \hat{d}_1)$, $\hat{\beta}_2 = (s_2, \hat{d}_2) \in \hat{\mathcal{B}}_c$.

Substituting $x_1 = \hat{x} + e_x$, $x_2 = \hat{x}$, $u_1 = u_2 = \kappa_N(\hat{x}, \hat{\beta})$, $\hat{d}_1 = \hat{d} + e_d$, and $\hat{d}_2 = \hat{d}$ into (8.76), we have $|\hat{x}^+ - \bar{x}^+| \leq \sigma_f(|e|) + |w| + |e_x^+|$. But $|\hat{\beta}^+ - \hat{\beta}| \leq |\Delta\beta| + |e_d| + |e_d^+|$, so

$$|(\hat{x}^+ - \bar{x}^+, \hat{\beta}^+ - \hat{\beta})| \leq \sigma_f(\tilde{d}) + 5|\tilde{d}|. \quad (8.79)$$

Substituting $x_1 = \hat{x}^+$, $x_2 = \hat{f}_c(\hat{x}, \hat{\beta})$, $\mathbf{u}_1 = \mathbf{u}_2 = \tilde{\mathbf{u}}(\hat{x}, \hat{\beta})$, $\hat{\beta}_1 = \hat{\beta}^+$, and $\hat{\beta}_2 = \hat{\beta}$ into (8.77) and (8.78) gives

$$\begin{aligned} |V_f(\hat{x}^+(N), \hat{\beta}^+) - V_f(\bar{x}^+(N), \hat{\beta})| &\leq \sigma_{V_f}(|\hat{x}^+ - \bar{x}^+, \hat{\beta}^+ - \hat{\beta}|) \\ &\leq \tilde{\sigma}_{V_f}(|\tilde{d}|) \end{aligned} \quad (8.80)$$

$$\begin{aligned} |V_N(\hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}^+) - V_N(\bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta})| &\leq \sigma_{V_N}(|\hat{x}^+ - \bar{x}^+, \hat{\beta}^+ - \hat{\beta}|) \\ &\leq \sigma_r(|\tilde{d}|) \end{aligned} \quad (8.81)$$

where $\tilde{\sigma}_{V_f}(\cdot) := \sigma_{V_f}(\sigma_f(\cdot) + 5(\cdot)) \in \mathcal{K}_\infty$, $\sigma_r(\cdot) := \sigma_{V_N}(\sigma_f(\cdot) + 5(\cdot)) \in \mathcal{K}_\infty$, and the second and fourth inequalities follow from (8.79).

Part (a). By definition (8.9) and (8.10), we have $\tilde{\mathbf{u}}(\hat{x}, \hat{\beta}) \in \mathcal{U}_N(\hat{x}^+, \hat{\beta}^+)$ if and only if $V_f(\hat{x}^+(N), \hat{\beta}^+) \leq c_f$. Thus, it suffices to construct $\delta_1 > 0$ (independently of $\hat{\beta}$ and \tilde{d}) for which $\hat{x} \in \mathcal{X}_N(\hat{\beta})$ implies $V_f(\hat{x}^+(N), \hat{\beta}^+) \leq c_f$. Since $\hat{x} \in \mathcal{X}_N(\hat{\beta})$, we already have $V_f(\bar{x}(N), \hat{\beta}) \leq c_f$, and by Assumptions 8.7 and 8.8,

$$\begin{aligned} V_f(\bar{x}^+(N), \hat{\beta}) &\leq V_f(\bar{x}(N), \hat{\beta}) - \ell(\bar{x}(N), \kappa_f(\bar{x}(N), \hat{\beta}), \hat{\beta}) \\ &\leq V_f(\bar{x}(N), \hat{\beta}) - \underline{\sigma}(Q)|\bar{x}(N) - x_s(\hat{\beta})|^2. \end{aligned}$$

Since $\hat{\mathcal{B}}_c$ is compact and $\bar{\sigma}, P_f$ are continuous functions, the maximum

$$a_{f,2} := \max_{\hat{\beta} \in \hat{\mathcal{B}}_c} \bar{\sigma}(P_f(\hat{\beta}))$$

exists and is finite, so

$$\frac{c_f}{2} \leq V_f(\bar{x}(N), \hat{\beta}) \leq a_{f,2}|\bar{x}(N) - x_s(\hat{\beta})|^2.$$

Then $|\bar{x}(N) - x_s(\hat{\beta})| \geq \sqrt{\frac{c_f}{2a_{f,2}}}$ and

$$V_f(\bar{x}^+(N), \hat{\beta}) \leq c_f - \frac{c_f \underline{\sigma}(Q)}{2a_{f,2}}. \quad (8.82)$$

On the other hand, if $V_f(\bar{x}(N), \hat{\beta}) \leq \frac{c_f}{2}$, then we have

$$V_f(\bar{x}^+(N), \hat{\beta}) \leq \frac{c_f}{2}. \quad (8.83)$$

Finally, combining (8.80), (8.82), and (8.83), we have

$$V_f(\hat{x}^+(N), \hat{\beta}^+) \leq c_f - \gamma_f + \tilde{\sigma}_{V_f}(|\tilde{d}|)$$

where $\gamma_f := \min \left\{ \frac{c_f}{2}, \frac{c_f \underline{\sigma}(Q)}{2a_{f,2}} \right\}$ was defined independently of $(\hat{\beta}, \tilde{d})$. Finally, taking $\delta_1 := \min \{ \delta_0, \tilde{\sigma}_{V_f}^{-1}(\gamma_f) \}$, we have $V_f(\hat{x}^+(N), \hat{\beta}^+) \leq c_f$ and $\tilde{\mathbf{u}}(\hat{x}, \hat{\beta}) \in \mathcal{U}_N(\hat{x}^+, \hat{\beta}^+)$.

Part (b). By (8.35), we have

$$\begin{aligned} V_N(\bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}) &\leq V_N^0(\hat{x}, \hat{\beta}) - \ell(\hat{x}, \kappa_N(\hat{x}, \hat{\beta}), \hat{\beta}) \\ &\leq V_N^0(\hat{x}, \hat{\beta}) - \underline{\sigma}(Q) |\bar{x}(N) - x_s(\hat{\beta})|^2. \end{aligned} \quad (8.84)$$

Combining (8.81) and (8.84) gives (8.41) with $a_3 := \underline{\sigma}(Q)$, which is positive since Q is positive definite.

Part (c). The proof of this part follows similarly that of part (a). Since $\hat{x} \in \mathcal{X}_N^\rho(\hat{\beta})$, we have $V_N^0(\hat{x}, \hat{\beta}) \leq \rho$. If $V_N^0(\hat{x}, \hat{\beta}) \geq \frac{\rho}{2}$, then, by Theorem 8.29(a), we have

$$\frac{\rho}{2} \leq V_N^0(\hat{x}, \hat{\beta}) \leq a_2 |\hat{x} - x_s(\hat{\beta})|^2$$

for some $a_2 > 0$. Therefore $|\hat{x} - x_s(\hat{\beta})| \leq \sqrt{\frac{\rho}{2a_2}}$ and

$$V_N(\bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}) \leq \rho - \frac{\rho\sigma(Q)}{2a_2}. \quad (8.85)$$

On the other hand, if $V_N^0(\hat{x}, \hat{\beta}) \leq \frac{\rho}{2}$, then

$$V_N(\bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}) \leq \frac{\rho}{2}. \quad (8.86)$$

Combining (8.41), (8.85), and (8.86) gives

$$V_N(\hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}) \leq \rho - \gamma + \tilde{\sigma}_{V_N}(|\tilde{d}|)$$

where $\gamma := \min \left\{ \frac{\rho}{2}, \frac{\rho\sigma(Q)}{2a_2} \right\}$. But $\tilde{\mathbf{u}}(\hat{x}, \hat{\beta})$ is feasible by part (a), so by optimality, we have

$$V_N^0(\hat{x}^+, \hat{\beta}^+) \leq V_N(\hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}) \leq \rho - \gamma + \tilde{\sigma}_{V_N}(|\tilde{d}|).$$

Thus, as long as $|\tilde{d}| \leq \delta := \min \{ \delta_1, \tilde{\sigma}_{V_N}^{-1}(\gamma) \}$, we have $V_N^0(\hat{x}^+, \hat{\beta}^+) \leq \rho$ and $\hat{x}^+ \in \mathcal{X}_N^\rho(\hat{\beta}^+)$. □

8.B.3 Proof of Proposition 8.35

Proof of Proposition 8.35. Proposition 8.46 gives (8.44a). By Proposition 7.49, there exists $\sigma_g \in \mathcal{K}_\infty$ such that

$$\begin{aligned} & |g(u_1, h(x_1, u_1, d_1) + v_1) - g(u_2, h(x_2, u_2, d_2) + v_2)| \\ & \leq \sigma_g(|(x_1 - x_2, u_1 - u_2, d_1 - d_2, v_1 - v_2)|) \end{aligned} \quad (8.87)$$

for all $x_1, x_2 \in \mathcal{X}_N^\rho(\beta)$, $u_1, u_2 \in \mathbb{U}$, $d_1, d_2 \in \mathbb{D}_c$, $v_1 \in \mathbb{V}_c(x_1, u_1, d_1)$, and $v_2 \in \mathbb{V}_c(x_2, u_2, d_2)$, where

$$\mathbb{D}_c := \{ d \in \mathbb{D} \mid (s_{\text{sp}}, d) \in \mathcal{B}_c \}$$

$$\mathbb{V}_c(x, u, d) := \{ v \in \delta\mathbb{B}^{n_y} \mid h(x, u, d) + v \in \mathbb{Y} \}$$

Fix $\hat{x} \in \mathcal{X}_N^\rho(\hat{\beta})$, $\hat{\beta} = (s_{\text{sp}}, \hat{d}) \in \mathcal{B}_c$, and $\tilde{d} = (e, e^+, \Delta s_{\text{sp}}, \tilde{w}) \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta\mathbb{B}^{n_{\tilde{d}}}$, where $e = (e_x, e_d)$ and $\tilde{w} = (w, w_d, v)$. Substituting $x_1 = \hat{x} + e_x$, $x_2 = \hat{x}$, $u_1 = u_2 = \kappa_N(\hat{x}, \hat{\beta})$, $d_1 = \hat{d} + e_d$, $d_2 = \hat{d}$, $v_1 = v$, and $v_2 = 0$ into (8.87) gives, independently of $(\hat{x}, \hat{\beta}, \tilde{d})$,

$$|\hat{g}_c(\hat{x}, \hat{\beta}, \tilde{d}) - g_c(\hat{x}, \hat{\beta})| \leq \sigma_g(|(e_x, e_d, v)|) \leq \sigma_g(|\tilde{d}|).$$

Then (8.44b) follows by the triangle inequality. Finally, if g and h are Lipschitz continuous on bounded sets, we can take $\sigma_g(\cdot) := c_g(\cdot)$ where $c_g > 0$ is the Lipschitz constant for $g(u, h(x, u, d) + v)$. \square

8.B.4 Proof of Proposition 8.40

To prove Proposition 8.40, we derive a bound on $|\tilde{w}|$.

Proposition 8.48. *Suppose Assumptions 8.4 to 8.6, 8.36, and 8.38 hold. For any compact $X \subseteq \mathbb{X}$ and $\mathcal{A}_c \subseteq \mathbb{R}^{n_r} \times \overline{\mathbb{Z}}_y \times \mathbb{W}$ such that $(s_{\text{sp}}, w_p) \in \mathcal{A}_c$ implies $(s_{\text{sp}}, 0) \in \mathcal{A}_c$, there exist functions $\sigma_w, \sigma_\alpha \in \mathcal{K}_\infty$ for which*

$$|\tilde{w}| \leq \sigma_w(|w_p|)|z - z_s(\beta)| + \sigma_\alpha(|\Delta\alpha|) \quad (8.88)$$

for all $z \in X \times \mathbb{U}$ and $\alpha = (s_{\text{sp}}, w_p)$, $\alpha^+ \in \mathcal{A}_c$, where $\beta := (s_{\text{sp}}, d_s(\alpha))$, $\tilde{w} := (w, w_d, v)$, $\Delta\alpha := \alpha^+ - \alpha$, and (8.50).

Proof. For ease of notation, we let $z = (x, u) \in X \times \mathbb{U}$, $\alpha = (s_{\text{sp}}, w_{\text{p}}) \in \mathcal{A}_c$, $\beta := (s_{\text{sp}}, d_s(\alpha))$, $\tilde{w} := (w, w_d, v)$, and

$$\Delta\tilde{w}(x, u, \alpha) := \begin{bmatrix} f_{\text{p}}(x + \Delta x_s(\alpha), u, w_{\text{p}}) - f(x, u, \hat{d}_s(\alpha)) - \Delta x_s(\alpha) \\ h_{\text{p}}(x + \Delta x_s(\alpha), u, w_{\text{p}}) - h(x, u, \hat{d}_s(\alpha)) \end{bmatrix}$$

throughout. We also note that, by definition of the SSTP (8.6) and the nominal model assumption (8.3), we have

$$\Delta\tilde{w}(z_s(\beta), \alpha) = 0, \quad \partial_z \Delta\tilde{w}(z, s_{\text{sp}}, 0) = 0. \quad (8.89)$$

Assume all functions continuously differentiable on $\mathbb{X} \times \mathbb{U}$ have been extended continuously differentiable functions on all of \mathbb{R}^{n+n_u} using appropriately defined partitions of unity (cf. (Lee, 2012, Lem. 2.26)).

