Contributions to Distributionally Robust and Distributed Stochastic Model Predictive Control

Beiträge zur verteilungsrobusten und verteilten stochastischen modellprädiktiven Regelung

Vom Fachbereich Elektrotechnik und Informationstechnik der
Rheinland-Pfälzischen Technischen Universität Kaiserslautern-Landau
zur Verleihung des akademischen Grades
Doktor der Ingenieurwissenschaften (Dr.-Ing.)
genehmigte Dissertation

von
Christoph Mark
geboren in Worms

D 386

Tag der mündlichen Prüfung:	02.02.2023
Dekan des Fachbereichs:	Prof. Dr.-Ing. Marco Rahm
Vorsitzender der Prüfungskommission:	Prof. Dr.-Ing. Daniel Görges
1. Berichterstatter:	Prof. Dr.-Ing. Steven Liu
2. Berichterstatter:	Prof. Dr. Marcello Farina

Acknowledgments

The results presented in this thesis are the outcome of my research activity at the Chair of Control Systems (LRS) at the University of Kaiserslautern. There are numerous people who have supported me and to whom I would like to express my gratitude at this point.

First and foremost, I would like to thank my supervisor Prof. Steven Liu for his continuous support during my doctoral studies and for providing me with a lot of freedom in my research. I also want to thank Prof. Marcello Farina for his interest in my work, the eye-opener that practical applications can still be theoretically challenging, as well as for being a member of my doctoral exam committee. Furthermore, I would like to thank Prof. Daniel Görges for chairing my doctoral exam committee and for his occasional advise over the years.

I would like to thank all the people I encountered at the LRS who contributed to such a stimulating and fun environment. These are Tim Steiner, Joe Ismail, Pedro Dos Santos, Sebastian Bard, Jonas Ulmen, Adrian Herbst, Chen Cai, Xiang Chen, Marco Guerreiro, Kashif Iqbal, Qingshan Pan, Yanhao He, Min Wu, as well as the "oldies" Felix Berkel, Alan Turnwald and Markus Bell. Thanks also goes to Swen Becker for the technical support and the daily "Nachbesprechung", and to Annegret Stabel for her administrative assistance.
A special thanks to my former office mate Tim for representing LRS with me on the fourth floor and pointing out the daily "hammer time", to Joe for the fruitful discussions on MPC and for the fun times we have spent together over the years, and to Pedro, our "Portuguese" who always lifted everyone's spirits with his great impersonations.

A special thanks also goes to my parents. I owe much of what I have accomplished to your constant support and encouragement to continue following my passions. Finally, I would like to thank my girlfriend Katarina for always believing in me, for helping me through the stressful times and for always reminding me that life is more than writing papers.

Contents

Notation

In the following, we define some of the basic notation and acronyms used throughout this thesis.

Abbreviations and acronyms

MPC	model predictive control
DMPC	distributed model predictive control
SMPC	stochastic model predictive control
DSMPC	distributed stochastic model predictive control
DR-MPC	distributionally robust model predictive control
FH-SOCP	finite-horizon stochastic optimal control problem
DRO	distributionally robust optimization
VaR	value-at-risk
CVaR	conditional value-a-risk
i.i.d.	independent and identically distributed
PRS	probabilistic reachable set
DR-PRS	distributionally robust probabilistic reachable set
ICC	individual chance constraint
JCC	joint chance constraint
LMI	linear matrix inequality
ADMM	alternating direction method of multipliers
w.r.t.	with respect to
CCU	central convex unimodal
LQG	linear quadratic gaussian
SDP	semidefinite program
SAA	sample average approximation
SOC	second-order cone

Sets

$\emptyset$	The empty set
$\mathbb{N}$	Set of natural numbers including 0
$\mathbb{R}$	Set of real numbers

$\mathbb{R}_{>0}$	Set of positive real numbers
$\mathbb{R}_{\geq 0}$	Set of nonnegative real numbers
$\mathbb{R}^n$	Set of n-dimensional vectors with real entries
$\mathbb{R}^{m \times n}$	Set of $(m \times n)$-dimensional matrices with real entries

Algebraic Operators

Let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ be vector, $p \geq 1$ be a real number, $P \in \mathbb{R}^{n \times n}$ a quadratic matrix, $Q \in \mathbb{R}^{n \times n}$ a symmetric matrix and $s \in \mathbb{R}$ a scalar.

$\lambda_{\min}(Q)$	Minimum eigenvalue of Q
$\lambda_{\max}(Q)$	Maximum eigenvalue of Q
$Q \succ 0$	Positive definite matrix, $Q \succ 0 \Leftrightarrow v^\top Q v > 0 \, \forall v \neq 0 \Leftrightarrow \lambda_{\min}(Q) > 0$
$Q \succeq 0$	Positive semidefinite matrix, $Q \succeq 0 \Leftrightarrow v^\top Q v \geq 0 \, \forall v \neq 0 \Leftrightarrow \lambda_{\min}(Q) \geq 0$
$\lvert s \rvert$	absolute value
$\lVert v \rVert_p$	l_p vector norm, $\lVert v \rVert_p = \left(\sum_{i=1}^n \lvert v_i \rvert^p \right)^{\frac{1}{p}}$
$\lVert v \rVert_2$	l_2 vector norm (Euclidean norm), $\lVert v \rVert_2 = \sqrt{v_1^2 + \ldots + v_n^2}$
$\lVert v \rVert_Q$	Weighted l_2 vector norm, $\lVert v \rVert_Q = \sqrt{v^\top Q v}$, $Q \succeq 0$
$\lVert v \rVert_\infty$	l_∞ vector norm, $\lVert v \rVert_\infty = \max_{i \in \{1,\ldots,n\}} \lvert v_i \rvert$
$\lVert P \rVert_2$	l_2 induced matrix norm (Spectral norm), $\lVert P \rVert_2 = \sqrt{\lambda_{\max}(P^\top P)}$
$\lVert P \rVert_{\mathrm{F}}$	Frobenius norm of a matrix, $\lVert P \rVert_{\mathrm{F}} = \sqrt{\mathrm{tr}(P^\top P)}$

Set Operators

Let $\mathbb{S}, \mathbb{S}_1, \ldots, \mathbb{S}_M \subseteq \mathbb{R}^n$ be sets and $p \in \mathbb{R}^n$ a vector.

$\lvert \mathbb{S} \rvert$	Cardinality of $\mathbb{S}$
$\mathbb{S}_1 \cup \mathbb{S}_2$	Union of sets, $\mathbb{S}_1 \cup \mathbb{S}_2 = \{s \lvert s \in \mathbb{S}_1 \vee s \in \mathbb{S}_2\}$
$\mathbb{S}_1 \cap \mathbb{S}_2$	Intersection of sets, $\mathbb{S}_1 \cap \mathbb{S}_2 = \{s \lvert s \in \mathbb{S}_1 \wedge s \in \mathbb{S}_2\}$
$\mathbb{S}_1 \oplus \mathbb{S}_2$	Minkowski sum, $\mathbb{S}_1 \oplus \mathbb{S}_2 = \{s_1 + s_2 \lvert s_1 \in \mathbb{S}_1, s_2 \in \mathbb{S}_2\}$
$\mathbb{S}_1 \ominus \mathbb{S}_2$	Pontryagin difference, $\mathbb{S}_1 \ominus \mathbb{S}_2 = \{s_1 \in \mathbb{S}_1 \lvert s_1 + s_2 \in \mathbb{S}_1, \forall s_2 \in \mathbb{S}_2\}$
$\mathbb{S}_1 \times \mathbb{S}_2$	Cartesian product of sets, $\mathbb{S}_1 \times \mathbb{S}_2 = \{(s_1, s_2) \lvert s_1 \in \mathbb{S}_1, s_2 \in \mathbb{S}_2\}$
$\prod_{i=1}^M \mathbb{S}_i$	Cartesian product of M sets, $\prod_{i=1}^M \mathbb{S}_i = \mathbb{S}_1 \times \ldots \times \mathbb{S}_M$
$\mathrm{dist}(p, \mathbb{S})$	Point to set distance from p to $\mathbb{S}$, $\mathrm{dist}(p, \mathbb{S}) \coloneqq \inf_{s \in \mathbb{S}} \lVert p - s \rVert$
$\mathbb{1}_{\mathbb{S}}(p)$	The indicator function $\mathbb{1}_{\mathbb{S}}(p) = 1$, if $p \in \mathbb{S}$ and $\mathbb{1}_{\mathbb{S}}(p) = 0$ if $p \notin \mathbb{S}$.

Probability theory

Let $\mu \in \mathbb{R}^n$, $\Sigma \succ 0$, A, B two events and x, y random vectors in $\mathbb{R}^n$.