Let Z_c denote the convex hull of $X \times \mathbb{U}$. For each $i \in \mathbb{I}_{1:n+n_y}$, $\partial_z \Delta\tilde{w}_i$ is continuous, and by Proposition 7.49, there exists $\sigma_i \in \mathcal{K}_{\infty}$ such that

$$|\partial_z \Delta\tilde{w}_i(z_1, \alpha_1) - \partial_z \Delta\tilde{w}_i(z_2, \alpha_2)| \leq \sigma_i(|(z_1 - z_2, \alpha_1 - \alpha_2)|)$$

for all $z_1, z_2 \in Z_c$ and $\alpha_1, \alpha_2 \in \mathcal{A}_c$. Substituting $z_1 = z_2 = z$, $\alpha_1 = \alpha$, and $\alpha_2 = (s_{\text{sp}}, 0)$ into the above inequality, we have

$$|\partial_z \Delta\tilde{w}_i(z, \alpha)| = |\partial_z \Delta\tilde{w}_i(z, \alpha) - \partial_z \Delta\tilde{w}_i(z, s_{\text{sp}}, 0)| \leq \sigma_i(|w_{\text{p}}|) \quad (8.90)$$

where the equality follows by (8.89). By Taylor's theorem (Apostol, 1974, Thm. 12.14), for each $i \in \mathbb{I}_{1:n+n_y}$, there exist $z_i(z, \alpha) \in Z_c$ and $t_i(z, \alpha) \in (0, 1)$ such that

$$\Delta\tilde{w}_i(z, \alpha) = \partial_z \Delta\tilde{w}_i(\tilde{z}_i(z, \alpha), \alpha)(z - z_s(\beta)) \quad (8.91)$$

where $\tilde{z}_i(z, \alpha) := t_i(z, \alpha)z_s(\beta) + (1 - t_i(z, \alpha))z_i(z, \alpha) \in Z_c$ by convexity of Z_c , and the zero-order term drops by (8.89). Combining (8.90) and (8.91),

$$|\Delta\tilde{w}(z, \alpha)| \leq \sum_{i=1}^{n+n_y} |\Delta\tilde{w}_i(z, \alpha)| \leq \sum_{i=1}^{n+n_y} \sigma_i(|w_P|)|z - z_s(\beta)| = \sigma_w(|w_P|)|z - z_s(\beta)| \quad (8.92)$$

where $\sigma_w := \sum_{i=1}^{n+n_y} \sigma_i$. By Proposition 7.49, since $x_{P,s}, x_s, d_s$ are continuous, there exist $\sigma_x, \sigma_d \in \mathcal{K}_\infty$ such that

$$|\Delta x_s(\alpha_1) - \Delta x_s(\alpha_2)| \leq \sigma_x(|\alpha_1 - \alpha_2|) \quad (8.93a)$$

$$|d_s(\alpha_1) - d_s(\alpha_2)| \leq \sigma_d(|\alpha_1 - \alpha_2|) \quad (8.93b)$$

for all $\alpha_1, \alpha_2 \in \mathcal{A}_c$. Finally, using (8.92) and (8.93) with $\alpha_1 = \alpha$ and $\alpha_2 = \alpha^+$ gives

$$\begin{aligned} |\tilde{w}| &\leq |\Delta\tilde{w}(z, \alpha)| + |\Delta x_s(\alpha^+) - \Delta x_s(\alpha)| + |d_s(\alpha^+) - d_s(\alpha)| \\ &\leq \sigma_w(|w_P|)|z - z_s(\beta)| + \sigma_\alpha(|\Delta\alpha|) \end{aligned}$$

with $\sigma_\alpha := \sigma_x + \sigma_d \in \mathcal{K}_\infty$. □

Proof of Proposition 8.40. With $\delta_w \in (0, \sigma_w^{-1}(\sqrt{\frac{c_3}{4c_4L_s^2}}))$, we can combine (8.18b), (8.74b), and (8.88) (from Assumption 8.15 and Propositions 8.45 and 8.48, respectively) and the identity $(a + b)^2 \leq 2a^2 + 2b^2$ to give

$$\begin{aligned} |\tilde{w}|^2 &\leq [\sigma_w(|w_P|)|z - z_s(\beta)| + \sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq 2[\sigma_w(|w_P|)]^2|z - z_s(\beta)|^2 + 2[\sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq 2[\sigma_w(|w_P|)]^2[(1 + c_u)|\hat{x} - x_s(\hat{\beta})| + L_s|e|]^2 + 2[\sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq 4[\sigma_w(|w_P|)]^2(1 + c_u)^2|\hat{x} - x_s(\hat{\beta})|^2 + 4[\sigma_w(|w_P|)]^2L_s^2|e|^2 + 2[\sigma_\alpha(|\Delta\alpha|)]^2 \end{aligned}$$

and therefore (8.52), where $\hat{c}_3 := c_3 - 4c_4[\sigma_w(\delta_w)]^2 L_s^2 > 0$, $\hat{\sigma}_w(\cdot) := 4c_4[\sigma_w(\cdot)]^2(1 + c_u)^2$, $\hat{\sigma}_\alpha(\cdot) := 2c_4[\sigma_\alpha(\cdot)]^2$, and $L_s > 0$ is the Lipschitz constant for z_s . \square

8.B.5 Proof of Proposition 8.41

To establish Proposition 8.41, we require two preliminary results.

Proposition 8.49. *Suppose Assumptions 8.4 to 8.8, 8.36, and 8.38 hold. Let $\rho, \delta_w > 0$. There exist $\tilde{c}_e > 0$ and $\tilde{\sigma}_w, \tilde{\sigma}_\alpha \in \mathcal{K}_\infty$ such that*

$$|\tilde{d}|^2 \leq \tilde{c}_e|(e, e^+)|^2 + \tilde{\sigma}_w(|w_P|)|\hat{x} - x_s(\hat{\beta})|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|) \quad (8.94)$$

so long as $\alpha = (s_{\text{sp}}, w_P) \in \mathcal{A}_c(\delta_w)$, $\Delta\alpha = (\Delta s_{\text{sp}}, \Delta w_P) \in \mathbb{A}_c(\alpha, \delta_w)$, $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$, $\tilde{d} = (e, e^+, \Delta s_{\text{sp}}, \tilde{w}) \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta_0 \mathbb{B}^{n_{\tilde{d}}}$, and $\hat{\beta} = (s_{\text{sp}}, \hat{d})$, given (8.17) and (8.50).

Proof. From Propositions 8.45 and 8.48 and (Rawlings and Ji, 2012, Eq. (1)),

$$\begin{aligned} |\tilde{w}|^2 &\leq [\sigma_w(|w_P|)|z - z_s(\beta)| + \sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq [\sigma_w(|w_P|)|z - z_s(\hat{\beta})| + L_s \sigma_w(|w_P|)|e| + \sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq [\sigma_w(|w_P|)|x - x_s(\hat{\beta})| + \sigma_w(|w_P|)|u - u_s(\hat{\beta})| + L_s \sigma_w(|w_P|)|e| + \sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq [(1 + c_u)\sigma_w(|w_P|)|\hat{x} - x_s(\hat{\beta})| + (L_s + 1)\sigma_w(|w_P|)|e| + \sigma_\alpha(|\Delta\alpha|)]^2 \\ &\leq 9(1 + c_u)^2[\sigma_w(|w_P|)]^2|\hat{x} - x_s(\beta)|^2 + 9(L_s + 1)^2[\sigma_w(|w_P|)]^2|e|^2 + 9[\sigma_\alpha(|\Delta\alpha|)]^2 \end{aligned}$$

where $L_s > 0$ is the Lipschitz constant for z_s and $c_u > 0$ and $\sigma_w, \sigma_\alpha \in \mathcal{K}_\infty$ satisfy (8.74b) and

(8.88). Therefore

$$\begin{aligned} |\tilde{d}|^2 &= |(e, e^+)|^2 + |\Delta s_{\text{sp}}|^2 + |\tilde{w}|^2 \\ &\leq 9(1 + c_u)^2 (\sigma_w(|w_{\text{P}}|))^2 |\hat{x} - x_s(\beta)|^2 \\ &\quad + (1 + 9(L_s + 1)^2 (\sigma_w(\delta_w))^2) |(e, e^+)|^2 + |\Delta\alpha|^2 + 9\sigma_\alpha (|\Delta\alpha|)^2 \end{aligned}$$

so (8.94) holds with $\tilde{c}_e := 1 + 9(L_s + 1)^2 [\sigma_w(\delta_w)]^2 > 0$ and $\tilde{\sigma}_w := 9(1 + c_u)^2 \sigma_w^2$, $\sigma_\alpha := \text{ID}^2 + 9\sigma_\alpha \in \mathcal{K}_\infty$. \square

Proposition 8.50. *Suppose Assumptions 8.4 to 8.8, 8.36, and 8.38 hold and let $\rho > 0$. There exist $a_{V_N,1} \in (0, \underline{\sigma}(Q))$ and $a_{V_N,2}, \delta > 0$ and $\sigma_{V_N} \in \mathcal{K}_\infty$ such that*

$$|V_N(\hat{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta}^+) - V_N(\bar{x}^+, \tilde{\mathbf{u}}(\hat{x}, \hat{\beta}), \hat{\beta})| \leq a_{V_N,1} |\hat{x} - x_s(\hat{\beta})|^2 + a_{V_N,2} |\tilde{d}|^2 \quad (8.95)$$

for all $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$ and $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta \mathbb{B}^{n_{\tilde{d}}}$, where $\bar{x}^+ := f_c(\hat{x}, \hat{\beta})$, $\hat{x}^+ := \hat{f}_c(\hat{x}, \hat{\beta}, \tilde{d})$, and $\hat{\beta}^+ := \hat{f}_{\beta,c}(\hat{\beta}, \tilde{d})$.

Proof. By continuity of P_f , there exists $\sigma_{P_f} \in \mathcal{K}_\infty$ such that

$$\|P_f(\beta_1) - P_f(\beta_2)\| \leq \sigma_{P_f} (|\beta_1 - \beta_2|) \quad (8.96)$$

for all $\beta_1, \beta_2 \in \hat{\mathcal{B}}_c$. Moreover, since $\hat{\mathcal{B}}_c$ is compact and P_f is continuous and positive definite, the maximum $\gamma := \max_{\hat{\beta} \in \hat{\mathcal{B}}_c} \bar{\sigma}(P_f(\hat{\beta}))$ exists and is finite and the minimum $\gamma_0 := \max_{\hat{\beta} \in \hat{\mathcal{B}}_c} \underline{\sigma}(P_f(\hat{\beta}))$ exists and is positive. Let $L_s > 0$ denote the Lipschitz constant for z_s on $\hat{\mathcal{B}}_c$. Throughout, we let $(\hat{x}, \hat{\beta}) \in \hat{\mathcal{S}}_N^\rho$, $\delta \hat{x} := \hat{x} - x_s(\hat{\beta})$, $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta})$, $\bar{x}^+ := f_c(\hat{x}, \hat{\beta})$, $\hat{x}^+ := \hat{f}_c(\hat{x}, \hat{\beta}, \tilde{d})$, $\hat{\beta}^+ := \hat{f}_{\beta,c}(\hat{\beta}, \tilde{d})$, $\tilde{\mathbf{u}} := \tilde{\mathbf{u}}(\hat{x}, \hat{\beta})$, $\bar{x}^+(k) := \phi(k; \bar{x}^+, \tilde{\mathbf{u}}, \hat{\beta})$, and $\hat{x}^+(k) := \phi(k; \hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta}^+)$.

By Assumption 8.38, we have

$$|\bar{x}^+ - \hat{x}^+| \leq L_f |e| + |w| + |e_x^+| \leq L'_f |\tilde{d}| \quad (8.97)$$

where $L_f > 0$ is the Lipschitz constant for f and $L'_f := L_f + 2$, by Assumption 8.36(b), we have

$$|z_s(\hat{\beta}^+) - z_s(\hat{\beta})| \leq L_s |\hat{\beta}^+ - \hat{\beta}| \leq L_s (|\Delta\beta| + |e_d| + |e_d^+|) \leq 3L_s |\tilde{d}| \quad (8.98)$$

and by Proposition 8.45, we have $c_x, c_u > 0$ such that

$$|\bar{x}^+(j) - x_s(\hat{\beta})| \leq c_x |\delta\hat{x}| \quad (8.99)$$

$$|\tilde{u}(k) - u_s(\hat{\beta})| \leq c_u |\delta\hat{x}| \quad (8.100)$$

for each $j \in \mathbb{I}_{0:N-1}$ and $k \in \mathbb{I}_{0:N-2}$.

By Assumptions 8.7 and 8.8, we have

$$\begin{aligned} \gamma_0 |\bar{x}^+(N) - x_s(\hat{\beta})|^2 &\leq V_f(\bar{x}^+(N-1), \hat{\beta}) \\ &\leq V_f(\bar{x}^+(N-1), \hat{\beta}) - \underline{\sigma}(Q) |\bar{x}^+(N-1) - x_s(\hat{\beta})|^2 \\ &\leq [\gamma - \underline{\sigma}(Q)] |\bar{x}^+(N-1) - x_s(\hat{\beta})|^2 \\ &\stackrel{(8.99)}{\leq} [\gamma - \underline{\sigma}(Q)] c_x^2 |\delta\hat{x}|^2. \end{aligned}$$

Therefore

$$|\bar{x}^+(N) - x_s(\hat{\beta})| \leq \gamma_f c_x |\delta\hat{x}| \quad (8.101a)$$

where $\gamma_f := \sqrt{\frac{\gamma - \underline{\sigma}(Q)}{\gamma_0}}$. Similarly, using the fact that $V_f(\bar{x}^+(N), \hat{\beta}) \geq 0$, we have

$$\begin{aligned} \underline{\sigma}(R)|\tilde{u}(N-1) - u_s(\hat{\beta})|^2 &\leq V_f(\bar{x}^+(N-1), \hat{\beta}) - \underline{\sigma}(Q)|\bar{x}^+(N-1) - x_s(\hat{\beta})|^2 \\ &\leq (\gamma - \underline{\sigma}(Q))|\bar{x}^+(N-1) - x_s(\hat{\beta})|^2 \\ &\stackrel{(8.99)}{\leq} (\gamma - \underline{\sigma}(Q))c_x^2|\delta\hat{x}|^2 \end{aligned}$$

and therefore

$$|\tilde{u}(N-1) - u_s(\hat{\beta})| \leq c_{u,f}|\delta\hat{x}| \quad (8.101b)$$

with $c_{u,f} := c_x \sqrt{\frac{\gamma - \underline{\sigma}(Q)}{\underline{\sigma}(R)}}$.