$\mathcal{Q}(\mu, \Sigma)$	Parametric distribution with mean μ and covariance matrix Σ
$\mathcal{N}(\mu, \Sigma)$	Multivariate normal distribution with mean μ and covariance matrix Σ
$x \sim \nu$	A random vector x that follows a distribution ν
$x \overset{d}{=} y$	Two random vectors x and y that have the same distribution are equal in distribution
$\mathbb{P}(A)$	The probability of occurrence of event A
$\mathbb{P}(A\|B)$	The conditional probability of occurrence of event A given B
$\mathbb{E}(x)$	The expected value of a random variable x
$\mathbb{E}(x\|A)$	The expectation of x conditioned on A
$\mathbb{E}_\nu(x)$	The expectation of a random variable x w.r.t. distribution ν
$\mathrm{var}(x)$	The variance of x is $\mathbb{E}((x - \mathbb{E}(x))(x - \mathbb{E}(x))^\top)$
$\mathrm{var}(x\|A)$	The variance of x conditioned on A is $\mathbb{E}((x - \mathbb{E}(x\|A))(x - \mathbb{E}(x\|A))^\top\|A)$

Other

Let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ be a vector, $A \in \mathbb{R}^{m \times n}$ a matrix, $x \in \mathbb{R}$ a scalar, $\mathbb{Z}$ the set of integers and $B_1, B_2, \ldots, B_n$ matrices

$[A]_{ij}$	The element in the i-th row and j-th column of matrix A
$[A]_i$	The i-th row of matrix A
$[v]_i$	The i-th element of vector v
$A^\dagger$	The pseudo inverse of a matrix A
$\lceil x \rceil$	The ceiling function $\min(c \in \mathbb{Z}, c \geq x)$
$\mathrm{blkdiag}_{j \in \{1,\ldots,n\}}(B_j)$	A block diagonal matrix with elements $B_1, \ldots, B_n$
$v = \mathrm{col}_{j \in \{1,\ldots,n\}}(v_j)$	A column vector with elements $v_1, \ldots, v_n$

Convex analysis

Let $x, y \in \mathbb{R}^n$ be vectors and $f : \mathbb{X} \to \mathbb{R}$ a real-valued function supported on $\mathbb{X} \subseteq \mathbb{R}^n$.

$(f(x))_+ = \max(0, f(x))$	The positive part of f
$f^*(\theta) \coloneqq \sup_{x \in \mathbb{X}} \theta^\top x - f(x)$	The convex conjugate of f
$\mathrm{dom}(f) = \{x \in \mathbb{X} \| f(x) < \infty\}$	The effective domain of f
$\|x\|_* \coloneqq \sup_{\|y\| \leq 1} x^\top y$	The dual norm of a norm $\|\cdot\|$

1 Introduction

1.1 Motivation

Model Predictive Control (MPC) is an advanced optimization-based control method that relies on repetitive solution of finite-horizon optimal control problems [97, 144]. In closed-loop operation, feedback is generated by implementing only the initial part of the optimized input sequence and repeating the optimization procedure at the next time instant with a shifted prediction horizon. The popularity of MPC is based on its ability to handle arbitrary system dynamics as well as constraints on states and inputs, which is attractive for many applications.

Most of the early work on MPC addresses the nominal regulation task, i.e., setpoint stabilization without disturbances, and is based on the stability theory presented by Mayne et al. [119]. Over the last couple of decades the theory has been extended into several directions, such as robust MPC [118], stochastic MPC (SMPC) [32], distributed MPC (DMPC) [130] and economic MPC [145].

The contribution of this thesis is twofold. In the first part, we extend the literature on SMPC for distributed systems (DSMPC), while the second part addresses a relatively new area called distributionally robust MPC (DR-MPC), which can be viewed as a more realistic approach to SMPC in case of partially known distributional information.

1.1.1 Stochastic MPC for distributed linear systems

The classical literature on MPC typically deals with the centralized setting, i.e., all state measurements of the plant are gathered locally, where then a centralized MPC optimization problem is solved. However, in the domain of large-scale distributed systems this task can be infeasible due to the following reasons:

1. The sensors are spatially distributed and no central coordinator exists.
2. The curse of dimensionality that is associated with a large state dimension [14].

The above two problems have led to the development of DMPC strategies, where the main goal is to decompose the distributed system into several subsystems, each of which has a local MPC controller that solves a smaller scale sub problem. The local controllers are able to share their information with others through a communication network so that a common

objective function can be minimized. Applications can be found in power systems [166], water supply systems [102], irrigation canals [131], building and energy hubs [101], platooning vehicles [172] and wind farms [161].

A DMPC can be roughly categorized into the following components [38]:

- Communication:
  During one sampling interval an *iterative* controller exchanges several times information with its neighbors, while a *non-iterative* controller receives and transmits information only once.
- Attitude:
  A DMPC is said to be *cooperative* if the objective is to minimize a global cost function and *non-cooperative* if each agent intents to minimize a local cost function.
- Update:
  A *sequential* DMPC updates only one subsystem input at each time instant, whereas a *parallel* DMPC updates all subsystem inputs simultaneously.

The vast majority of publications in the field of DMPC considers the nominal case, which is typically an invalid assumption in reality. In many applications it can be observed that the underlying dynamics are corrupted by additive/multiplicative disturbances, leading to the need of stochastic or robust approaches. While robust approaches assume an a-priori known bound of the worst-case disturbance to satisfy the constraints robustly [118], stochastic approaches make use of an underlying model of the disturbance, e.g., the probability distribution, to relax the constraints as so-called chance constraints [121]. Unlike hard constraints, chance constraints only need to be verified with a predefined probability, allowing for a certain frequency of constraint violations. It should be noted that for some applications it is not possible to find a worst-case bound for the disturbance, such as control tasks involving the ambient temperature [76], rendering robust approaches inapplicable to the specific task.

The first part of this thesis deals with uncertainties in form of additive/multiplicative stochastic noise, where the first and second moment and/or the probability distribution are known. In the literature there are basically two approaches how to treat stochastic disturbances. On the one hand, there are *scenario-based approaches* [79, 150], which, at every time instant, sample sufficiently many disturbance realizations in order to approximate a stochastic optimal control problem. The inherent sampling technique makes the scenario-based methods applicable for systems with arbitrary disturbances. However, due to their heavy computational load these methods are still limited to small-scale systems and thus have not been investigated for the purpose of distributed systems. On the other hand, we have *analytical approximation methods* [60, 76, 109, 110], which assume a parametric probability distribution in order to reformulate the stochastic optimal control problem based on the moments of the disturbance. Note that the latter approach reduces the online complexity to a nominal MPC optimization problem, enabling the development of fast and scalable distributed controllers.

The development of the last decades has shown that the complexity of modern control systems is continuously increasing [130]. In addition, these complex interconnections challenge current control algorithms with uncertain dynamics or falsified measurements due to measurement or process noise. Therefore, there is an increasing need for distributed control algorithms capable of dealing with stochastic uncertainties and for systematic approaches on how to propagate stochastic uncertainties in networked systems [121]. In part one of this thesis, we investigate iterative, cooperative and parallel DMPC controllers for stochastic distributed linear systems subject to additive or multiplicative noise, where each proposed DMPC controller respects the following research goals:

- Quantifying the effect of additive or multiplicative uncertainty in a distributed way.
- Distributed controller synthesis and distributed online operation.

In the following, we discuss chapter-wise the usability and purpose of each considered class of control problems.

(i) **Tracking problem with additive noise:**
In Chapter 3, we study a DSMPC for tracking of piece-wise constant output references. This class is important for many practical applications that require online adaption of setpoints unequal to zero [104]. Typical examples include distributed power systems [166], water supply systems [102], energy systems [101] or wind farms [161].

(ii) **Output-feedback regulation with additive noise:**
In Chapter 4, we consider a distributed regulation task under output-feedback with additive noise. This chapter is motivated by the practical aspect that in many control systems the state vector is not fully measurable, which confronts us with the new challenge of how to deal with different sources of uncertainty, i.e., process and measurement noise [61]. Examples can be found in distributed power systems [165] or coordination problems [141].

(iii) **State-feedback regulation with multiplicative noise:**
Chapter 5 deals with a state feedback approach for distributed systems subject to multiplicative noise. Multiplicative noise models are less prominent than their additive counterparts, but they are a very useful tool for representing complex dynamical systems that are difficult to model [68]. A practical example is the control of wind turbines [30], which can be extended to a larger scale by considering wind farms [161].

1.1.2 Distributionally Robust MPC

Distributionally robust MPC is a relatively new area of research that arises from a practical aspect with respect to the applicability of SMPC. In particular, SMPC assumes that the underlying probability distribution or moments are known exactly, which in practice is a rather restrictive assumption, i.e., the true distribution is rarely known and must be estimated from limited data [126]. This is especially problematic if the process of generating

data, i.e., sampling from the true distribution, is costly or time-consuming. In DR-MPC, the assumption of exact knowledge of the distribution is removed by optimizing the stochastic optimal control problem over a class of probability distributions contained in a so-called ambiguity set.

The theory behind DR-MPC traces back to distributionally robust optimization (DRO) [143], which roughly distinguishes the ambiguity sets into two classes. The first class is denoted as *moment-based ambiguity sets* [49], where we assume that the first and second moment of the disturbance are (probabilistically) bounded in a set $\mathcal{P}$ with high confidence. In a control context, these sets can be useful if we model process or measurement noise based on limited samples, cf. Chapter 7. The second class of ambiguity sets are so-called *distance-based ambiguity sets*, which include all probability distributions within a radius ϵ of a nominal distribution, measured in the space of probability distributions. These distance measures include, e.g., the Wasserstein metric [126], Kullback-Leibler divergence [83], total variation distance [52] and several variations of those, see [143] for a recent review. The main advantage of distance-based ambiguity sets is their ability to capture non-parametric probability distributions, which, however, leads to more computationally demanding optimization problems.

DRO has been applied successfully in many practical scenarios, such as optimal power flow [72], portfolio optimization [49], economic dispatch [134] and robot motion control [75], which outperforms the classical stochastic optimization methods based on empirical data in terms of robustness to sample errors [126].

A risk-averse point of view An alternative interpretation of DRO is given by risk-averse optimization [157], where under some mild conditions, i.e., real-valued costs, convex and bounded ambiguity set, DRO is equivalent to minimizing a coherent risk measure [5, 142]. In fact, the dual representation of every coherent risk measure is given by a DRO problem [156]. The notion of risk plays an important role in modern control systems, especially in the data-driven regime, where distributional uncertainty is inevitably present. Therefore, to incorporate risk awareness in the control design, one typically formulates so-called distributionally robust chance constraints, i.e., chance constraints that have to hold for all possible distributions contained in an ambiguity set. However, the resulting feasible set is generally non-convex, which has been addressed in various ways in the literature, e.g., [147] proposes the use of a coherent risk measure called conditional Value-at-Risk (CVaR), which serves as an inner approximation of the chance constrained set. This approach is widely used in the literature in case of Wasserstein ambiguity sets, cf. [45, 111, 112, 126]. An alternative is given by a second-order cone (SOC) constraint reformulation in case of moment-based ambiguity sets, cf. [49, 103, 117]. Note that under some mild conditions on the constraint function, it can be shown that the feasible sets of distributionally robust CVaR constraints and distributionally robust chance constraints are equal [177, Thm. 2.2].

In control applications, the use of CVaR constraints is in fact more reasonable than chance constraints, as the CVaR penalizes not only the frequency of constraint violations, as in

chance constraints, but also the magnitude [45]. For example, a large constraint violation is potentially more harmful to the control system compared to a small one.

In the following, we discuss the usability of each considered class of ambiguity sets.

(i) **DR-MPC with Wasserstein ambiguity sets:**
In Chapter 6, we study two DR-MPCs with Wasserstein ambiguity sets for state feedback regulation with additive uncertainty and unknown distributions. In Section 6.2, we consider the general setting of potentially correlated stochastic processes. This is of practical interest since many control applications do not satisfy the usual assumption of independent and identically distributed (i.i.d.) random variables. Examples are control tasks that include ambient temperature forecasts [76, 112], wind speed data [7] or specifically in power systems by computing an optimal power flow in presence of renewable energy sources [72]. In Section 6.3, we impose the typical i.i.d. assumption on the additive noise. In contrast to Section 6.2, this results in a controller several orders of magnitude faster, i.e., the controller has the complexity of a nominal MPC, which dramatically extends the range of practical applications.

(ii) **DR-MPC with moment-based ambiguity set:**
In Chapter 7, we consider a DR-MPC with moment-based ambiguity sets for linear systems subject to additive i.i.d. sub-Gaussian noise. Note that many distributions of practical interest are sub-Gaussian, e.g., any bounded random variable [167, Sec. 2.5]. From a computational viewpoint, this approach is comparable to the Wasserstein approach proposed in Section 6.3 and allows for control applications with fast dynamics. Finally, in Chapter 8, we apply the DR-MPC as a supervisory controller for a wind farm.

1.2 Related work

In this section, we review some of the related literature on SMPC and DR-MPC.

1.2.1 Stochastic MPC for linear systems

In the following, we provide a brief overview of the related literature on SMPC and compare different aspects, such as recursive feasibility, chance constraints and closed-loop guarantees. This overview is by no means exhaustive and we refer the reader to some recent review papers [57, 121] for a more in-depth comparison. We have selected some key publications to highlight some aspects of SMPC and also included the author's publications.

Recursive feasibility

Recursive feasibility is one of the most important properties of an MPC optimization problem, which ensures that the optimization problem will remain feasible for every possible initial state [144]. In context of SMPC, this issue can be addressed in several ways [57], e.g., under the assumption of a bounded uncertainty, one can use a robust MPC inspired constraint tightening [106]. In presence of an unbounded uncertainty, such a robust bound cannot be found and thus the feasibility issue is more complicated. One possible solution is proposed by Farina et al. [60, 61], where a binary initialization constraint is introduced with the intention to choose between the feedback initialization (state measurement) or a backup strategy, i.e., the shifted optimal solution from the previous time step. By doing so, the chance constraints are inherently verified as conditional probabilistic constraints and thus, no closed-loop guarantees can be given.

An alternative, the so-called *direct-feedback* approach based on probabilistic reachable sets (PRS), was proposed by [77], where an additional central convex unimodality assumption of the noise distribution ensures conservative satisfaction of the chance constraint in closed-loop. Conservatism is introduced by conditioning the chance constraints on the current state, which are essentially implemented as hard constraints [78], cf. Remark 2.3. The same authors proposed a second alternative initialization scheme based on *indirect-feedback*, where the chance constraints are conditioned on the closed-loop error [76]. Recursive feasibility is trivially verified since the MPC optimization problem is always initialized with shifted nominal state, while feedback is introduced only in the cost function. The authors of [93] and [153] use a PRS-based approach in conjunction with an interpolated initial constraint that constrains the initial state on a line between the state measurement and the predicted nominal state from the previous time step. Similar to [77], this ensures closed-loop chance constraint satisfaction and recursive feasibility. Finally, Cannon et al. [31] ensures recursive feasibility with the concept of invariance with probability p, where an alternative optimization problem is solved in case of infeasibility.

Other notable work includes scenario-based SMPC algorithms [151], where recursive feasibility is either assumed or enforced by soft constraints, leading to scenario optimization guarantees [28]. Exceptions to this are the works [79, 128], where the authors make use of an indirect-feedback initialization to decouple the stochastic uncertainty from the nominal MPC optimization problem. In this case, recursive feasibility can be established and the chance constraints hold in closed-loop with high probability.

Chance constraint

The main advantage of SMPC is the use of chance constraints, which can be separated in individual chance constraints (ICC) or joint chance constraints (JCC), cf. Section 2.2 or [57] for a recent review paper. The vast majority of proposed SMPC controllers considers ICCs [29, 31, 33, 59, 60, 115], which are in general easier to deal with compared to JCCs [77, 109, 110, 137]. In analytical approximation frameworks, the ICCs can be reformulated via concentration inequalities, such as the Chebyshev-Cantelli inequality [60], or in case of JCCs by making use of the two-sided Chebyshev inequality [77]. Another possibility to treat JCCs is to use the union bound to approximate the JCCs by a set of ICCs, which typically results in conservative constraint sets [106]. In most of the literature, constraint satisfaction in analytical frameworks is only enforced in prediction, i.e., the chance constraints are conditioned on the most recent state feedback. This, however, does not ensure that the closed-loop system verifies the chance constraints, which typically requires stronger assumptions, cf. [77, 93, 153]. Finally, the authors of [79, 128] propose scenario optimization techniques to approximate arbitrary constraints via non-symmetric PRS (ICCs) that result in non-conservative closed-loop constraint satisfaction.

Distributed stochastic MPC

In the following, we compare some of the key publications in SMPC for distributed systems, where Table 1.1 focuses on the SMPC aspects and Table 1.2 on the different DMPC architectures. In both tables, the double horizontal line separates the related publications (upper part) from the author's publications (lower part).

From Table 1.1, it can be seen that most related work focuses on the state feedback case considering ICCs. Moreover, the majority of proposed DMPCs verifies chance constraints only in prediction. In general, analytical approximation methods are widely used for DSMPC [46–48, 56, 58], with the exception of [128], which uses a hybrid approach that unifies scenario-based PRS with analytical approximations to ensure recursive feasibility. Notably, this allows the authors to give scenario-based closed-loop chance constraint guarantees. Finally, [148] uses a scenario-based approach for systems subject to arbitrary additive and multiplicative uncertainty, as well as JCCs. Note that recursive feasibility cannot be established and thus, the chance constraints are verified only in prediction.