Due to continuous differentiability of f , we have

$$\begin{aligned} |\hat{x}^+(k) - \bar{x}^+(k)| &= |f(\hat{x}^+(k-1), \tilde{u}(k), \hat{d}^+) - f(\bar{x}^+(k-1), \tilde{u}(k), \hat{d})| \\ &\leq L_f|\hat{x}^+(k-1) - \bar{x}^+(k-1)| + L_f|\hat{d}^+ - \hat{d}| \end{aligned}$$

where $L_f > 0$ is the Lipschitz constant for f . Applying this inequality recursively, for all $k \in \mathbb{I}_{0:N}$, we have

$$|\hat{x}^+(k) - \bar{x}^+(k)| \leq L_f^k|\hat{x}^+ - \bar{x}^+| + L_f(k)|\hat{d}^+ - \hat{d}| \leq L'_f(k)|\tilde{d}| \quad (8.102)$$

where $L_f(k) := \sum_{i=1}^k L_f^i$ and $L'_f(k) := L_f^k L_f + 3L_f(k)$, and we have used (8.97) and the fact that $|\hat{d}^+ - \hat{d}| \leq |w_d| + |e_d| + |e_d^+| \leq 3|\tilde{d}|$. Moreover,

$$|\hat{x}^+(k) - x_s(\hat{\beta})| \stackrel{(8.99), (8.102)}{\leq} c_x|\delta\hat{x}| + L'_f(k)|\tilde{d}| \quad (8.103)$$

and

$$|\hat{x}^+(N) - x_s(\hat{\beta})| \stackrel{(8.101), (8.102)}{\leq} c_x\gamma_f|\delta\hat{x}| + L'_f(N)|\tilde{d}|. \quad (8.104)$$

Using the inequalities, $\|\xi\|_{M_1}^2 - \|\xi\|_{M_2}^2 \leq \|M_1 - M_2\| \|\xi\|^2$, (8.96), and $|\hat{\beta}^+ - \hat{\beta}| \leq |\Delta\beta| + |e_d| + |e_d^+| \leq 3|\tilde{d}|$, we have

$$V_f(\hat{x}^+(N), \hat{\beta}^+) \leq |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta})}^2 + \sigma_{P_f}(3|\tilde{d}|) |\hat{x}^+(N) - x_s(\hat{\beta}^+)|^2.$$

Using the identity $|\xi_1 + \xi_2|^2 \leq 2|\xi_1|^2 + 2|\xi_2|^2$, we have

$$V_f(\hat{x}^+(N), \hat{\beta}^+) \leq |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta})}^2 + \sigma_{P_f,x}(|\tilde{d}|) |\delta\hat{x}|^2 + \sigma_{P_f,d}(|\tilde{d}|) |\tilde{d}|^2. \quad (8.105)$$

where $\sigma_{P_f,x}(\cdot) := 2c_x^2\gamma_f^2\sigma_{P_f}(3(\cdot))$ and $\sigma_{P_f,d}(\cdot) := 2(L'_f(N))^2\sigma_{P_f}(3(\cdot))$.

For the remainder of this part, we let $\lambda > 0$ (to be defined) and use the identity $2ab \leq \lambda a^2 + \lambda^{-1}b^2$. Expanding quadratics and using the identities (8.98)–(8.100), we have

$$\begin{aligned} & \left| |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta})}^2 - |\hat{x}^+(N) - x_s(\hat{\beta})|_{P_f(\hat{\beta})}^2 \right| \\ & \leq 6\gamma L_s |\hat{x}^+(N) - x_s(\hat{\beta})| |\tilde{d}| + 9\gamma L_s^2 |\tilde{d}|^2 \\ & \leq 6\gamma L_s c_x \gamma_f |\delta\hat{x}| |\tilde{d}| + (6\gamma L_s L'_f(N) + 9\gamma L_s^2) |\tilde{d}|^2 \\ & \leq 3\gamma \lambda L_s c_x \gamma_f |\delta\hat{x}|^2 + (6\gamma L_s L'_f(N) + 9\gamma L_s^2 + 3\lambda^{-1}\gamma L_s c_x \gamma_f) |\tilde{d}|^2 \\ & \leq \lambda \hat{L}_1(N) |\delta\hat{x}|^2 + \hat{L}_2(N, \lambda) |\tilde{d}|^2 \end{aligned} \quad (8.106)$$

where $\hat{L}_1(N) := 3\gamma L_s c_x \gamma_f$ and $\hat{L}_2(N, \lambda) := 6\gamma L_s L'_f(N) + 9\gamma L_s^2 + 3\lambda^{-1}\gamma L_s c_x \gamma_f$. Similarly,

for each $k \in \mathbb{I}_{0:N-1}$,

$$\begin{aligned}
& \left| |\hat{x}^+(k) - x_s(\hat{\beta}^+)|_Q^2 - |\hat{x}^+(k) - x_s(\hat{\beta})|_Q^2 \right| \\
& \leq 6\sigma(Q)L_s|\hat{x}^+(k) - x_s(\hat{\beta})|\|\tilde{d}\| + 9\sigma(Q)L_s^2\|\tilde{d}\|^2 \\
& \leq 6\sigma(Q)L_sc_x|\delta\hat{x}|\|\tilde{d}\| + (6\sigma(Q)L_sL'_f(k) + 9\sigma(Q)L_s^2)\|\tilde{d}\|^2 \\
& \leq 3\lambda\sigma(Q)L_sc_x|\delta\hat{x}|^2 + (6\sigma(Q)L_sL'_f(k) + 9\sigma(Q)L_s^2 + 3\lambda^{-1}\gamma L_sc_x)\|\tilde{d}\|^2 \\
& \leq \lambda\hat{L}_1(k)|\delta\hat{x}|^2 + \hat{L}_2(k, \lambda)\|\tilde{d}\|^2
\end{aligned} \tag{8.107}$$

where $\hat{L}_1(k) := 3\sigma(Q)L_sc_x$ and $\hat{L}_2(k, \lambda) := 6\sigma(Q)L_sL'_f(k) + 9\sigma(Q)L_s^2 + 3\lambda^{-1}\gamma L_sc_x$, and

$$\begin{aligned}
& \left| |\tilde{u}(k) - u_s(\hat{\beta}^+)|_R^2 - |\tilde{u}(k) - u_s(\hat{\beta})|_R^2 \right| \\
& \leq 6\sigma(R)L_s|\tilde{u}(k) - u_s(\hat{\beta})|\|\tilde{d}\| + 9\sigma(R)L_s^2\|\tilde{d}\|^2 \\
& \leq 6\sigma(R)L_sc_u(k)|\delta\hat{x}|\|\tilde{d}\| + 9\sigma(R)L_s^2\|\tilde{d}\|^2 \\
& \leq 3\lambda\sigma(R)L_sc_u(k)|\delta\hat{x}|^2 + (9\sigma(R)L_s^2 + 3\lambda^{-1}\sigma(R)L_sc_u(k))\|\tilde{d}\|^2 \\
& \leq \lambda\tilde{L}_1(k)|\delta\hat{x}|^2 + \tilde{L}_2(k, \lambda)\|\tilde{d}\|^2
\end{aligned} \tag{8.108}$$

where $\tilde{L}_1(k) := 3\sigma(R)L_sc_u(k)$, $\tilde{L}_2(k, \lambda) := 9\sigma(R)L_s^2 + 3\lambda^{-1}\sigma(R)L_sc_u(k)$, and $c_u(k) = c_u$ if $k \in \mathbb{I}_{0:N-2}$ and $c_u(N-1) = c_{u,f}$.

For the uniform $\hat{\beta}$ bound, we have

$$\begin{aligned}
& |V_N(\hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta}) - V_N(\bar{x}^+, \tilde{\mathbf{u}}, \hat{\beta})| \\
& \leq \sum_{k=0}^{N-1} 2\bar{\sigma}(Q)|\hat{x}^+(k) - \bar{x}^+(k)||\bar{x}^+(k) - x_s(\hat{\beta})| + \bar{\sigma}(Q)|\hat{x}^+(k) - \bar{x}^+(k)|^2 \\
& \quad + 2\gamma|\hat{x}^+(N) - \bar{x}^+(N)||\bar{x}^+(N) - x_s(\hat{\beta})| + \gamma|\hat{x}^+(N) - \bar{x}^+(N)|^2 \\
& \leq \sum_{k=0}^{N-1} 2\bar{\sigma}(Q)c_x L'_f(k)|\delta\hat{x}||\tilde{d}| + \bar{\sigma}(Q)(L'_f(k))^2|\tilde{d}|^2 \\
& \quad + 2\gamma c_x \gamma_f L'_f(N)|\delta\hat{x}||\tilde{d}| + \gamma(L'_f(N))^2|\tilde{d}|^2 \\
& \leq \sum_{k=0}^{N-1} \lambda\bar{\sigma}(Q)c_x L'_f(k)|\delta\hat{x}|^2 + (\bar{\sigma}(Q)(L'_f(k))^2 + \lambda^{-1}\bar{\sigma}(Q)c_x L'_f(k))|\tilde{d}|^2 \\
& \quad + \lambda\gamma c_x \gamma_f L'_f(N)|\delta\hat{x}|^2 + (\gamma(L'_f(N))^2 + \lambda^{-1}\gamma c_x \gamma_f L'_f(N))|\tilde{d}|^2 \\
& \leq \sum_{k=0}^{N-1} \lambda L_1(k)|\delta\hat{x}|^2 + L_2(k, \lambda)|\tilde{d}|^2 + \lambda L_1(N)|\delta\hat{x}|^2 + L_2(N, \lambda)|\tilde{d}|^2
\end{aligned}$$

where $L_1(k) := \bar{\sigma}(Q)c_x L'_f(k)$ and $L_2(k, \lambda) := \bar{\sigma}(Q)(L'_f(k))^2 + \lambda^{-1}\bar{\sigma}(Q)c_x L'_f(k)$ for each $k \in \mathbb{I}_{0:N-1}$, and $L_1(N) := \gamma c_x \gamma_f L'_f(N)$ and $L_2(N, \lambda) := \gamma(L'_f(N))^2 + \lambda^{-1}\gamma c_x \gamma_f L'_f(N)$.

Finally, we compile the above results,

$$\begin{aligned}
& \left| |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta}^+)}^2 - |\bar{x}^+(N) - x_s(\hat{\beta})|_{P_f(\hat{\beta})}^2 \right| \\
& \stackrel{(8.105)}{\leq} \left| |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta})}^2 - |\bar{x}^+(N) - x_s(\hat{\beta})|_{P_f(\hat{\beta})}^2 \right| + \sigma_{P_f, x}(|\tilde{d}|)|\delta\hat{x}|^2 + \sigma_{P_f, d}(|\tilde{d}|)|\tilde{d}|^2 \\
& \stackrel{(8.106)}{\leq} \left| |\hat{x}^+(N) - x_s(\hat{\beta}^+)|_{P_f(\hat{\beta})}^2 - |\bar{x}^+(N) - x_s(\hat{\beta})|_{P_f(\hat{\beta})}^2 \right| \\
& \quad + (\sigma_{P_f, x}(|\tilde{d}|) + \lambda\hat{L}_1(N))|\delta\hat{x}|^2 + (\sigma_{P_f, d}(|\tilde{d}|) + \hat{L}_2(N, \lambda))|\tilde{d}|^2 \tag{8.109}
\end{aligned}$$

and therefore

$$\begin{aligned}
& |V_N(\hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta}^+) - V_N(\hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta})| \\
& \stackrel{(8.107)-(8.109)}{\leq} \sum_{k=0}^{N-1} \lambda(\hat{L}_1(k) + \tilde{L}_1(k)) |\delta \hat{x}|^2 + (\hat{L}_2(k, \lambda) + \tilde{L}_2(k, \lambda)) |\tilde{d}|^2 \\
& \quad + (\sigma_{P_f, x}(|\tilde{d}|) + \lambda \hat{L}_1(N)) |\delta \hat{x}|^2 + (\sigma_{P_f, d}(|\tilde{d}|) + \hat{L}_2(N, \lambda)) |\tilde{d}|^2
\end{aligned}$$

Finally (8.95) holds so long as $|\tilde{d}| \leq \delta$, with

$$\begin{aligned}
a_{V_N, 1} & := \sigma_{P_f, x}(\delta) + \lambda \left(L_1(N) + \hat{L}_1(N) + \sum_{k=0}^{N-1} \bar{L}_1(k) \right) \\
a_{V_N, 2} & := \sigma_{P_f, d}(\delta) + L_2(N, \lambda) + \hat{L}_2(N, \lambda) + \sum_{k=0}^{N-1} \bar{L}_2(k, \lambda)
\end{aligned}$$

where $\bar{L}_1(k) := L_1(k) + \hat{L}_1(k) + \tilde{L}_1(k)$ and $\bar{L}_2(k, \lambda) := L_2(k, \lambda) + \hat{L}_2(k, \lambda) + \tilde{L}_2(k, \lambda)$. To ensure $a_{V_N, 1} < \underline{\sigma}(Q)$, we can simply choose $\lambda \in \left(0, \frac{\underline{\sigma}(Q) - \sigma_{P_f, x}(\delta)}{L_1(N) + \hat{L}_1(N) + \sum_{k=0}^{N-1} \bar{L}_1(k)} \right)$ and $\delta \in (0, \sigma_{P_f, x}^{-1}(\underline{\sigma}(Q)))$. \square

Proof of Proposition 8.41. For convenience, we define $\tilde{\mathbf{u}} := \tilde{\mathbf{u}}(\hat{x}, \hat{\beta})$. From Propositions 8.49 and 8.50, we have $a_{V_N, 1} \in (0, \underline{\sigma}(Q))$, $a_{V_N, 2}, \tilde{c}_e, \delta, \delta_w > 0$, and $\tilde{\sigma}_w, \tilde{\sigma}_\alpha \in \mathcal{K}_\infty$ such that

$$|V_N(\hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta}^+) - V_N(\bar{x}^+, \tilde{\mathbf{u}}, \hat{\beta})| \leq (a_{V_N, 1} + \tilde{\sigma}_w(|w_P|)) |\delta \hat{x}|^2 + a_{V_N, 2} c_e |(e, e^+)|^2 + \tilde{\sigma}_\alpha(|\Delta \alpha|)$$

so long as $\tilde{d} \in \tilde{\mathbb{D}}_c(\hat{x}, \hat{\beta}) \cap \delta \mathbb{B}^{n_{\tilde{d}}}$, $\alpha \in \mathcal{A}_c(\delta_w)$, and $\Delta \alpha \in \mathbb{A}_c(\alpha, \delta_w)$. Without loss of generality, assume $\delta_w < \tilde{\sigma}_w^{-1}(\underline{\sigma}(Q) - a_{V_N, 1})$. By Proposition 8.33, we can choose $\delta > 0$ such that $\tilde{\mathbf{u}} \in$

$\mathcal{U}_N(\hat{x}^+, \hat{\beta}^+)$, so

$$\begin{aligned} V_N^0(\hat{x}^+, \hat{\beta}^+) &\leq V_N(\hat{x}^+, \tilde{\mathbf{u}}, \hat{\beta}^+) \\ &\leq V_N(\bar{x}^+, \tilde{\mathbf{u}}, \hat{\beta}) + (a_{V_N,1} + \tilde{\sigma}_w(\delta_w))|\delta\hat{x}|^2 + a_{V_N,2}c_e|(e, e^+)|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|) \\ &\leq V_N^0(\hat{x}, \hat{\beta}) - (\underline{\sigma}(Q) - a_{V_N,1} - \tilde{\sigma}_w(\delta_w))|\delta\hat{x}|^2 + a_{V_N,2}c_e|(e, e^+)|^2 + \tilde{\sigma}_\alpha(|\Delta\alpha|). \end{aligned}$$

where the first inequality follows by optimality and the third inequality follows by (8.35).