Next, we compare some of the papers regarding their distributed structure. From Table 1.2,

it can be seen that some of the previous work on DSMPC considers the non-iterative setting to reduce the communication overhead of the online algorithm [46–48, 56, 58]. As a result, the MPC optimization becomes suboptimal, which deteriorates closed-loop performance. Furthermore, the couplings of the subsystems were usually considered only in the constraints or in the dynamics. In addition, the design procedures proposed by [46–48] require a central coordinating node, such that the controller ingredients can be synthesized. In comparison, we have proposed several DSMPC controllers that consider dynamic and constraint couplings, with emphasis on a distributed design, cf. Section 1.3 for more details. Furthermore, we consider an iterative DSMPC with parallel updating local MPCs. The main drawback of [109, 110, 113, 115, 128] is the increased communication load.

Table 1.1: Comparison regarding SMPC.

	Disturbance			Chance constraint		
Paper	Additive	Multiplicative	Feedback	Type	Guarantee	MPC type
[46]	bounded	bounded	output	ICC	Predictive	Analytical
[47]	bounded	-	state	ICC	Predictive	Analytical
[56]	unbounded	-	state	ICC	Predictive	Analytical
[58]	unbounded	-	state	ICC	Predictive	Analytical
[48]	bounded	bounded	state	ICC	Predictive	Analytical
[148]	unbounded	unbounded	state	JCC	Predictive	Scenario
[128]	unbounded	-	state	ICC	Closed-loop	Scenario
[109, 113]	unbounded	-	state	JCC	Closed-loop	Analytical
[110]	unbounded	-	output	JCC	Closed-loop	Analytical
[115]	-	unbounded	state	ICC	Predictive	Analytical

Table 1.2: Comparison regarding distributed MPC architecture.

Paper	Communication	Update	Design	Coupling
[46]	non-iterative	sequential	central	Constraints
[47]	non-iterative	sequential	central	Constraints
[56]	non-iterative	parallel	distributed	Dynamics
[58]	non-iterative	parallel	distributed	Constraints
[48]	non-iterative	sequential	central	Constraints
[148]	non-iterative	parallel	-	Dynamics
[128]	iterative	parallel	distributed	Dynamics, Constraints
[109, 113]	iterative	parallel	distributed	Dynamics, Constraints
[110]	iterative	parallel	distributed	Dynamics, Constraints
[115]	iterative	parallel	distributed	Dynamics, Constraints

1.2.2 Distributionally robust MPC

In the following, we discuss some of the related work on DR-MPC.

Ambiguity set

The ambiguity set is one of the most distinctive elements of a DR-MPC since its choice directly impacts the complexity of the underlying optimization problem. The literature proposes several different approaches, where a large part considers moment-based ambiguity sets [103, 117, 162, 164]. These sets model distributional uncertainty in the first and/or second moment by introducing upper bounds for the mean and the covariance matrix [49]. An approach closely related is proposed by Coppens et al. [43], where the ambiguity set is formulated for conically representable risk [35]. Another distinguishing feature in the selection of moment-based ambiguity sets is whether the set is data-dependent or whether the size is fixed a priori. Many of the DR-MPCs with moment-based ambiguity sets assume that the true distribution is contained in the ambiguity set with probability one, e.g., [85, 103, 108, 162, 164], while only few publications consider data-driven ambiguity sets [43, 117], where the true distribution belongs to the ambiguity set with high confidence. Thus, the latter two approaches can reduce conservatism by successively adjusting the ambiguity radius by collecting more data.

Besides moment-based ambiguity sets, the use of Wasserstein ambiguity sets is also prominent [45, 74, 108, 111, 112, 174, 176], where the idea is to optimize over all distributions close to a nominal distribution. Some other notable work considers Dirichlet Process Mixture Models [133] or the total-variation distance [52].

Theoretical properties

Most publications on DR-MPC only focus on reformulating the MPC optimization problem, with theoretical guarantees such as recursive feasibility or convergence to an average performance bound not discussed in detail or not at all. Since DR-MPC is basically an extension of SMPC, the same principles apply to ensure recursive feasibility, cf. the discussion in the previous section. To briefly summarize: [103] uses a binary initialization scheme, [111, 112] use indirect-feedback, [108] uses a soft-constrained set, [43] relies on state feedback initialization due to bounded disturbances and finally, [117] uses an interpolated initial constraint. Regarding the chance constraints, several approximation methods have been proposed in the literature, while the closed-loop constraint guarantees are usually not addressed. For comparison, we refer to Table 1.3.

In addition to the classic MPC frameworks, several recent publications tend to incorporate online learning, e.g., for Markovian switching systems [154], iterative tasks [176], motion planning [73] and autonomous driving [155].

Table 1.3: Comparison of related literature on DR-MPC.

			Chance constraint		MPC guarantees		
Paper	Dynamics	Disturbance	Reformulation	Guarantee	Rec. feasible	Perf. bound	Ambiguity set
[164]	linear	additive, i.i.d.	CVaR	-	-	-	Moment-based
[108]	linear	additive, bounded ,i.i.d.	Robust	-	yes	-	Moment-based
[85]	linear	additive, i.i.d.	CVaR, conic	-	-	-	Moment-based
[43]	linear	additive, bounded, i.i.d	conic	predictive	yes	-	Moment-based
[173]	linear	additive, bounded, i.i.d.	conic	predictive	yes	-	Wasserstein
[103]	linear	additive, i.i.d.	CVaR, conic	predictive	yes	yes	Moment-based
[45]	Non-param.	i.i.d.	CVaR	predictive	-	-	Wasserstein
[133]	linear	additive, bounded, i.i.d	CVaR	predictive	yes	-	DPMM
[52]	linear	additive, discrete, i.i.d.	CVaR	predictive	-	-	Total-variation
[174]	nonlinear	additive, bounded, i.i.d.	conic	-	-	-	Wasserstein
[162]	linear	additive, i.i.d	conic	predictive	yes	yes	Moment-based
[124]	linear	additive, i.i.d	CVaR	predictive	-	-	Wasserstein
[111]	linear	additive, i.i.d.	CVaR/VaR	closed-loop	yes	-	Wasserstein
[112]	linear	additive, correlated	CVaR	predictive	yes	-	Wasserstein
[117]	linear	additive, i.i.d.	conic	predictive	yes	yes	Moment-based

1.3 Contributions and outline

In the following, we detail the outline of this thesis and summarize the contributions.

Chapter 2: Preliminaries

In this chapter, we provide an introduction to some notions of probability theory and stochastic discrete time-invariant systems. In addition, we discuss common SMPC designs and detail a PRS-based SMPC design throughout this section.

Part I: Distributed Stochastic Model Predictive Control

Chapter 3: Tracking of piece-wise constant references

In this chapter, we present a DSMPC framework for tracking of piece-wise constant references for dynamically coupled distributed systems with neighbor-to-neighbor communication and (unbounded) additive stochastic noise. Chance constraints are treated with the concept of distributed PRS, for which we propose two analytical design methods. The first one relies on the solution of a central linear matrix inequality (LMI), while the second one uses an iterative distributed update scheme that requires only neighbor-to-neighbor information exchange. The DSMPC optimization problem simultaneously optimizes the inputs and the steady-state tracking targets, where we use the consensus alternating direction method of multipliers (ADMM) [22] to obtain a distributed algorithm. Due to the unboundedness of the noise distribution, recursive feasibility of the MPC problem cannot be achieved by constraint tightening, e.g., as in robust MPC [39]. Therefore, on the one hand, a backup strategy (the shifted optimal solution) is necessary to ensure this fundamental property, while on the other hand, the convergence threshold of the ADMM algorithm makes the state trajectory inaccurate, which affects the feasibility of the shifted optimal solution. The novelty of the proposed approach is the incorporation of the ADMM threshold in the controller design (similar to but different from [92]), which ensures feasibility of the backup strategy and satisfaction of the chance constraints in closed-loop. The proposed method is validated for two numerical examples of a distributed system with 50 double integrators in a chain graph and a four-tank benchmark system. This chapter is based on [113].

Chapter 4: Output-feedback regulation with additive noise

In this chapter, we develop a stochastic output-feedback MPC scheme for distributed systems with additive noise. The underlying DSMPC optimization problem is reduced to a quadratic program, which we opt to solve via distributed optimization (similar to Chapter 3). Chance constraints are treated with the concept of distributed PRS, which we extend

from the state feedback to the output-feedback case and for which we propose two analytical design methods. Under an exact feasibility assumption, the MPC optimization problem is proven to be recursively feasible with guaranteed (conservative) closed-loop chance constraint satisfaction and asymptotic convergence to an average cost bound. This assumption can be removed by considering a slightly modified version of the DSMPC algorithm from Chapter 3. In Section 4.3, we extend the direct output-feedback design to the indirect output-feedback case, where we propose an alternative distributed scenario-based PRS design. Using results from scenario optimization [28], we can construct distributed non-symmetric PRS that result in non-conservative closed-loop constraint satisfaction. A numerical example is used to contrast the three distributed PRS designs and discuss the effects on control performance. The results of this chapter have partly been presented in [110].