Thus, (8.53) holds with $\tilde{a}_3 := \underline{\sigma}(Q) - a_{V_N,1} - \tilde{\sigma}_w(\delta_w) > 0$ and $\tilde{a}_4 := a_{V_N,2}c_e > 0$. \square

8.C Steady-state target problem assumptions

In this appendix, we prove Lemmas 8.32 and 8.37.

8.C.1 Proof of Lemma 8.32

Proof of Lemma 8.32. First, note that M_1 full row rank implies $n_r \leq n_u$. Consider the function

$$\mathbf{f}_1(z_s, \beta) := \begin{bmatrix} f(x_s, u_s, d) - x_s \\ g(u, h(x_s, u_s, d)) - r_{\text{sp}} \end{bmatrix}$$

and define the objective and Lagrangian

$$\begin{aligned} \phi(z_s, \beta) &:= \ell_s(u_s - u_{\text{sp}}, y_s(z_s, \beta) - y_{\text{sp}}) \\ \mathcal{L}(z_s, \beta, \lambda) &:= \phi(z_s, \beta) + \lambda^\top \mathbf{f}_1(z_s, \beta) \end{aligned}$$

where $z_s := (x_s, u_s)$, $y_s(z_s, \beta) := h(x_s, u_s, d)$, and $\beta := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d)$. The first-order derivatives of the Lagrangian are

$$\begin{aligned}\partial_{z_s} \mathcal{L}(z_s, \beta, \lambda) &= \partial_{z_s} \phi(z_s, \beta) + [\partial_{z_s} \mathbf{f}_1(z_s, \beta)]^\top \lambda \\ \partial_\lambda \mathcal{L}(z_s, \beta, \lambda) &= \mathbf{f}_1(z_s, \beta).\end{aligned}$$

The goal of the proof is to use the implicit function theorem on $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \beta, \lambda)$ to establish Lipschitz continuity of the SSTP solution map $z_s(\cdot)$. We already have $\partial_{(z_s, \lambda)} \mathcal{L}(0, 0, 0) = 0$ by assumption. Next, we aim to show $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \beta, \lambda) = 0$ is a necessary and sufficient condition for solving (8.6).

First, we have the partial derivatives $\partial_{z_s} \mathbf{f}_1(0, 0) = M_1$, which is full row rank by assumption. By continuity of $\partial_{z_s} \mathbf{f}_1$, there exist constants $\varepsilon_1, \delta_1 > 0$ such that $\partial_{z_s} \mathbf{f}_1(z_s, \beta)$ is full row rank for all $|z_s| \leq \varepsilon_1$ and $|\beta| \leq \delta_1$. Then, so long as (z_s, β) are kept sufficiently small, the linear independence constraint qualification holds, and $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \beta, \lambda) = 0$ is a necessary condition for solving (8.6).

Consider the following second-order derivatives:

$$\begin{aligned}\partial_{z_s}^2 \mathcal{L}(0, 0, 0) &= M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_3 \\ \partial_{z_s} \partial_\lambda \mathcal{L}(0, 0, 0) &= \partial_{z_s} \mathbf{f}_1(0, 0) = M_1 \\ \partial_\lambda^2 \mathcal{L}(0, 0, 0) &= 0\end{aligned}$$

where $M_3 := \begin{bmatrix} 0 & I \\ C & D \end{bmatrix}$.⁸ We have $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \beta, \lambda) = 0$ is a sufficient condition for solving (8.6) if

$$d^\top \partial_{z_s}^2 \mathcal{L}(z_s, \beta, \lambda) d > 0$$

⁸The second-order derivatives of $y_s(z_s, \beta)$ and $\mathbf{f}_1(z_s, \beta)$ vanish since $\partial_{(u, y)} \ell_s(0, 0)$ and $y_s(0, 0) = 0$ (by assumption) and we have set $\lambda = 0$.

for all $d \in \mathcal{N}(\partial_{z_s} \mathbf{f}_1(z_s, \beta)) \setminus \{0\}$. We require the following intermediate result.

Lemma 8.51. *For each $A = A^\top \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{m \times n}$, we have $x^\top Ax > 0$ for all $x \in \mathcal{N}(B) \setminus \{0\}$ if and only if $\begin{bmatrix} A \\ B \end{bmatrix}$ is full column rank.*

Proof. First, note that $\mathcal{N}(A + B^\top B) = \mathcal{N}(\begin{bmatrix} A \\ B \end{bmatrix})$, so $\begin{bmatrix} A \\ B \end{bmatrix}$ is full column rank is equivalent to $A + B^\top B$ being positive definite.

(\Rightarrow) Suppose $x^\top Ax > 0$ for all $x \in \mathcal{N}(B) \setminus \{0\}$. Then $x^\top (A + B^\top B)x \geq x^\top Ax > 0$ for all $x \in \mathcal{N}(B) \setminus \{0\}$ and $x^\top (A + B^\top B)x \geq x^\top B^\top Bx > 0$ for all $x \notin \mathcal{N}(B)$, so $A + B^\top B$ is positive definite.

(\Leftarrow) Suppose $A + B^\top B$ is positive definite. Then $x^\top Ax = x^\top (A + B^\top B)x > 0$ for all $x \in \mathcal{N}(B) \setminus \{0\}$. \square

Thus, it suffices to show

$$\begin{bmatrix} \partial_{z_s}^2 \mathcal{L}(z_s, \beta, \lambda) \\ \partial_{z_s} \mathbf{f}_1(z_s, \beta) \end{bmatrix} \quad (8.110)$$

is full column rank. Since $\partial_{(u,y)}^2 \ell_s(0, 0)$ is positive definite, $\mathcal{N}(\partial_{z_s}^2 \mathcal{L}(0, 0, 0)) = \mathcal{N}(M_3)$. Then

with $M_4 := \begin{bmatrix} \partial_{z_s}^2 \mathcal{L}(0, 0, 0) \\ \partial_{z_s} \mathbf{f}_1(0, 0) \end{bmatrix}$ we have

$$\begin{aligned} \mathcal{N}(M_4) &= \mathcal{N}(\partial_{z_s}^2 \mathcal{L}(0, 0, 0)) \cap \mathcal{N}(\partial_{z_s} \mathbf{f}_1(0, 0)) \\ &= \mathcal{N}(M_3) \cap \mathcal{N}(M_1) \\ &= \mathcal{N}\left(\begin{bmatrix} M_3 \\ M_1 \end{bmatrix}\right) = \{0\} \end{aligned}$$

where the last equality follows from the fact that

$$\begin{bmatrix} M_3 \\ M_1 \end{bmatrix} = \begin{bmatrix} 0 & I \\ A-I & B \\ H_y C & H_u + H_y D \end{bmatrix}$$

is full column rank, as it is the row permutation of a block triangular matrix with full col-

umn rank diagonal blocks I and $\begin{bmatrix} A-I \\ C \end{bmatrix}$.⁹ Therefore M_4 is full column rank, and because (8.110) is continuous, there exist $\varepsilon_2, \delta_2, \gamma_2 > 0$ for which (8.110) is full column rank for all $|z_s| \leq \varepsilon_2$, $|\beta| \leq \delta_2$, and $|\lambda| \leq \gamma_2$. Therefore, so long as (z_s, β, λ) are kept sufficiently small, $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \beta, \lambda) = 0$ is in fact a necessary and sufficient condition for solving (8.6).

Now we are able to solve (8.6). We have the derivatives

$$\partial_{(z_s, \lambda)}^2 \mathcal{L}(0, 0, 0) = \begin{bmatrix} M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_3 & M_1^\top \\ M_1 & 0 \end{bmatrix}.$$

According to (Magnus and Neudecker, 2019, Thm. 3.21), we have the nullspace relationship

$$\mathcal{N}(\partial_{(z_s, \lambda)}^2 \mathcal{L}(0, 0, 0)) = \mathcal{N}\left(\begin{bmatrix} V_0 \\ W_0 \end{bmatrix}\right) \quad (8.111)$$

where

$$V_0 := M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_3 + M_1^\top M_1 = \begin{bmatrix} M_3 \\ M_1 \end{bmatrix}^\top \begin{bmatrix} \partial_{(u, y)}^2 \ell_s(0, 0) & \\ & I \end{bmatrix} \begin{bmatrix} M_3 \\ M_1 \end{bmatrix}$$

$$W_0 := M_1 V_0^+ M_1^\top.$$

Recall $\begin{bmatrix} M_3 \\ M_1 \end{bmatrix}$ is full column rank and $\partial_{(u, y)}^2 \ell_s(0, 0)$ is invertible, so V_0 is invertible. Likewise, M_1 full row rank and V_0 invertible implies that W_0 is invertible. Finally, $\begin{bmatrix} V_0 \\ W_0 \end{bmatrix}$ is invertible, and by (8.111), $\partial_{(z_s, \lambda)}^2 \mathcal{L}(0, 0, 0)$ is invertible. By the implicit function theorem (Rudin, 1976, Thm. 9.24) there exist $\delta_3 > 0$ and continuously differentiable functions $\mathbf{g}_1 : \mathbb{R}^{n_\beta} \rightarrow \mathbb{R}^{n+n_u}$ and $\mathbf{g}_\lambda : \mathbb{R}^{n_\beta} \rightarrow \mathbb{R}^{n+n_r}$ such that $\mathbf{g}_1(0) = 0$, $\mathbf{g}_\lambda(0) = 0$, and $\partial_{(\alpha, \lambda)} \mathcal{L}(\mathbf{g}_1(\beta), \beta, \mathbf{g}_\lambda(\beta)) = 0$ for all $|\beta| \leq \delta_3$.

⁹Full column rank of $\begin{bmatrix} A-I \\ C \end{bmatrix}$ for all $|\lambda| \geq 1$ follows from detectability of (A, C) .

For convenience, we define the functions

$$\begin{aligned}\mathbf{g}_1(\beta) &:= (x_s(\beta), u_s(\beta)) \\ \tilde{c}(\beta) &:= \max_{1 \leq i \leq n_c} c_i(u_s(\beta), h(x_s(\beta), u_s(\beta), d)) + b_i\end{aligned}$$

for each $\beta = (r_{\text{sp}}, z_{\text{sp}}, d) \in \mathcal{B}$, which are continuous because \mathbf{g}_1 , h , and c are continuous. Moreover, \mathbb{X}, \mathbb{U} contain neighborhoods of the origin and $\tilde{c}(0) < 0$ by assumption, so there exists $\delta_3 > 0$ for which $z_s(\beta) \in \mathbb{X} \times \mathbb{U}$ and $\tilde{c}(\beta) \leq 0$ for all $|\beta| \leq \delta_3$. Let $\delta < \delta_4 := \min\{\delta_1, \delta_2, \delta_3\}$, $\delta_0 := \delta_4 - \delta$, $\mathcal{B}_c := \delta \mathbb{B}^{n_\beta}$, and $\overline{\mathcal{B}}_c := \delta_4 \mathbb{B}^{n_\beta}$. Defining $\hat{\mathcal{B}}_c$ as in Assumption 8.31(i), we have $|\hat{\beta}| \leq |\beta| + |e_d| \leq \delta + \delta_0 = \delta_4$ for each $\hat{\beta} = (s_{\text{sp}}, \hat{d}) \in \hat{\mathcal{B}}_c$, and therefore $\mathcal{B}_c \subseteq \hat{\mathcal{B}}_c \subseteq \overline{\mathcal{B}}_c \subseteq \mathcal{B}$. Moreover, $(x_s(\hat{\beta}), u_s(\hat{\beta})) \in \mathcal{Z}_O(r_{\text{sp}}, \hat{d})$ and $(x_s(\hat{\beta}), u_s(\hat{\beta}))$ uniquely solve (8.6) and are continuously differentiable for each $\hat{\beta} = (s_{\text{sp}}, \hat{d}) \in \hat{\mathcal{B}}_c$. Finally, Assumption 8.31 is satisfied by z_s , $\mathcal{B}_c \subseteq \mathcal{B}$, and $\delta_0 > 0$. \square