Chapter 5: Regulation problem with multiplicative noise

We propose a DSMPC algorithm for distributed linear systems subject to individual chance constraints and multiplicative noise. Similar to the previous chapters, we use dual decomposition to obtain a fully parallelizable DSMPC optimization problem that scales with the system dimension and can be solved efficiently with distributed consensus ADMM. In addition, we provide a fully distributed synthesis method for distributed linear feedback controllers and the distributed terminal ingredients. Thus, the MPC synthesis and the online MPC algorithm both do not rely on a central coordination node. Recursive feasibility of the MPC optimization problem is ensured by adopting two alternative control policies. In a numerical example, we illustrate the scalability of the proposed DSMPC algorithm for different system sizes and compare the closed-loop performance for different parameters. This chapter is based on [115].

In summary, the main contributions of Part I are the following:

- Development of a DSMPC algorithm using inexact dual consensus ADMM for tracking of piece-wise constant output references – This controller can easily be amended to the regulation task, since this forms a special case of tracking.
- We propose a direct and indirect output-feedback DSMPC algorithm that is amendable to distributed optimization.
- Development of various analytical and scenario-based distributed PRS design methods.
- We extend the existing literature on SMPC with multiplicative uncertainty to the distributed case, where we propose distributed design procedures and a fully parallelizable DSMPC algorithm.

Part II: Distributionally Robust Model Predictive Control

Chapter 6: Wasserstein distributionally robust Model Predictive Control

In this chapter, we present two DR-MPC strategies with Wasserstein ambiguity sets for linear systems subject to additive stochastic uncertainty. In Section 6.2, we initially propose a scenario-based indirect-feedback DR-MPC that allows for use of correlated stochastic processes. We investigate linear and nonlinear tube controllers and derive theoretical guarantees for recursive feasibility, distributionally robust performance and distributionally robust chance constraint satisfaction. The performance of the controller for different Wasserstein radii is demonstrated on a four-room temperature regulation task with correlated ambient temperature.

In Section 6.2, the distributional assumptions are strengthened, which allows us to derive an analytical DR-MPC scheme. By using an indirect-feedback initialization, the closed-loop error is decoupled from the prediction dynamics, enabling the formulation of a distributionally robust PRS for the closed-loop error. This inherently renders the chance constraints conditioned on the closed-loop error, resulting in closed-loop chance constraint satisfaction. We provide details for arbitrary ambiguity sets and give a concrete design in case of Wasserstein ambiguity sets. In a numerical example, we highlight the out-of-sample performance and confidence of the DR-PRS for various sample sizes and Wasserstein radii. This chapter is based on [111] and [112].

Chapter 7: Moment-based distributionally robust Model Predictive Control

In this chapter, we propose a DR-MPC with data-driven moment-based ambiguity sets for linear systems with additive i.i.d. noise. Under a sub-Gaussian assumption on the noise distribution, we derive an explicit number of samples to ensure a user-defined confidence level, such that the true distribution belongs to the ambiguity set with high probability. In contrast, for Wasserstein ambiguity sets (Chapter 6) this is not possible, since the Wasserstein radius still depends on some unknown parameters of the true distribution. Therefore, we can estimate the data-driven Wasserstein radius empirically using machine learning tools, while in the case of moment-based ambiguity sets, the additional sub-Gaussian assumption allows us to be more rigorous. We use a simplified affine disturbance feedback parameterization to analytically reformulate the distributionally robust cost function, while the chance constraints are cast as SOC constraints. The MPC optimization problem is proven to be recursively feasible, while the closed-loop performance converges to an asymptotic average bound. In a numerical example, we compare closed-loop performance and constraint satisfaction for different sample sizes. In addition, we quantitatively investigate the impact of unmodeled disturbances on constraint satisfaction and performance. This chapter is based on [117].

Chapter 8: Distributionally robust control of a wind farm

In this chapter, we apply the moment-based DR-MPC scheme from Chapter 7 to a wind farm, where the goal is to supervise and coordinate several wind turbines with the common goal of power tracking and fatigue load minimization. We will introduce an auto-regressive moving average model to predict locally the turbulent wind speed, while the optimization problem is robustified using a moment-based ambiguity set. These sets are parameterized for different mean wind speed and turbulence characteristics of the current wind profile, which changes on a slow-scale within ten minutes up to one hour. The control algorithm is tested using the simulation environment SimWindFarm [71], which serves as a nonlinear reference model in Matlab Simulink. This chapter is based on [116].

In summary, the main contributions of Part II are the following:

- Development of scenario-based and analytical indirect-feedback DR-MPC schemes for correlated and i.i.d. additive disturbances using Wasserstein ambiguity sets.
- Introduction of distributionally robust PRS for general ambiguity sets, where we propose a design method using Wasserstein ambiguity sets.
- We extend the literature on moment-based DR-MPC.
- We propose a DR-MPC as a supervisory controller for wind farms.

1.4 Publications of the Author

The following are the author's publications that were published or submitted during his time as a doctoral candidate.

Journal publications

[113] C. Mark and S. Liu. "Stochastic Distributed Predictive Tracking Control Under Inexact Minimization". In: *IEEE Transactions on Control of Network Systems* 8.4 (2021), pp. 1892–1904.

[115] C. Mark and S. Liu. "A stochastic MPC scheme for distributed systems with multiplicative uncertainty". In: *Automatica* 140 (2022), p. 110208.

[117] C. Mark and S. Liu. "Recursively Feasible Data-Driven Distributionally Robust Model Predictive Control With Additive Disturbances". In: *IEEE Control Systems Letters* 7 (2023), pp. 526–531.

Conference publications

[109] C. Mark and S. Liu. "Distributed Stochastic Model Predictive Control for dynamically coupled Linear Systems using Probabilistic Reachable Sets". In: *Proc. European Control Conf. (ECC)*. 2019, pp. 1362–1367.

[110] C. Mark and S. Liu. "A stochastic output-feedback MPC scheme for distributed systems". In: *Proc. American Control Conf. (ACC)*. extended version: arXiv:2001.10838. 2020, pp. 1937–1942.

[111] C. Mark and S. Liu. "Stochastic MPC with Distributionally Robust Chance Constraints". In: *Proc. 21st IFAC World Congress*. extended version: arXiv:2005.00313. 2020, pp. 7136–7141.

[114] C. Mark and S. Liu. "Stochastic Model Predictive Control for tracking of distributed linear systems with additive uncertainty". In: *Proc. European Control Conf. (ECC)*. extended version: arXiv:2103.01087. 2021, pp. 216–221.

Preprints

[112] C. Mark and S. Liu. "Data-driven distributionally robust model predictive control: An indirect feedback approach". In: *arXiv preprint arXiv:2109.09558* (2021). Submitted to International Journal of Robust and Nonlinear Control (2022).

[116] C. Mark and S. Liu. "Distributionally robust model predictive control for wind farms". In: *arXiv preprint arXiv:2303.03276* (2023). Accepted for presentation at the 22nd IFAC World Congress.

2 Preliminaries

In this chapter, we will discuss some preliminary results on stochastic MPC, which we will use frequently in this thesis. In Section 2.1, we introduce some concepts from probability theory, which formalizes the idea of a probability space, random variables and expected values. Afterwards, in Section 2.2, we introduce nonlinear stochastic systems followed by a brief discussion on stochastic MPC for nonlinear systems to illustrate the theoretical challenges. The chapter concludes with the Section 2.3, where we focus on linear stochastic systems for which rigorous theoretical results can be obtained. In particular, we derive and contrast two popular stochastic MPC approaches.

2.1 Probability theory

A probability space is defined by the triplet $(\Omega, \mathcal{A}, \mathbb{P})$, where Ω is the sample space, $\mathcal{A}$ the σ-algebra on Ω and $\mathbb{P}$ the probability measure on $(\Omega, \mathcal{A})$. Throughout this thesis, we consider the σ-algebra to be the Borel σ-algebra.

Definition 2.1 (Borel σ-algebra, [89, Def. 1.21])**.** *Let (Ω, τ) be a topological space. The σ-algebra*

$$\mathcal{B}(\Omega) \coloneqq \sigma(\tau)$$

that is generated by the open sets is called the Borel σ-algebra on Ω. The elements $A \in \mathcal{B}(\Omega)$ are called Borel measurable sets.