8.C.2 Proof of Lemma 8.37

Proof of Lemma 8.37. Recall from the proof of Lemma 8.32 that M_1 full row rank implies $n_r \leq n_u$. Moreover, M_2 invertible implies $n_d = n_y$. Consider the functions

$$\begin{aligned}\mathbf{f}_2(z_s, \mathbf{x}_s, \alpha) &:= \begin{bmatrix} f_P(x_{P,s}, u_s, w_P) - x_{P,s} \\ h_P(x_{P,s}, u_s, w_P) - h(x_s, u_s, d_s) \end{bmatrix} \\ \mathbf{f}(z_s, \mathbf{x}_s, \alpha) &:= \begin{bmatrix} \mathbf{f}_1(z_s, \beta) \\ \mathbf{f}_2(z_s, \mathbf{x}_s, \alpha) \end{bmatrix}\end{aligned}$$

where $z_s := (x_s, u_s)$, $\mathbf{x}_s := (x_{P,s}, d_s)$, $\alpha := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, w_P)$, $\beta := (r_{\text{sp}}, u_{\text{sp}}, y_{\text{sp}}, d_s)$, and \mathbf{f}_1 is defined in the proof of Lemma 8.32. Defining ϕ and \mathcal{L} as in the proof of Lemma 8.32, we seek

to use the implicit function theorem on

$$\begin{aligned} \mathbf{h}(z_s, \mathbf{x}_s, \lambda, \alpha) &:= \begin{bmatrix} \partial_{(z_s, \lambda)} \mathcal{L}(z_s, \mathbf{x}_s, \beta, \lambda) \\ \mathbf{f}_2(z_s, \mathbf{x}_s, \beta) \end{bmatrix} \\ &= \begin{bmatrix} \partial_{z_s} \phi(z_s, \mathbf{x}_s, \beta) + [\partial_{z_s} \mathbf{f}_1(z_s, \mathbf{x}_s, \beta)]^\top \lambda \\ \mathbf{f}(\alpha, \beta) \end{bmatrix} \end{aligned}$$

which is the combination of the stationary point condition for the Lagrangian of (8.6) with the steady-state disturbance problem (8.47). We already have $\mathbf{h}(0, 0, 0, 0) = 0$ by assumption. From the proof of Lemma 8.32, there exists $\delta_1 > 0$ such that, for all $|(\mathbf{x}_s, \alpha)| \leq \delta_1$, $\partial_{(z_s, \lambda)} \mathcal{L}(z_s, \mathbf{x}_s, \alpha, \lambda) = 0$ is a necessary and sufficient condition for solving (8.6). Thus, if we keep $|(\mathbf{x}_s, \alpha)| \leq \delta_1$ sufficiently small, then $\mathbf{h}(z_s, \mathbf{x}_s, \lambda, \alpha) = 0$ is necessarily and sufficient for simultaneously solving (8.6) and (8.47).

Defining the invertible matrices

$$T_1 := \begin{bmatrix} I_n & 0 & 0 & 0 \\ 0 & I_{n_r} & 0 & 0 \\ I_n & 0 & -I_n & 0 \\ 0 & 0 & 0 & I_{n_y} \end{bmatrix}, \quad T_2 := \begin{bmatrix} I_n & 0 & 0 & 0 \\ 0 & I_{n_u} & 0 & 0 \\ I_n & 0 & -I_n & 0 \\ 0 & 0 & 0 & I_{n_d} \end{bmatrix},$$

We have

$$T_1 \partial_{(z_s, \mathbf{x}_s)} \mathbf{f}(0, 0, 0) T_2 = \begin{bmatrix} M_1 & * \\ 0 & M_2 \end{bmatrix}.$$

We can write the derivatives

$$\partial_{(z_s, \mathbf{x}_s, \lambda)} \mathbf{h}(0, 0, 0, 0) = \begin{bmatrix} M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_5 & M_1^\top \\ \partial_{(z_s, \mathbf{x}_s)} \mathbf{f}(0, 0, 0) & 0 \end{bmatrix}$$

where M_3 is defined as in the proof of Lemma 8.32, and $M_5 := \begin{bmatrix} 0 & I & 0 & 0 \\ C & D & 0 & C_d \end{bmatrix}$. Note that $M_5 T_2 = M_5$ and $M_5 = \begin{bmatrix} M_3 & * \end{bmatrix}$. Define the invertible matrices

$$T_3 := \begin{bmatrix} I_{n+n_u} & \\ & T_1 \end{bmatrix}, \quad T_4 := \begin{bmatrix} T_2 & \\ & I_{n+n_d} \end{bmatrix}, \quad P := \begin{bmatrix} I_{n+n_u} & 0 & 0 \\ 0 & 0 & I_{n+n_d} \\ 0 & I_{n+n_r} & 0 \end{bmatrix}.$$

Then we can write

$$T_3 \partial_{(z_s, \mathbf{x}_s, \lambda)} \mathbf{h}(0, 0, 0, 0) T_4 P = \begin{bmatrix} M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_3 & M_1^\top & * \\ M_1 & 0 & * \\ 0 & 0 & M_2 \end{bmatrix}. \quad (8.112)$$

But M_2 is invertible by assumption, and $\begin{bmatrix} M_3^\top \partial_{(u, y)}^2 \ell_s(0, 0) M_3 & M_1^\top \\ M_1 & 0 \end{bmatrix}$ was shown to be invertible in the proof of Lemma 8.32, so $\partial_{(z_s, \mathbf{x}_s, \lambda)} \mathbf{h}(0, 0, 0, 0)$ must be invertible. By the implicit function theorem (Rudin, 1976, Thm. 9.24) there exist $\delta_2 > 0$ and continuously differentiable functions $\mathbf{g} : \mathbb{R}^{n_\alpha} \rightarrow \mathbb{R}^{2n+n_u+n_d}$ and $\mathbf{g}_\lambda : \mathbb{R}^{n_\alpha} \rightarrow \mathbb{R}^{2n+n_r+n_y}$ (where $\mathcal{A} := \mathbb{R}^{n_r} \times \overline{\mathbb{Z}}_y \times \mathbb{W}$) such that $\mathbf{g}(0) = 0$, $\mathbf{g}_\lambda(0) = 0$, and $\partial_{(z_s, \mathbf{x}_s, \lambda)} \mathcal{L}(\mathbf{g}(\alpha), \alpha, \mathbf{g}_\lambda(\alpha)) = 0$ for all $|\alpha| \leq \delta_2$.

As in the proof of Lemma 8.32, we define the functions

$$\begin{aligned} \mathbf{g}(\alpha) &:= (x_s(\alpha), u_s(\alpha), x_{P,s}(\alpha), d_s(\alpha)) \\ \tilde{c}(\alpha) &:= \max_{1 \leq i \leq n_c} c_i(u_s(\alpha), h_P(x_{P,s}(\alpha), u_s(\alpha), w_P)) + b_i \end{aligned}$$

for each $\alpha = (r_{\text{sp}}, z_{\text{sp}}, w_P) \in \mathbb{R}^{n_\alpha}$, which are continuous because \mathbf{g} , h_P , and c are continuous. From Lemma 8.32, we already have a set $\mathcal{B}_c \subseteq \mathcal{B}$ containing a neighborhood of the origin and continuously differentiable functions (with a slight abuse of notation) $(x_s, u_s) : \mathcal{B} \rightarrow \mathbb{X} \times \mathbb{U}$ that uniquely solve (8.6) (and satisfies Assumption 8.31). Since $\mathbb{X}, \mathbb{U}, \mathbb{D}, \mathcal{B}_c$ contain neighborhoods of the origin, there must exist $\delta_3 > 0$ such that $\mathbf{g}(\alpha) \in \mathbb{X} \times \mathbb{U} \times \mathbb{X} \times \mathbb{D}$, $\beta = (r_{\text{sp}}, z_{\text{sp}}, d_s(\alpha)) \in \mathcal{B}_c$, $|(x_s(\alpha), \alpha)| \leq \delta_2$, and $\tilde{c}(\beta) \leq 0$ for all $|\alpha| \leq \delta_3$. Therefore $(x_s(\alpha), u_s(\alpha))$ are also the unique solutions to (8.6) with $\beta = (r_{\text{sp}}, z_{\text{sp}}, d_s(\alpha))$, i.e., $(x_s(\alpha), u_s(\alpha)) = (x_s(\beta), u_s(\beta))$, and all parts of Assumption 8.36 are satisfied with $(x_s, u_s) : \mathcal{B} \rightarrow \mathbb{X} \times \mathbb{U}$, $(x_{P,s}, d_s) : \mathcal{A}_c \rightarrow \mathbb{X} \times \mathbb{D}$, $\mathcal{A}_c := \delta \mathbb{B}^{n_\alpha}$, and $\delta := \min \{ \delta_1, \delta_2, \delta_3 \} > 0$. \square

Chapter 9

Conclusion

When we purchase a new piece of technology, our expectation and reality is that the product works perfectly, right out of the box, in a “turnkey” fashion. For plant operators dealing with control systems, this is less of a reality and more of a distant dream. The deployment of control systems of all types—from simple PID loops, to high-level production scheduling with model predictive control—is plagued with time-consuming tuning steps that use ad hoc methods to produce suboptimal performance. This work has been driven by the philosophy that data-based design is the most rigorous and optimal way to acquire an estimator, including for offset-free control.

9.1 Summary

Part I: Identification

In Part I, we presented methods for identifying models with integrating disturbances. In Chapter 3 we developed maximum likelihood identification for offset-free control applications. We contributed two key improvements upon the prior methodology. First, we incorporated high-level design and modeling constraints by way of LMI region-based eigenvalue constraints and sparsity structuring. Second, we convert the constraints into a well-posed

nonlinear program by introducing constraint back-offs to strict inequalities, and substituting positive semidefinite matrices for their Cholesky factors. These methods rigorously approximate the maximum likelihood problem for the ill-posed constraint set. In Chapter 4, we extend a broad class of standard linear identification methods to allow for disturbance model identification. We focus on simple methods with closed-form solutions and provide a straightforward method for industrial practitioners to adopt data-driven estimator tuning methods.

Part II: Application

In Part II, we apply the methods presented in the previous two chapters to the offset-free control of two systems: a benchmark temperature controller, and an operational, industrial-scale chemical reactor at Eastman Chemical's plant in Kingsport, TN. All of the case studies in this part use real-world data collected from these systems. We show that the eigenvalue constraints of Chapter 3 can improve estimator performance, and prevent the identification of unstable or otherwise poorly tuned estimators. We establish that superior closed-loop performance can be achieved by data-driven estimator designs, even over well-designed MPCs (in this case, the MPC was on-line for at least 20 years Caveness and Downs (2005)). Our data-based tuning achieved a 38% reduction in setpoint tracking error on Eastman's chemical reactor.

Part III: Theory

In Part III, we advance the theory on stability of MPC subject to plant-model mismatch. We start with a simple analysis of unconstrained linear optimal control for nonlinear systems in Chapter 6. We consider nonlinear standard MPC with multiplicative errors in Chapter 7. Finally, we extend the theory to the general, nonlinear offset-free MPC in Chapter 8. This theory is the first of its kind in establishing closed-loop stability and offset-free performance

of nonlinear MPCs with integrating disturbances, and of MPCs of all kinds subject to plant-model mismatch.

9.2 Future work

Data-driven control

The approach of Parts I and II is *indirect* data-driven control, where model parameters are first estimated, and a controller designed based on it. A potential alternative is the direct data-driven control approach, where the controller itself is designed according to data (Berberich et al., 2021; Dorfler et al., 2022; Berberich et al., 2022a; Yuan and Cortés, 2022; Bianchin et al., 2023). Data-driven model predictive controllers have been suggested by some authors (Berberich et al., 2021, 2022a), and while the designs are generally robust to disturbances, there are no designs that can boast, rigorously, the kind of offset-free performance we are afforded by offset-free MPC. The main limitation to these approaches is the reliance on Willem’s Fundamental Lemma (Willems et al., 2005), which assumes the data is generated from a plant of the model class and does not allow structured models. We also remark the models considered in this thesis have far more general noise models than those considered in direct data-driven control works. Despite these limitations, there is still a possibility of improving the theory with designs based on likelihood functions (Yin et al., 2023) or by bridging direct and indirect methodologies (Dorfler et al., 2022).

Performance monitoring

A promising area of future research is on the monitoring of MPC performance, model updating, and fault diagnosis. With maximum likelihood estimates of the system parameters, we are in good shape to continue developing the statistical performance monitoring algorithm

of Zagrobelny et al. (2013). Model fault diagnosis for offset-free MPC has mostly relied on heuristic methods to detect and diagnose problems (Harrison and Qin, 2009; Pannocchia and De Luca, 2012). The first steps towards rigorous re-identification are likely case studies on real processes with poorly behaving controllers. A simple application for such an algorithm is the self-tuning PID loop, where the on-line acquisition of a simple process model enables the real-time tuning of the PID loop to new disturbance information.

Economic model predictive control

We mostly ignored the cross-over between economic and offset-free MPC designs Pannocchia (2018); Faulwasser and Pannocchia (2019); Vaccari et al. (2021). However, this literature addresses some important concerns about applying “optimal” control algorithms to real-world systems. We put “optimal” in quotes here due to the futility of the situation for real-world systems: due to plant-model mismatch, we can never hope to achieve optimal performance, only to achieve some level of suboptimality. This is exacerbated on complex systems such as chemical plants, where unknown physics and black-box modeling is routine. Pannocchia (2018) originally suggested a gradient-updating scheme to allow for asymptotically optimal performance. However, these approaches are plagued by the same theoretical limitations as was offset-free MPC (before this work). Another option is to consider linear approximations of the nonlinear economic MPC (Zanon et al., 2016, 2017). While this simplifies the approach and makes it more attractive to a practical audience, it does not address plant-model mismatch.

A simpler and more comprehensive theory

Our stability results are similar but different to the standard inherent robustness results of the control literature (De Nicolao et al., 1996; Scokaert et al., 1997; Grimm et al., 2004; Pannoc-

chia et al., 2011; Allan et al., 2017). We have found the theory of Chapters 7 and 8 considerably narrower in terms of the allowed class of MPC designs. Of course, inherent robustness is a much weaker requirement than stability despite persistent disturbances. However, our analysis is limited to the quadratic cost case, whereas inherent robustness holds in general for coercive cost functions. There may be additional conditions on the cost functions that provide the kind of dissipation inequality we require to demonstrate stability despite mismatch.

A drawback of the theory of Chapter 8 is the reliance on Lyapunov functions for the estimator. We know of no approaches that guarantee the existence of Lyapunov functions for optimization-based estimators, although some authors have provided close alternatives, such as the so-called Q function and related incremental input-output-to-state stability (i-IOSS) notion (Allan and Rawlings, 2019, 2021; Allan et al., 2021), and N -step-ahead Lyapunov functions (Schiller et al., 2023). Moreover, the nature of uncontrollable integrating disturbances invalidates most i-IOSS assertions. Disturbances have been handled separately (usually in the form of parameter drift) by Muntwiler et al. (2023); Schiller and Müller (2023).

Bibliography

- H. Akaike. A new look at statistical model identification. *IEEE Trans. Auto. Cont.*, 19(6): 716–723, 1974.
- 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, 2019.
- D. A. Allan and J. B. Rawlings. Robust stability of full information estimation. *SIAM J. Cont. Opt.*, 59(5):3472–3497, 2021.
- 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: 10.1016/j.sysconle.2017.03.005.
- D. A. Allan, J. B. Rawlings, and A. R. Teel. Nonlinear detectability and incremental input/output-to-state stability. *SIAM J. Cont. Opt.*, 59(4):3017–3039, 2021.
- B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall, Englewood Cliffs, N. J., 1979.
- B. D. O. Anderson, J. B. Moore, and R. M. Hawkes. Model approximations via prediction error identification. *Automatica*, 14(6):615–622, 1978. ISSN 0005-1098. doi: 10.1016/0005-1098(78)90051-1.
- T. W. Anderson. Asymptotic distribution of the reduced rank regression estimator under general conditions. *Ann. Stat.*, 27(4):1141–1154, Aug 1999. ISSN 0090-5364, 2168-8966. doi: 10.1214/aos/1017938918.
- T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, New York, third edition, 2003.
- T. M. Apostol. *Mathematical analysis*. Addison-Wesley, 1974.
- T. J. Arnold and J. B. Rawlings. Tractable calculation and estimation of the optimal weighting for autocovariance least squares. *IEEE Trans. Auto. Cont.*, 67(11):6045–60, 2022. doi: <https://doi.org/10.1109/TAC.2021.3124193>.

- K. Åström and P. Eykhoff. System identification—A survey. *Automatica*, 7(2):123–162, 1971. ISSN 00051098. doi: 10.1016/0005-1098(71)90059-8.
- K. Åström and C. Källström. Identification of ship steering dynamics. *Automatica*, 12(1):9–22, 1976. ISSN 00051098. doi: 10.1016/0005-1098(76)90064-9.
- K. J. Åström. Maximum Likelihood and Prediction Error Methods. *IFAC Proceedings Volumes*, 12(8):551–574, Sep 1979. ISSN 1474-6670.
- K. J. Åström. Special section system identification tutorial: Maximum likelihood and prediction error methods. *Automatica*, 16(5):551–574, 1980.
- D. Bauer. Order estimation for subspace methods. *Automatica*, 37(10):1561–1573, 2001.
- K. Baumgärtner, R. Reiter, and M. Diehl. Moving Horizon Estimation with Adaptive Regularization for Ill-Posed State and Parameter Estimation Problems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 2165–2171, 2022.
- J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer. Data-Driven Model Predictive Control With Stability and Robustness Guarantees. *IEEE Trans. Auto. Cont.*, 66(4):1702–1717, 2021. ISSN 1558-2523. doi: 10.1109/TAC.2020.3000182.
- J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer. Linear Tracking MPC for Nonlinear Systems—Part II: The Data-Driven Case. *IEEE Trans. Auto. Cont.*, 67(9):4406–4421, 2022a. ISSN 1558-2523. doi: 10.1109/TAC.2022.3166851.
- J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer. Linear Tracking MPC for Nonlinear Systems—Part I: The Model-Based Case. *IEEE Trans. Auto. Cont.*, 67(9):4390–4405, 2022b. ISSN 1558-2523. doi: 10.1109/TAC.2022.3166872.
- J. Berberich, C. W. Scherer, and F. Allgöwer. Combining Prior Knowledge and Data for Robust Controller Design. *IEEE Trans. Auto. Cont.*, 68(8):4618–4633, 2023. ISSN 1558-2523. doi: 10.1109/TAC.2022.3209342.
- O. Bermond and J.-F. Cardoso. Approximate likelihood for noisy mixtures. In *International Workshop on Independent Component Analysis (ICA '99)*, volume 99, pages 325–330, Aussois, France, 1999.
- D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, New York, 1978.
- G. Betti, M. Farina, and R. Scattolini. A Robust MPC Algorithm for Offset-Free Tracking of Constant Reference Signals. *IEEE Trans. Auto. Cont.*, 58(9):2394–2400, 2013. ISSN 1558-2523. doi: 10.1109/TAC.2013.2254011.

- G. Bianchin, M. Vaquero, J. Cortés, and E. Dall’Anese. Online Stochastic Optimization for Unknown Linear Systems: Data-Driven Controller Synthesis and Analysis. *IEEE Trans. Auto. Cont.*, pages 1–15, 2023. ISSN 1558-2523. doi: 10.1109/TAC.2023.3323581.
- J. F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research. Springer Science & Business Media, New York, NY, 2000.
- S. Boyd, L. E. Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in Control Theory*. SIAM, Philadelphia, 1994.
- A. E. Bryson and Y. Ho. *Applied Optimal Control*. Hemisphere Publishing, New York, 1975.
- S. Burer, R. D. Monteiro, and Y. Zhang. Solving a class of semidefinite programs via nonlinear programming. *Math. Prog.*, 93(1):97–122, 2002a. ISSN 1436-4646.
- S. Burer, R. D. Monteiro, and Y. Zhang. Interior-Point Algorithms for Semidefinite Programming Based on a Nonlinear Formulation. *Comp. Optim. Appl.*, 22(1):49–79, 2002b. doi: 10.1023/A:1014834318702.
- W. M. Canney. The future of advanced process control promises more benefits and sustained value. *Oil & Gas J.*, 101(16):48–54, Apr 2003. ISSN 00301388.
- M. H. Caviness and J. J. Downs. Reactor control using infinite horizon model predictive control. AIChE Spring Meeting, Atlanta, GA, April 2005.
- T. Chen and L. Ljung. Implementation of algorithms for tuning parameters in regularized least squares problems in system identification. *Automatica*, 49(7):2213–2220, 2013. ISSN 0005-1098. doi: 10.1016/j.automatica.2013.03.030.
- T. Chen, M. S. Andersen, L. Ljung, A. Chiuso, and G. Pillonetto. System Identification Via Sparse Multiple Kernel-Based Regularization Using Sequential Convex Optimization Techniques. *IEEE Trans. Auto. Cont.*, 59(11), 2014. ISSN 1558-2523. doi: 10.1109/TAC.2014.2351851.
- M. Chilali and P. Gahinet. H_∞ design with pole placement constraints: an LMI approach. *IEEE Trans. Auto. Cont.*, 41(3):358–367, 1996. ISSN 1558-2523. doi: 10.1109/9.486637.
- M. Chilali, P. Gahinet, and P. Apkarian. Robust pole placement in LMI regions. *IEEE Trans. Auto. Cont.*, 44(12):2257–2269, 1999.
- A. Chiuso. On the Asymptotic Properties of Closed-Loop CCA-Type Subspace Algorithms: Equivalence Results and Role of the Future Horizon. *IEEE Trans. Auto. Cont.*, 55(3):634–649, 2010. ISSN 1558-2523.
- D. W. Clarke. Adaptive generalized predictive control. In *Proceedings of the Fourth International Conference on Chemical Process Control*, pages 395–417, 1991.

- D. W. Clarke, C. Mohtadi, and P. S. Tuffs. Generalized predictive control—Part I. The basic algorithm. *Automatica*, 23(2):137–148, 1987a.
- D. W. Clarke, C. Mohtadi, and P. S. Tuffs. Generalized predictive control—Part II. Extensions and interpretations. *Automatica*, 23(2):149–160, 1987b.
- F. Clarke, Y. Ledyae, and R. Stern. Asymptotic Stability and Smooth Lyapunov Functions. *J. Diff. Eq.*, 149(1):69–114, 1998. doi: 10.1006/jdeq.1998.3476.
- C. R. Cutler and B. L. Ramaker. Dynamic matrix control—a computer control algorithm. In *Proceedings of the Joint Automatic Control Conference*, 1980.
- P. Daoutidis, W. Tang, and A. Allman. Decomposition of control and optimization problems by network structure: Concepts, methods, and inspirations from biology. *AIChE J.*, 65(10): e16708, 2019. doi: 10.1002/aic.16708.
- M. L. Darby and M. Nikolaou. MPC: Current practice and challenges. *Control Eng. Pract.*, 20(4):328 – 342, 2012.
- 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, Mar 1996.
- D. Deenen, E. Maljaars, L. Sebeke, B. De Jager, E. Heijman, H. Gröll, and W. P. M. H. Heemels. Offset-free model predictive control for enhancing MR-HIFU hyperthermia in cancer treatment. *IFAC–P. Online*, 51(20):191–196, 2018. ISSN 24058963. doi: 10.1016/j.ifacol.2018.11.012.
- M. Denham. Canonical forms for the identification of multivariable linear systems. *IEEE Trans. Auto. Cont.*, 19(6):646–656, 1974.
- M. Diehl, K. Mombaur, and D. Noll. Stability Optimization of Hybrid Periodic Systems via a Smooth Criterion. *IEEE Trans. Auto. Cont.*, 54(8):1875–1880, 2009. ISSN 1558-2523. doi: 10.1109/TAC.2009.2020669.
- V. Digalakis, J. Rohlicek, and M. Ostendorf. ML estimation of a stochastic linear system with the EM algorithm and its application to speech recognition. *IEEE Trans. Speech and Audio Proc.*, 1(4):431–442, 1993. ISSN 1558-2353. doi: 10.1109/89.242489.
- F. Dorfler, J. Coulson, and I. Markovskiy. Bridging Direct & Indirect Data-Driven Control Formulations via Regularizations and Relaxations. *IEEE Trans. Auto. Cont.*, 68(2):883–897, 2022. ISSN 1558-2523. doi: 10.1109/TAC.2022.3148374.
- R. Douc and E. Moulines. Asymptotic properties of the maximum likelihood estimation in misspecified hidden Markov models. *Ann. Stat.*, 40(5):2697–2732, 2012.
- J. C. Doyle. Guaranteed margins for LQG regulators. *IEEE Trans. Auto. Cont.*, 23:756–757, 1978.

- P. Falugi. Model predictive control for tracking randomly varying references. *Int. J. Control*, 88(4):745–753, 2015. ISSN 0020-7179. doi: 10.1080/00207179.2014.972464.
- 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.
- T. Faulwasser and G. Pannocchia. Toward a Unifying Framework Blending Real-Time Optimization and Economic Model Predictive Control. *Ind. Eng. Chem. Res.*, 58(30):13583–13598, 2019. ISSN 0888-5885. doi: 10.1021/acs.iecr.9b00782.
- J. Fiala, M. Kočvara, and M. Stingl. PENLAB: A MATLAB solver for nonlinear semidefinite optimization, 2013.
- S. Formentin and A. Chiuso. Control-oriented regularization for linear system identification. *Automatica*, 127:109539, 2021. ISSN 0005-1098. doi: 10.1016/j.automatica.2021.109539.
- P. Gahinet, A. Laub, C. Kenney, and G. Hoyer. Sensitivity of the stable discrete-time Lyapunov equation. *IEEE Trans. Auto. Cont.*, 35(11):1209–1217, 1990. ISSN 1558-2523. doi: 10.1109/9.59806.
- G. Galuppini, L. Magni, and A. Ferramosca. Nonlinear MPC for Tracking Piecewise-Constant Reference Signals: the Positive Semidefinite Stage Cost Case. *IFAC-P. Online*, 56(1):210–215, 2023. ISSN 2405-8963. doi: 10.1016/j.ifacol.2023.02.036.
- S. Gibson and B. Ninness. Robust maximum-likelihood estimation of multivariable dynamic systems. *Automatica*, 41(10):1667–1682, 2005.
- G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.
- G. Gong and F. J. Samaniego. Pseudo Maximum Likelihood Estimation: Theory and Applications. *Ann. Stat.*, 9(4):861–869, 1981. ISSN 0090-5364.
- 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. Nominally robust model predictive control with state constraints. *IEEE Trans. Auto. Cont.*, 52(10):1856–1870, Oct 2007.
- L. Grüne, E. D. Sontag, and F. R. Wirth. Asymptotic stability equals exponential stability, and ISS equals finite energy gain — if you twist your eyes. *Sys. Cont. Let.*, 38(2):127–134, 1999. ISSN 0167-6911. doi: 10.1016/S0167-6911(99)00053-5.
- A. Haber and M. Verhaegen. Sparse solution of the Lyapunov equation for large-scale interconnected systems. *Automatica*, 73:256–268, 2016. ISSN 00051098. doi: 10.1016/j.automatica.2016.06.002.