In context of this thesis, we consider continuous random variables on $\mathbb{R}^n$. Thus, we specify the set of open subsets of $\mathbb{R}^n$ as $\mathcal{O} = \{A \subset \mathbb{R}^n \mid A \text{ is open}\}$, such that $\mathcal{B}(\mathbb{R}^n) = \sigma(\mathcal{O})$ generates the Borel σ-algebra, cf. [89, Thm 1.23].

In the following, we formalize the concept of a random variable, which can be understood as a measurable map from Ω to a space of possible observations $\mathbb{W}$. The probability of the possible outcomes are described by the distribution of the corresponding random variable, which is the pre-image of the probability measure $\mathbb{P}$ under the measurable map w.

Definition 2.2 (Random variables [89, Def. 1.102])**.** *Let $(\mathbb{W}, \mathcal{F})$ be a measurable space and let $w : \Omega \to \mathbb{W}$ be a measurable map.*

The map w is called a random variable. If $(\mathbb{W}, \mathcal{F}) = (\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$, then w is a real random variable. For any set $F \in \mathcal{F}$, we denote $\{w \in \mathcal{F}\} \coloneqq w^{-1}(F)$ and $\mathbb{P}(w \in F) \coloneqq \mathbb{P}(w^{-1}(F))$.

Definition 2.3 (Distribution of a random variable [89, Def. 1.103])**.** *Let $w : \Omega \to \mathbb{W}$ be a random variable that maps from the probability space $(\Omega, \mathcal{A}, \mathbb{P})$ to a measurable space $(\mathbb{W}, \mathcal{F})$.*

1. *The probability measure $\mathbb{P}_{\mathrm{w}} = \mathbb{P} \circ w^{-1}$ is called the distribution (push-forward measure) of w.* [1]
2. *We write $w \sim \mathbb{P}_{\mathrm{w}}$ and say that w has distribution $\mathbb{P}_{\mathrm{w}}$.*

Expected value

In the following, we introduce the basic concepts of expectations of random variables, where we start with the fundamental definition of μ-integrable functions.

Definition 2.4 (μ-integrable function [89, Def. 4.7])**.** *A measurable function $f : \Omega \to \mathbb{R} \cup \{-\infty, +\infty\}$ is called μ-integrable if $\int_\Omega |f(\omega)| d\mu(\omega) < \infty$. We write*

$$\mathcal{L}^1(\mu) \coloneqq \mathcal{L}^1(\Omega, \mathcal{A}, \mu) \coloneqq \left\{ f : \Omega \to \mathbb{R} \cup \{-\infty, +\infty\} : f \text{ is measurable and } \int |f| d\mu < \infty \right\}.$$

where the integral is a generalized Lebesgue integral, cf. [89, Def. 4.4].

Equipped with the basic integral formulation, we can define the expected value of a random variable $w \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbb{P})$ with $w \sim \mathbb{P}_{\mathrm{w}}$ and $w : \Omega \to \mathbb{W}$ as follows

$$\mathbb{E}_{\mathbb{P}}(w) \coloneqq \int_\Omega w(\omega) d\mathbb{P}(\omega) = \int_{w(\Omega)} \epsilon \, d\mathbb{P}_{\mathrm{w}}(\epsilon),$$

where the second equality uses the change-of-variables [16, Thm 16.13] with $w(\Omega) = \mathbb{W}$ being the image of Ω under the random variable w, i.e., $w(\Omega) \coloneqq \{\epsilon \,|\, \epsilon = w(\omega), \omega \in \Omega\}$. This can be generalized to the expected value under a measurable map $g : \mathbb{W} \to \mathbb{R}^n$, such that

$$\mathbb{E}_{\mathbb{P}}(g(w)) \coloneqq \int_\Omega g(w(\omega)) d\mathbb{P}(\omega) = \int_{\mathbb{W}} g(\epsilon) d\mathbb{P}_{\mathrm{w}}(\epsilon). \tag{2.1}$$

Throughout this thesis, we consider continuous random variables that have a probability density function $f_{\mathrm{w}} : \Omega \to [0, \infty]$. Thus, we can simplify (2.1) so that

$$\mathbb{E}_{\mathbb{P}}(g(w)) = \int_{\mathbb{W}} g(\epsilon) f_{\mathrm{w}}(\epsilon) d\epsilon. \tag{2.2}$$

[1] The symbol $\circ$ denotes the composition of two functions.

Remark 2.1. *In the first part of this thesis, we consider random variables w for which the true distribution (push-forward measure) $\mathbb{P}_{\mathrm{w}}$, cf. Def. 2.3, is known exactly, or at least the first two moments of the random variable are known. In the second part, we remove this assumption and simply assume that the true distribution belongs with high probability to a so-called ambiguity set $\mathcal{P}$, i.e., $\mathbb{P}(\mathbb{P}_{\mathrm{w}} \in \mathcal{P}) \geq 1 - \beta$ for some confidence level $\beta \in (0, 1)$.*

Consider a random variable $w \in \mathcal{L}^1(\Omega, \mathcal{A}, \mathbb{P})$ and a sub σ-algebra $\mathcal{F} \subset \mathcal{A}$. We define the conditional expectation as follows.

Definition 2.5 (Conditional expectation [89, Def. 8.11])**.** *A random variable $y = \mathbb{E}(w \,|\mathcal{F})$ is called a conditional expectation of w given $\mathcal{F}$, if:*

(i) y is $\mathcal{F}$-measurable.

(ii) For any $A \in \mathcal{F}$, we have $\int_A w(\omega) d\mathbb{P}(\omega) = \int_A y(\omega) d\mathbb{P}(\omega)$.

For $B \in \mathcal{A}$, the conditional probability of B given $\mathcal{F}$ is defined as $\mathbb{P}(B \mid \mathcal{F}) \coloneqq \mathbb{E}(\mathbb{1}_B \mid \mathcal{F})$, where $\mathbb{1}_B$ is the indicator function of the set B.

2.2 Stochastic systems

Consider the discrete-time dynamical system

$$x(k+1) = f(x(k), u(k), w(k)) \quad \forall k \in \mathbb{N}, \tag{2.3}$$

where $x \in \mathbb{X} \subseteq \mathbb{R}^n$ is the state, $u \in \mathbb{U} \subseteq \mathbb{R}^m$ the control input, $w \in \mathbb{W} \subseteq \mathbb{R}^n$ the stochastic uncertainty (random variable) and $f : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \to \mathbb{X}$ the dynamic.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space for an infinite sequence $w_\infty : \Omega \to \mathbb{W}^\infty$ of random variables $w(k)$, i.e., $w_\infty = \{w(k)\}_{k=0}^{\infty}$ is a stochastic process, and define the subsequence $w_k : \Omega \to \mathbb{W}^k$ of w_∞ as $w_k = \{w(t)\}_{t=0}^{k-1}$. Let $(\mathcal{F}_0, \mathcal{F}_1, \ldots)$ denote the natural filtration of the sequence w_∞, where the sub σ-algebra $F_k \subset \mathcal{F}$ contains all sets $\{\omega \in \Omega | w_k(\omega) \in F_k\}$ for $F_k \in \mathcal{F}_k \coloneqq \mathcal{B}(\mathbb{W}^k)$. In view of this, also the state $x(k)$ for $k \geq 1$ is a stochastic process defined on $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\mathcal{F}_k$, i.e., the state $x(k)$ is $\mathcal{F}_k$-measurable for all $k \in \mathbb{N}$.

i.i.d. assumption Throughout this thesis, we pose different assumptions on the random sequence w_∞. However, for the sake of introduction we focus only on the independent and identically distributed case, where each random variable $w(k) : \Omega \to \mathbb{W}$ is i.i.d. in time for all $k \in \mathbb{N}$, i.e., each random variable has a known and equivalent probability measure $\mu_{\mathrm{w}} : \mathcal{B}(\mathbb{W}) \to [0, 1]$ defined such that $\mu_{\mathrm{w}}(F) = \mathbb{P}(\{\omega \in \Omega : w(k; \omega) \in F\})$ for all $F \in \mathcal{B}(\mathbb{W})$. In addition, we assume that the expectation is zero and the second moment of $w(k)$ is finite with known covariance matrix $\Sigma_{\mathrm{w}} = \mathrm{var}(w(k)) \succ 0$.

In SMPC, we need to predict the future evolution of the state process in expectation. To this end, we assume that at time $k = 0$ the initial state $x(0)$ is deterministic and define the conditional expectation subject to the filtration $\mathcal{F}_k$ at time k as

$$\mathbb{E}(x(k+t+1) \mid \mathcal{F}_k) = \mathbb{E}(f(x(k+t), u(k+t), w(k+t)) \mid x(k)) \quad \forall t \in \mathbb{N}.$$

The expression defines a t-step ahead prediction with information available at time k, i.e., conditioned on the filtration $\mathcal{F}_k$ or similarly on $x(k)$.