- J. Hale. *Ordinary Differential Equations*. Robert E. Krieger Publishing Company, second edition, 1980.
- E. J. Hannan and M. Deistler. *The Statistical Theory of Linear Systems*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 2012. ISBN 978-1-61197-218-4. doi: 10.1137/1.9781611972191.
- C. A. Harrison and S. J. Qin. Discriminating between disturbance and process model mismatch in model predictive control. *J. Proc. Cont.*, 19(10):1610–1616, 2009.
- J. P. Hespanha. *Linear Systems Theory*. Princeton University Press, Princeton and Oxford, second edition, 2018.
- G. A. Hicks and W. H. Ray. Approximation methods for optimal control synthesis. *Can. J. Chem. Eng.*, 49:522–528, Aug 1971.
- B. L. Ho and R. E. Kalman. Efficient construction of linear state variable models from input/output functions. *Regelungstechnik*, 14:545–548, 1966.
- A. E. Hoerl and R. W. Kennard. Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*, 12(1):55–67, Feb 1970a. ISSN 0040-1706. doi: 10.1080/00401706.1970.10488634.
- A. E. Hoerl and R. W. Kennard. Ridge Regression: Applications to Nonorthogonal Problems. *Technometrics*, 12(1):69–82, Feb 1970b. ISSN 0040-1706. doi: 10.1080/00401706.1970.10488635.
- R. Huang, S. C. Patwardhan, and L. T. Biegler. Offset-free Nonlinear Model Predictive Control Based on Moving Horizon Estimation for an Air Separation Unit. *IFAC Proceedings Volumes*, 43(5):631–636, 2010. ISSN 14746670. doi: 10.3182/20100705-3-BE-2011.00105.
- M. Jansson. Subspace Identification and ARX Modeling. *IFAC Proceedings Volumes*, 36(16):1585–1590, 2003. ISSN 1474-6670.
- Z.-P. Jiang and Y. Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37:857–869, 2001.
- Z.-P. Jiang and Y. Wang. A converse Lyapunov theorem for discrete-time systems with disturbances. *Sys. Cont. Let.*, 45:49–58, 2002.
- T. A. Johansen. On Tikhonov regularization, bias and variance in nonlinear system identification. *Automatica*, 33(3):441–446, 1997. ISSN 0005-1098. doi: 10.1016/S0005-1098(96)00168-9.
- J. B. Jørgensen, J. K. Huusom, and J. B. Rawlings. Finite horizon MPC for systems in innovation form. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 1896–1903, 2011.

- A. B. W. Jr. On the numerical solution of initial/boundary-value problems in one space dimension. *SIAM J. Numer. Anal.*, 19(4):683–697, 1982.
- C. G. Källström and K. J. Åström. Experiences of System Identification Applied to Ship Steering Dynamics. *IFAC Proceedings Volumes*, 12(8):173–184, 1979. ISSN 1474-6670. doi: 10.1016/S1474-6670(17)65419-3.
- S. Kameswaran and L. T. Biegler. Simultaneous dynamic optimization strategies: Recent advances and challenges. *Comput. Chem. Eng.*, 30:1560–1575, Sep 2006.
- C. M. Kellett. A compendium of comparison function results. *Math. Contr. Sign. Syst.*, 26(3): 339–374, 2014. doi: 10.1007/s00498-014-0128-8.
- C. M. Kellett and A. R. Teel. On the robustness of \mathcal{KL} -stability for difference inclusions: Smooth discrete-time Lyapunov functions. *SIAM J. Cont. Opt.*, 44(3):777–800, 2005. doi: 10.1137/S0363012903435862.
- H. K. Khalil. *Nonlinear Systems*. Prentice-Hall, Upper Saddle River, NJ, third edition, 2002.
- M. Kheradmandi and P. Mhaskar. Model predictive control with closed-loop re-identification. *Comput. Chem. Eng.*, 109:249–260, 2018.
- M. Kočvara and M. Stingl. PENNON: Software for linear and nonlinear matrix inequalities, 2015.
- J. Köhler, M. A. Müller, and F. Allgöwer. A nonlinear tracking model predictive control scheme for dynamic target signals. *Automatica*, 118:109030, 2020. ISSN 0005-1098. doi: 10.1016/j.automatica.2020.109030.
- S. J. Kuntz and J. B. Rawlings. Maximum likelihood estimation of linear disturbance models for offset-free model predictive control. In *American Control Conference*, pages 3961–3966, Atlanta, GA, June 8–10, 2022.
- S. J. Kuntz and J. B. Rawlings. On the unified theory of linear Gaussian estimation: solution methods, applications, and extensions. Technical Report 2023–01, TWCCC Technical Report, December 2023a. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2023-01.pdf>.
- S. J. Kuntz and J. B. Rawlings. On linear multivariate regression with singular error covariances. Technical Report 2023–02, TWCCC Technical Report, December 2023b. URL <https://sites.engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2023-02.pdf>.
- S. J. Kuntz and J. B. Rawlings. Maximum Likelihood Identification of Linear Models with Integrating Disturbances for Offset-Free Control. *IEEE Trans. Auto. Cont.*, 2024a. Submitted 6/5/2024, Revised 11/4/2024.

- S. J. Kuntz and J. B. Rawlings. Maximum likelihood identification of uncontrollable linear time-invariant models for offset-free control. *arXiv preprint arXiv:2406.03760*, 2024b.
- S. J. Kuntz and J. B. Rawlings. Beyond inherent robustness: strong stability of MPC despite plant-model mismatch. *IEEE Trans. Auto. Cont.*, 2024c. Submitted 11/22/2024.
- S. J. Kuntz and J. B. Rawlings. Offset-free model predictive control: stability under plant-model mismatch. *IEEE Trans. Auto. Cont.*, 2025. In progress.
- S. J. Kuntz, J. J. Downs, S. M. Miller, and J. B. Rawlings. An industrial case study on the combined identification and offset-free control of a chemical process. *Comput. Chem. Eng.*, 179, 2023. doi: <https://doi.org/10.1016/j.compchemeng.2023.108429>.
- O. Y. Kushel. Geometric properties of LMI regions, 2019. URL <http://arxiv.org/abs/1910.10372>. arXiv:1910.10372 [math].
- H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems*. John Wiley and Sons, New York, 1972. ISBN 0-471-51110-2.
- W. E. Larimore. System identification, reduced-order filtering and modeling via canonical variate analysis. In *Proceedings of the American Control Conference*, pages 445–451, San Fransisco, CA, 1983.
- W. E. Larimore. Canonical variate analysis in identification, filtering, and adaptive control. In *Proceedings of the 29th Conference on Decision and Control*, pages 596–604, 1990.
- W. E. Larimore. Optimal reduced rank modeling, prediction, monitoring, and control using canonical variate analysis. In *Proceedings of the IFAC 1997 International Symposium on Advanced Control*, pages 61–66, Banff, Canada, June 1997.
- W. E. Larimore. Maximum likelihood subspace identification for linear, nonlinear, and closed-loop systems. In *Proceedings of the 2005 American Control Conference*, pages 2305–2319. IEEE, 2005.
- M. Lazar and W. P. M. H. Heemels. Predictive control of hybrid systems: Input-to-state stability results for sub-optimal solutions. *Automatica*, 45(1):180–185, 2009.
- J. H. Lee and Z. H. Yu. Tuning of model predictive controllers for robust performance. *Comput. Chem. Eng.*, 18(1):15–37, 1994.
- J. H. Lee, M. Morari, and C. E. García. State-space interpretation of model predictive control. *Automatica*, 30(4):707–717, 1994.
- J. M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, NY, 2012. ISBN 978-1-4419-9981-8 978-1-4419-9982-5. doi: 10.1007/978-1-4419-9982-5.

- N. Lehtomaki, N. Sandell, and M. Athans. Robustness results in linear-quadratic Gaussian based multivariable control designs. *IEEE Trans. Auto. Cont.*, 26(1):75–93, 1981. doi: 10.1109/TAC.1981.1102565.
- P. Li, I. Postlethwaite, and M. Turner. Parameter Estimation Techniques for Helicopter Dynamic Modelling. In *2007 American Control Conference*, pages 2938–2943, 2007. doi: 10.1109/ACC.2007.4282197.
- D. Limon, T. Alamo, F. Salas, and E. F. Camacho. On the stability of MPC without terminal constraint. *IEEE Trans. Auto. Cont.*, 51(5):832–836, May 2006.
- D. Limon, I. Alvarado, T. Alamo, and E. F. Camacho. MPC for tracking piecewise constant references for constrained linear systems. *Automatica*, 44:2382–2387, 2008.
- D. Limon, A. Ferramosca, I. Alvarado, and T. Alamo. Nonlinear MPC for Tracking Piece-Wise Constant Reference Signals. *IEEE Trans. Auto. Cont.*, 63(11):3735–3750, 2018. ISSN 0018-9286, 1558-2523, 2334-3303. doi: 10.1109/TAC.2018.2798803.
- L. Ljung. Convergence analysis of parametric identification methods. *IEEE Trans. Auto. Cont.*, 23(5):770–783, 1978. ISSN 1558-2523. doi: 10.1109/TAC.1978.1101840.
- L. Ljung. *System Identification: Theory for the User*. Prentice Hall, New Jersey, second edition, 1999.
- P. Lundström, J. H. Lee, M. Morari, and S. Skogestad. Limitations of dynamic matrix control. *Comput. Chem. Eng.*, 19:409–421, 1995.
- U. Maeder and M. Morari. Offset-free reference tracking with model predictive control. *Automatica*, 46(9):1469–1476, 2010. ISSN 00051098. doi: 10.1016/j.automatica.2010.05.023.
- J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. John Wiley & Sons Ltd., Hoboken, NJ, third edition, 2019.
- P. Marquis and J. P. Broustail. SMOC, a bridge between state space and model predictive controllers: Application to the automation of a hydrotreating unit. In T. J. McAvoy, Y. Arkun, and E. Zafiriou, editors, *Proceedings of the 1988 IFAC Workshop on Model Based Process Control*, pages 37–43, Oxford, 1988. Pergamon Press.
- S. Massei and L. Saluzzi. On the data-sparsity of the solution of Riccati equations with applications to feedback control, 2024. URL <http://arxiv.org/abs/2408.16569>. arXiv:2408.16569 [cs, math].
- R. D. McAllister. *The Stochastic Robustness of Model Predictive Control and Closed-Loop Scheduling*. PhD thesis, University of California, Santa Barbara, August 2022. URL <https://engineering.ucsb.edu/~jbraw/jbrweb-archives/theses/mcallister.pdf>.

- R. D. McAllister and J. B. Rawlings. The stochastic robustness of nominal and stochastic model predictive control. *IEEE Trans. Auto. Cont.*, 2022a. doi: <https://doi.org/10.1109/TAC.2022.3226712>. Online Early Access.
- R. D. McAllister and J. B. Rawlings. Inherent stochastic robustness of model predictive control to large and infrequent disturbances. *IEEE Trans. Auto. Cont.*, 67(10):5166–5178, 2022b. doi: <https://doi.org/10.1109/TAC.2021.3122365>.
- R. D. McAllister and J. B. Rawlings. On the inherent distributional robustness of stochastic and nominal model predictive control. *IEEE Trans. Auto. Cont.*, 69(2):741–754, 2024. doi: <https://doi.org/10.1109/TAC.2023.3273420>.
- T. McKelvey, A. Helmersson, and T. Ribarits. Data driven local coordinates for multivariable linear systems and their application to system identification. *Automatica*, 40(9):1629–1635, 2004. ISSN 0005-1098.
- H. Melgaard, E. Hendricks, and H. Madsen. Continuous Identification of a Four-Stroke SI Engine. In *1990 American Control Conference*, pages 1876–1881, 1990. doi: 10.23919/ACC.1990.4791053.
- D. N. Miller and R. A. De Callafon. Subspace identification with eigenvalue constraints. *Automatica*, 49(8):2468–2473, 2013. ISSN 00051098. doi: 10.1016/j.automatica.2013.04.028.
- M. Morari and U. Maeder. Nonlinear offset-free model predictive control. *Automatica*, 48(9):2059–2067, 2012.
- N. Motee and A. Jadbabaie. Distributed Multi-Parametric Quadratic Programming. *IEEE Trans. Auto. Cont.*, 54(10):2279–2289, 2009. ISSN 1558-2523. doi: 10.1109/TAC.2009.2014916.
- N. Motee and Q. Sun. Sparsity and Spatial Localization Measures for Spatially Distributed Systems. *SIAM J. Cont. Opt.*, 55(1):200–235, Jan 2017. ISSN 0363-0129. doi: 10.1137/15M1049294.
- S. Muntwiler, J. Köhler, and M. N. Zeilinger. MHE under parametric uncertainty – Robust state estimation without informative data, 2023. URL <http://arxiv.org/abs/2312.14049>. arXiv:2312.14049 [cs, eess].
- K. R. Muske and T. A. Badgwell. Disturbance modeling for offset-free linear model predictive control. *J. Proc. Cont.*, 12(5):617–632, 2002.
- B. J. Odelson, A. Lutz, and J. B. Rawlings. Application of autocovariance least-squares methods to laboratory data. Technical Report 2003–03, TWMCC, Department of Chemical Engineering, University of Wisconsin-Madison, September 2003. URL <https://engineering.ucsb.edu/~jbrow/jbrweb-archives/tech-reports/twmcc-2003-03.pdf>.
- B. J. Odelson, M. R. Rajamani, and J. B. Rawlings. A new autocovariance least-squares method for estimating noise covariances. *Automatica*, 42(2):303–308, Feb 2006.