Chance constraints In many control applications, the system dynamic (2.3) needs to satisfy a set of constraints, which in the most general way can be written as

$$g(x(k), u(k), w(k)) \leq 0, \tag{2.4}$$

where $g : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \to \mathbb{R}^r$ denotes the constraint function and the inequality holds element-wise. In view of the random variable $w(k)$, satisfaction of (2.4) in form of hard constraints is not always possible, e.g., if the support set $\mathbb{W} = \mathbb{R}^n$ is unbounded. In this case, we introduce chance constraints of the form

$$\mathbb{P}(g(x(k), u(k), w(k)) \leq 0) \geq p, \tag{2.5}$$

where $p \in (0, 1)$ denotes the level of constraint satisfaction. At this point, a first distinction of the chance constraint (2.5) is possible:

(i) If $g \mapsto \mathbb{R}$ is a scalar function, then the chance constraint (2.5) is called an individual chance constraint (ICC).

(ii) If $g \mapsto \mathbb{R}^r$ is a vector-valued function, i.e., $r \geq 2$, then the chance constraint (2.5) is called a joint chance constraint (JCC).

Although JCCs are more natural from a control perspective, their evaluation is in general cumbersome and involves solving a multivariate integral. In addition, the feasible set is typically non-convex, which requires further convex approximations, cf. [132]. In case of ICCs, we can impose several chance constraints at once by defining the vector-valued function $g(x(k), u(k), w(k)) = [g_1(x(k), u(k), w(k)), \ldots, g_r(x(k), u(k), w(k))]^\top$ with $g_i : \mathbb{X} \times \mathbb{U} \times \mathbb{W} \to \mathbb{R}$ for $i = 1, \ldots, r$ and interpret the chance constraint (2.5) element-wise.

2.2.1 Discussion of stochastic MPC for nonlinear systems

In the following, we briefly discuss some of the related literature on SMPC for nonlinear systems of the form (2.3) with potentially nonlinear constraints (2.4).

Uncertainty propagation One of the major issues in nonlinear systems is that the superposition principle does not hold. Thus, propagating the uncertainty through the system

dynamics cannot be performed as in linear systems. Several methods have been proposed over the last decade, with some of the early work considering polynomial chaos expansion [122], Gaussian Process models [23] or the Fokker-Planck equation [24], see also a recent survey paper [121] and the references therein. Recently, another line of research has enabled a different view on nonlinear SMPC with the concept of incremental stability [95], which allows to derive a constraint tightening based on the inverse cumulative distribution function of the bounded random variable [152]. A related approach is proposed by [149], which uses Lipschitz arguments to bound the uncertainty.

Stability The second issue is related to asymptotic stability of the closed-loop system, where most of the original work on nonlinear SMPC relies on stochastic stability results developed by Kushner [100]. Since then, much effort has been put into refining and extending stochastic Lyapunov functions to achieve global asymptotic stability, cf. [63, 66]. The concept of input-to-state stability (ISS) was also extended to the stochastic case (SISS), where SISS-Lyapunov functions were established [163]. The previous work on stochastic stability theory, however, requires a continuous closed-loop system, which in presence of a discontinuous SMPC feedback is violated [120]. However, by considering a bounded random variable, the authors of [149] show that the closed-loop system is ISS using an ISS-Lyapunov function, while Schlüter and Allgöwer [152] consider an incremental Lyapunov function to prove practical asymptotic stability of the closed-loop system.

Existence and measurability In a recent paper [120], the authors draw attention to the fact that some fundamental properties, such as measurability or existence of solutions, are not currently present in the literature. This question targets the stochastic properties of the system, such as, e.g., the expected value of the cost function (2.10) is well-defined or even if it exists. In particular, the measurability condition of g in equation (2.1) must be fulfilled, such that the optimal state trajectory remains measurable. McAllister and Rawlings [120] provide some theoretical results on Borel measurability of the optimal value function and the optimal control law mapping, while furthermore targeting the stochastic asymptotic stability problem of the closed-loop system, where they introduce a definition of robust asymptotic stability in expectation.

Chance constraints The extension of general nonlinear constraints (2.4) to chance constraints is once again a non-trivial task, even in a static nonlinear optimization problems [1]. Some solution approaches are based on sparse-grid integration [64] or via mixed-integer nonlinear programs [1]. Thus, most of the work on SMPC for nonlinear systems uses convex JCCs or ICCs, which renders the resulting optimization problem computationally tractable.

Due to the aforementioned technical issues, most of the work on SMPC targets linear systems with linear constraints, for which rigorous theoretical properties can be proven. Some of the recent advances in SMPC for linear systems are presented below.

2.3 Stochastic MPC for linear systems

The goal of this section is to give a brief overview of the state of the art in SMPC, while describing in detail a particular approach used in different chapters of this thesis.

In the following, we consider linear time-invariant systems of the form

$$x(k+1) = Ax(k) + Bu(k) + w(k) \tag{2.6}$$

with system matrix $A \in \mathbb{R}^{n\times n}$, input matrix $B \in \mathbb{R}^{n\times m}$ and a zero-mean i.i.d. noise $w(k)$ with known covariance matrix $\Sigma_{\mathrm{w}} \succ 0$. The system is subject to a JCC for the state

$$\mathbb{P}(x(k) \in \bar{\mathbb{X}}) \geq p \quad \forall k \in \mathbb{N}, \tag{2.7}$$

where $\bar{\mathbb{X}} \subseteq \mathbb{R}^n$ is a convex set containing the origin. Input chance constraints can be imposed in a similar way, but for brevity we omit this in the introduction. The control objective is to minimize the expected average infinite horizon cost

$$J_\infty = \lim_{t\to\infty} \mathbb{E}\left(\frac{1}{t}\sum_{k=0}^{t-1} l(x(k), u(k))\right), \tag{2.8}$$

where $l : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_{\geq 0}$ denotes a non-negative stage cost function, subject to the state constraints (2.7) at all times $k \in \mathbb{N}$.

Receding horizon optimization In the following, we approximate the infinite horizon cost (2.8) in a receding horizon fashion, resulting in the concept of SMPC. The basic idea of a SMPC controller is to solve a finite-horizon stochastic optimal control problem (FH-SOCP) over a prediction horizon $N \in \mathbb{N}$, implementing only the first control input and repeating the steps with a shifted time window. To make predictions with MPC, we define the predictive dynamics (open-loop dynamics), which has the same structure as its closed-loop surrogate (2.6), i.e.,

$$x(t+1|k) = Ax(t|k) + Bu(t|k) + w(t|k), \tag{2.9}$$

where $x(t|k)$ and $u(t|k)$ denote t-step ahead predictions made at time k. The predicted disturbance $w(t|k)$ is equal in distribution to $w(t+k)$, while the statistics of $w(k)$ for all $k \in \mathbb{N}$ are known by assumption. The finite horizon cost function is defined as follows

$$J(x(\cdot|k), u(\cdot|k)) = \mathbb{E}\left(V_{\mathrm{f}}(x(N|k)) + \sum_{t=0}^{N-1} l(x(t|k), u(t|k)) \middle| x(0|k)\right), \tag{2.10}$$

where $V_{\mathrm{f}} : \mathbb{X} \to \mathbb{R}_{\geq 0}$ denotes a terminal cost function that mimics the infinite horizon tail for $t > N$. The FH-SOCP is defined as

$$\begin{aligned}
\min_{x(\cdot|k),u(\cdot|k)} \quad & J(x(\cdot|k), u(\cdot|k)) && && \text{(2.11a)}\\
\text{s.t.} \quad & x(t+1|k) = Ax(t|k) + Bu(t|k) + w(t|k) && \forall t \in \{0,\dots,N-1\} && \text{(2.11b)}\\
& \mathbb{P}(x(t|k) \in \bar{\mathbb{X}} \mid x(0|k)) \geq p && \forall t \in \{0,\dots,N-1\} && \text{(2.11c)}\\
& x(N|k) \in \mathbb{X}_{\mathrm{f}}, && && \text{(2.11d)}
\end{aligned}$$

where $\mathbb{X}_\mathrm{f} \subseteq \bar{\mathbb{X}}$ denotes a terminal set that is used to ensure stability [57]. In the following, we outline two approaches to address the FH-SOCP, namely an analytical approximation method and a scenario-based approach.

2.3.1 Analytical approximation method

The key elements of an analytical approximation method include a proper control parameterization, a cost function that allows for an analytical evaluation of the expectation operator and a deterministic reformulation of the chance constraints.

Control parameterization In the FH-SOCP (2.11), optimizing over the control input $u(\cdot|k)$ would result in an infinite dimensional stochastic programming problem, since the input sequence $u(\cdot|k)$ depends on $w(\cdot|k)$. These types of problems are difficult to solve, which renders an online implementation of the optimization problem computationally intractable.