- R. K. Olsson, K. B. Petersen, and T. Lehn-Schiøler. State-Space Models: From the EM Algorithm to a Gradient Approach. *Neural Comput.*, 19(4):1097–1111, 2007. ISSN 0899-7667. doi: 10.1162/neco.2007.19.4.1097.
- S. Oymak and N. Ozay. Revisiting Ho–Kalman-Based System Identification: Robustness and Finite-Sample Analysis. *IEEE Trans. Auto. Cont.*, 67(4):1914–1928, 2022. ISSN 1558-2523. doi: 10.1109/TAC.2021.3083651.
- G. Pannocchia. An economic MPC formulation with offset-free asymptotic performance. *IFAC–P. Online*, 51(18):393–398, 2018. ISSN 2405-8963. doi: 10.1016/j.ifacol.2018.09.332.
- G. Pannocchia and A. De Luca. Performance degradation diagnosis and remedies in offset-free MPC. In *2012 American Control Conference (ACC)*, pages 429–434, Jun 2012. doi: 10.1109/ACC.2012.6314957.
- G. Pannocchia and J. B. Rawlings. Disturbance models for offset-free MPC control. *AIChE J.*, 49(2):426–437, 2003.
- G. Pannocchia, J. B. Rawlings, and S. J. Wright. Conditions under which suboptimal nonlinear MPC is inherently robust. *Sys. Cont. Let.*, 60:747–755, 2011.
- G. Pannocchia, M. Gabiccini, and A. Artoni. Offset-free MPC explained: novelties, subtleties, and applications. *IFAC–P. Online*, 48(23):342–351, 2015. ISSN 2405-8963. doi: 10.1016/j.ifacol.2015.11.304.
- J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren. Benchmark temperature microcontroller for process dynamics and control. *Comput. Chem. Eng.*, 135, Apr 2020. ISSN 0098-1354. doi: {10.1016/j.compchemeng.2020.106736}.
- K. B. Petersen and O. Winther. Explaining slow convergence of EM in low noise linear mixtures. Technical Report 2005-2, Informatics and Mathematical Modelling, Technical University of Denmark, 2005. URL <https://core.ac.uk/download/pdf/13755227.pdf>.
- K. B. Petersen, O. Winther, and L. K. Hansen. On the Slow Convergence of EM and VBEM in Low-Noise Linear Models. *Neural Comput.*, 17(9):1921–1926, 2005. ISSN 0899-7667. doi: 10.1162/0899766054322991.
- L. N. Petersen, N. K. Poulsen, H. H. Niemann, C. Utzen, and J. B. Jørgensen. Comparison of three control strategies for optimization of spray dryer operation. *J. Proc. Cont.*, 57, 2017. ISSN 0959-1524. doi: 10.1016/j.jprocont.2017.05.008.
- D. Piga, M. Forgione, S. Formentin, and A. Bemporad. Performance-Oriented Model Learning for Data-Driven MPC Design. *IEEE Ctl. Sys. Let.*, 3(3):577–582, 2019. ISSN 2475-1456. doi: 10.1109/LCSYS.2019.2913347.
- S. J. Qin. An overview of subspace identification. *Comput. Chem. Eng.*, 30:1502–1513, 2006.

- S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Eng. Pract.*, 11(7):733–764, 2003.
- H. Raghavan, A. K. Tangirala, B. R. Gopaluni, and S. L. Shah. Identification of chemical processes with irregular output sampling. *Control Eng. Pract.*, 14(5):467–480, 2006. ISSN 0967-0661. doi: 10.1016/j.conengprac.2005.01.015.
- M. R. Rajamani and J. B. Rawlings. Estimation of the disturbance structure from data using semidefinite programming and optimal weighting. *Automatica*, 45(1):142–148, 2009.
- M. R. Rajamani, J. B. Rawlings, and S. J. Qin. Achieving state estimation equivalence for misassigned disturbances in offset-free model predictive control. *AIChE J.*, 55(2):396–407, Feb 2009.
- C. R. Rao. Unified theory of linear estimation. *Sankhyā A*, 33(4):371–394, 1971.
- 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 M. J. Risbeck. On the equivalence between statements with epsilon-delta and K-functions. Technical Report 2015–01, TWCCC Technical Report, December 2015. URL <https://engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2015-01.pdf>.
- J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Design, and Computation*. Nob Hill Publishing, Santa Barbara, CA, 2nd, paperback edition, 2020. 770 pages, ISBN 978-0-9759377-5-4.
- R. A. Redner and H. F. Walker. Mixture Densities, Maximum Likelihood and the EM Algorithm. *SIAM Rev.*, 26(2):195–239, 1984. ISSN 0036-1445. doi: 10.1137/1026034.
- T. Ribarits, M. Deistler, and B. Hanzon. An analysis of separable least squares data driven local coordinates for maximum likelihood estimation of linear systems. *Automatica*, 41(3): 531–544, 2005. ISSN 00051098. doi: 10.1016/j.automatica.2004.11.014.
- J. Richalet, A. Rault, J. L. Testud, and J. Papon. Model predictive heuristic control: Applications to industrial processes. *Automatica*, 14:413–428, 1978.
- R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Springer-Verlag, 1998.
- W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, New York, third edition, 1976.
- L. O. Santos and L. T. Biegler. A tool to analyze robust stability for model predictive control. *J. Proc. Cont.*, 9:233–245, 1999.
- L. O. Santos, L. T. Biegler, and J. A. A. M. Castro. A tool to analyze robust stability for constrained nonlinear MPC. *J. Proc. Cont.*, 18(3):383–390, 2008. ISSN 0959-1524. doi: 10.1016/j.jprocont.2007.10.004.

- J. D. Schiller and M. A. Müller. A moving horizon state and parameter estimation scheme with guaranteed robust convergence. *IFAC-P. Online*, 56(2):6759–6764, 2023.
- J. D. Schiller, S. Muntwiler, J. Köhler, M. N. Zeilinger, and M. A. Müller. A Lyapunov Function for Robust Stability of Moving Horizon Estimation. *IEEE Trans. Auto. Cont.*, 68(12):7466–7481, 2023. ISSN 1558-2523. doi: 10.1109/TAC.2023.3280344.
- P. O. M. Sokaert, J. B. Rawlings, and E. S. Meadows. Discrete-time stability with perturbations: Application to model predictive control. *Automatica*, 33(3):463–470, 1997.
- D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle. *Process Dynamics and Control*. John Wiley and Sons, New York, fourth edition, 2017.
- R. H. Shumway and D. S. Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *J. Time Series Anal.*, 3:253–264, 1982.
- R. H. Shumway and D. S. Stoffer. *Time Series Analysis and its Applications*. Springer International Publishing, Cham, Switzerland, fourth edition, 2017.
- L. M. Silverman. Discrete Riccati equations: Alternative algorithms, asymptotic properties, and system theory interpretations. In C. T. Leondes, editor, *Control and Dynamic Systems*, volume 12 of *Control and Dynamic Systems*, pages 313–386. Academic Press, 1976.
- L. Simpson, A. Ghezzi, J. Asprion, and M. Diehl. An Efficient Method for the Joint Estimation of System Parameters and Noise Covariances for Linear Time-Variant Systems. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 4524–4529, Dec 2023.
- L. Simpson, J. Asprion, S. Muntwiler, J. Köhler, and M. Diehl. Parallelizable Parametric Nonlinear System Identification via tuning of a Moving Horizon State Estimator, 2024. arXiv:2403.17858 [math].
- J. Sjöberg, T. McKelvey, and L. Ljung. On the Use of Regularization in System Identification. *IFAC Proceedings Volumes*, 26(2):75–80, 1993. ISSN 1474-6670. doi: 10.1016/S1474-6670(17)48226-7.
- R. Soloperto, J. Köhler, and F. Allgöwer. A Nonlinear MPC Scheme for Output Tracking Without Terminal Ingredients. *IEEE Trans. Auto. Cont.*, 68(4):2368–2375, 2023. ISSN 1558-2523. doi: 10.1109/TAC.2022.3173494.
- E. Sontag and Y. Wang. Lyapunov Characterizations of Input to Output Stability. *SIAM J. Cont. Opt.*, 39(1):226–249, 2000. ISSN 0363-0129. doi: 10.1137/S0363012999350213.
- 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. Notions of input to output stability. *Sys. Cont. Let.*, 38(4-5):235–248, 1999. ISSN 01676911. doi: 10.1016/S0167-6911(99)00070-5.

- J.-G. Sun. Perturbation Theory for Algebraic Riccati Equations. *SIAM J. Matrix Anal. and Appl.*, 19(1):39–65, 1998. ISSN 0895-4798. doi: 10.1137/S0895479895291303.
- Z. Sun, Y. Zhao, and S. J. Qin. Improving industrial MPC performance with data-driven disturbance modeling. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 1922–1927, Dec 2011. doi: 10.1109/CDC.2011.6161469.
- W. Tang, P. Carrette, Y. Cai, J. M. Williamson, and P. Daoutidis. Automatic decomposition of large-scale industrial processes for distributed MPC on the Shell–Yokogawa Platform for Advanced Control and Estimation (PACE). *Comput. Chem. Eng.*, 178:108382, 2023.
- L. W. Taylor. On-Orbit Systems Identification of Flexible Spacecraft. *IFAC Proceedings Volumes*, 18(5):511–516, 1985. ISSN 1474-6670. doi: 10.1016/S1474-6670(17)60611-6.
- A. N. Tikhonov. On the solution of ill-posed problems and the method of regularization. In *Doklady Akademii Nauk*, volume 151, pages 501–504. Russian Academy of Sciences, 1963.
- D. N. Tran, C. M. Kellett, and P. M. Dower. Input-to-state stability with respect to two measurement functions: Discrete-time systems. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 1817–1822, 2015. doi: 10.1109/CDC.2015.7402474.
- J. Umenberger, J. Wågberg, I. R. Manchester, and T. B. Schön. Maximum likelihood identification of stable linear dynamical systems. *Automatica*, 96:280–292, 2018.
- M. Vaccari and G. Pannocchia. A Modifier-Adaptation Strategy Towards Offset-Free Economic MPC. *Processes*, 5(1):2, 2017. doi: 10.3390/pr5010002.
- M. Vaccari, D. Bonvin, F. Pelagagge, and G. Pannocchia. Offset-Free Economic MPC Based on Modifier Adaptation: Investigation of Several Gradient-Estimation Techniques. *Processes*, 9(5):901, 2021. ISSN 2227-9717. doi: 10.3390/pr9050901.
- P. Van Overschee and B. De Moor. N4SID: subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica*, 30(1):75–93, 1994.
- P. Van Overschee and B. De Moor. A unifying theorem for three subspace system identification algorithms. *Automatica*, 31(12):1853–1864, Dec 1995.
- J. Vanbiervliet, B. Vandereycken, W. Michiels, S. Vandewalle, and M. Diehl. The Smoothed Spectral Abcissa for Robust Stability Optimization. *SIAM J. Optim.*, 20(1):156–171, 2009. ISSN 1052-6234. doi: 10.1137/070704034.
- M. Verhaegen. Identification of the deterministic part of MIMO state space models given in innovations form from input-output data. *Automatica*, 30(1):61–74, 1994.
- A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Prog.*, 106(1):25–57, 2006.

- M. Wallace, B. Das, P. Mhaskar, J. House, and T. Salsbury. Offset-free model predictive control of a vapor compression cycle. *J. Proc. Cont.*, 22(7):1374–1386, 2012. ISSN 0959-1524. doi: 10.1016/j.jprocont.2012.06.011.
- M. Wallace, P. Mhaskar, J. House, and T. I. Salsbury. Offset-Free Model Predictive Control of a Heat Pump. *Ind. Eng. Chem. Res.*, 54(3):994–1005, 2015. ISSN 0888-5885. doi: 10.1021/ie5017915.
- T. Westerlund. A digital quality control system for an industrial dry process rotary cement kiln. *IEEE Trans. Auto. Cont.*, 26(4):885–890, 1981. ISSN 0018-9286. doi: 10.1109/TAC.1981.1102738.
- H. White. Maximum Likelihood Estimation of Misspecified Dynamic Models. In T. K. Dijkstra, editor, *Misspecification Analysis*, Lecture Notes in Economics and Mathematical Systems, pages 1–19, Berlin, Heidelberg, 1984. Springer.
- J. C. Willems, P. Rapisarda, I. Markovskiy, and B. L. M. De Moor. A note on persistency of excitation. *Sys. Cont. Let.*, 54(4):325–329, 2005.
- M. Yin, A. Iannelli, and R. S. Smith. Maximum Likelihood Estimation in Data-Driven Modeling and Control. *IEEE Trans. Auto. Cont.*, 68(1):317–328, 2023. ISSN 1558-2523. doi: 10.1109/TAC.2021.3137788.
- C. Yousfi and R. Tournier. Steady-state optimization inside model predictive control. In *Proceedings of American Control Conference*, page 1866, 1991.
- Z. Yuan and J. Cortés. Data-Driven Optimal Control of Bilinear Systems. *IEEE Ctl. Sys. Let.*, 6:2479–2484, 2022. ISSN 2475-1456. doi: 10.1109/LCSYS.2022.3164983.
- M. A. Zagrobelny. *MPC Performance Monitoring and Disturbance Model Identification*. PhD thesis, University of Wisconsin–Madison, December 2014. URL <https://engineering.ucsb.edu/~jbrow/jbrweb-archives/theses/zagrobelny.pdf>.
- M. A. Zagrobelny and J. B. Rawlings. Identifying the uncertainty structure using maximum likelihood estimation. In *American Control Conference*, pages 422–427, Chicago, IL, July 1–3, 2015.
- M. A. Zagrobelny, L. Ji, and J. B. Rawlings. Quis custodiet ipsos custodes? *Annual Rev. Control*, 37:260–270, 2013.
- M. Zamani, B. Ninness, and J. C. Agüero. On Identification of Networked Systems with Time-invariant Topology. *IFAC–P. Online*, 48(28):1184–1189, 2015. doi: 10.1016/j.ifacol.2015.12.292.
- M. Zanon, S. Gros, and M. Diehl. A tracking mpc formulation that is locally equivalent to economic mpc. *J. Proc. Cont.*, 45:30–42, 2016.

- M. Zanon, S. Gros, and M. Diehl. A periodic tracking MPC that is locally equivalent to periodic economic MPC. *IFAC-P. Online*, 50(1):10711–10716, 2017.
- C. Zhang and M. Fu. A revisit to the gain and phase margins of linear quadratic regulators. *IEEE Trans. Auto. Cont.*, 41(10):1527–1530, 1996. doi: 10.1109/9.539438.
- Y. Zhu. System identification for process control: Recent experience and outlook. *IFAC Proceedings Volumes*, 39(1):20–32, Jan 2006. ISSN 1474-6670. doi: 10.3182/20060329-3-AU-2901.00003.