Therefore, we need to restrict the set of control policies so that the MPC makes only nominal (uncertainty-free) predictions. The literature commonly proposes affine parameterization, such as affine disturbance feedback [96, 135, 171] or affine state/error-feedback [29, 77, 98]. Note that these two parameterizations are actually equivalent [65], with the difference that affine state/error-feedback typically uses fixed feedback gains, while in disturbance feedback the feedback gains are optimized online. However, dynamic feedback gains can also be used in the case of state/error-feedback, but this results in the set of admissible control parameters (dynamic feedback gains and nominal inputs) being non-convex [65, Prop. 1], while in the case of affine disturbance feedback the set is convex [65, Prop. 2]. In the following, we consider the class of error-feedback parameterizations of the form

$$u(t|k) \coloneqq v(t|k) + K(x(t|k) - z(t|k)), \tag{2.12}$$

where $K \in \mathbb{R}^{m\times n}$ is a fixed stabilizing feedback gain for the matrix pair (A, B), z the nominal state and v the nominal input. Substituting (2.12) into (2.9) and using linear superposition, we can separate the state x into a deterministic (nominal) part z and a stochastic error part e, such that $x = z + e$. The dynamics are governed by

$$z(t+1|k) = Az(t|k) + Bv(t|k) \tag{2.13}$$
$$e(t+1|k) = A_\mathrm{K}e(t|k) + w(t|k), \tag{2.14}$$

where $A_\mathrm{K} = A + BK$ denotes the closed-loop matrix and $e(0|k) = x(0|k) - z(0|k)$.

Chance constraints In the literature on SMPC, various chance constraint reformulation methods are proposed. In general, we can distinguish between two subcategories, namely scenario-based and analytical approximation methods. In this section, we present the latter approach, where we reformulate the chance constraints (2.11c) as deterministic constraints and treat the uncertainty similarly to robust MPC by tightening the constraints, cf. [57,

60, 77, 106]. Within these analytical methods, there are minor variations in the derivation of the constraint tightening, which are omitted here for brevity.

We focus on one particular constraint tightening technique based on probabilistic reachable sets [77], where we reformulate the chance constraints (2.11c) by using the superposition $x = z + e$, such that

$$\mathbb{P}(x(t|k) \in \bar{\mathbb{X}} \mid x(0|k)) = \mathbb{P}(z(t|k) + e(t|k) \in \bar{\mathbb{X}} \mid x(0|k)) \geq p. \tag{2.15}$$

This is equal to $\exists \mathcal{R} \subseteq \mathbb{R}^n : z(t|k) \in \mathbb{Z} \coloneqq \bar{\mathbb{X}} \ominus \mathcal{R}$ and $\mathbb{P}(e(t|k) \in \mathcal{R} \mid x(0|k)) \geq p$, where the set $\mathcal{R}$ is also known as a PRS.

Definition 2.6 ([77] PRS)**.** *A set $\mathcal{R}$ is said to be a probabilistic reachable set (PRS) of probability level p for system* (2.14) *with $e(0|k) = 0$ if $\mathbb{P}(e(t|k) \in \mathcal{R} \mid e(0|k)) \geq p \quad \forall t \in \mathbb{N}$.*

For completeness, we also introduce a k-step PRS, which will be used later in this thesis.

Definition 2.7 ([77] t-step PRS)**.** *A set $\mathcal{R}(t)$ with $t \in \mathbb{N}$ is said to be a t-step PRS of probability level p for system* (2.14) *with $e(0|k) = 0$ if $\mathbb{P}(e(t|k) \in \mathcal{R}(t) \mid e(0|k)) \geq p$.*

Remark 2.2. *Note that the chance constraint* (2.15) *is conditioned on the initial value $x(0|k)$, while the PRS is conditioned on $e(0|k) = 0$. Therefore, we need to initialize $x(0|k)$ with $x(k)$ such that the error $e(0|k) = 0$, which can be enforced in terms of* (2.13) - (2.14) *by setting $z(0|k) = x(k)$, resulting in $e(0|k) = x(k) - z(0|k) = 0$. This immediately verifies the chance constraints in closed-loop, since the distribution of $x(k+1)$ given $x(k)$ is equal to the distribution of $x(1|k)$, i.e.,*

$$\begin{aligned}\mathbb{P}(x(1|k) \in \bar{\mathbb{X}} \mid x(k)) &= \mathbb{P}(z(1|k) + e(1|k) \in \bar{\mathbb{X}} \mid x(k)) \\ &\overset{(2.14),e(0|k)=0}{=} \mathbb{P}(z(1|k) + \underbrace{w(0|k)}_{=w(k)} \in \bar{\mathbb{X}} \mid x(k)) = \mathbb{P}(x(k+1) \in \bar{\mathbb{X}} \mid x(k)) \geq p.\end{aligned} \tag{2.16}$$

However, initializing $z(0|k) = x(k)$ in the presence of unbounded additive uncertainty may lead to a loss of feasibility of the MPC optimization problem, which advocates the use of a backup initialization strategy as proposed by [57, 60, 77, 109]. Typically, the shifted optimal solution $z(0|k) = z(1|k-1)$ from time $k-1$ is used as a guaranteed feasible backup solution, i.e., $x(0|k) = \mathbb{E}(x(k) \mid x(k-1))$. The problem we now face is that by using an initial value different from $x(k)$, the error $e(0|k) \neq 0$ leads to the loss of closed-loop chance constraint guarantees (without further distributional specifications), since the PRS $\mathcal{R}$ is only valid w.r.t. $e(0|k) = 0$, cf. Definition 2.6. In other words, the second equality of (2.16) *does not hold, and therefore we can conclude that the chance constraint* (2.15) *is only verified in prediction*

$$\mathbb{P}(x(1|k) \in \bar{\mathbb{X}} \mid x(k-1)) = \mathbb{P}(z(1|k) + e(1|k) \in \mathbb{Z} \oplus \mathbb{R} \mid x(k-1)) \geq p, \tag{2.17}$$

while the distribution of $x(k+1)$ given $x(k)$ is not equal to the distribution of $x(1|k)$. The above initialization scheme is termed direct feedback in the related literature [78], since it directly initializes the nominal prediction dynamics (2.13).

In contrast, consider a bounded random variable $w \in \mathbb{W}$, where one can use a robust MPC inspired constraint tightening [37] to robustly ensure that $z(0|k) = x(k)$ is always feasible for any realization of $w \in \mathbb{W}$, cf. [43, 106, 107]. In this case, the error $e(0|k)$ is always zero, which implies that the chance constraints are satisfied for the closed-loop system for all k.

In general, a PRS is a non-unique set that can be designed in different ways, e.g., [77, 109] define convex symmetric PRS via mean-variance information (see Chapter 3), while [76] considers non-symmetric mean-variance PRS for Gaussian random variables. Alternatively, [79, 128] propose a design of non-symmetric PRS via scenario-optimization, cf. Section 4.3.2. The latter approach is extended to the distributionally robust case in [111] and will be covered in Section 6.3.1.

In the following, we focus on a PRS design with mean-variance information of the error e, which is typically used in analytical frameworks [77, 109]. In particular, for a zero-mean random variable w and $e(0|k) = 0$, it holds that also the predicted error is zero-mean, i.e.,

$$\mathbb{E}(e(t+1|k)) = A_{\mathrm{K}}\mathbb{E}(e(t|k)) + \mathbb{E}(w(t|k)) = 0 \quad \forall t \in \mathbb{N},$$

while the predicted variance evolves as

$$\mathrm{var}(e(t+1|k)) = A_{\mathrm{K}}\mathrm{var}(e(t|k))A_{\mathrm{K}}^{\top} + \Sigma_{\mathrm{w}} \quad \forall t \in \mathbb{N}, \tag{2.18}$$

which, due to Schur stability of A_{K} and positive definiteness of Σ_{w} is guaranteed to converge to a stationary covariance matrix $\Sigma_{\mathrm{e}} = \lim_{t\to\infty} \mathrm{var}(e(t|k))$. Since the distribution of $w(k)$ is equal for all $k \in \mathbb{N}$, also the variance evolution for $t \geq 0$ is equivalent for all times k. Thus, this can be computed offline. The stationary covariance matrix allows us to define a mean-variance PRS, e.g., via Chebyshev's inequality

$$\mathcal{R} = \{e \in \mathbb{R}^n \mid e^{\top}(\Sigma_{\mathrm{e}})^{-1}e \leq \gamma\},$$

where $\gamma = n/(1-p)$, which holds for arbitrary probability distributions. However, if the disturbance is normally distributed, then $\gamma = \mathcal{X}_n^2(p)$ yields the tightest probability bound, where $\mathcal{X}_n^2(p)$ is the inverse cumulative distribution function of the Chi-squared distribution at probability level p with n degrees of freedom.

Under an additional central convex unimodality assumption for the distribution of w, it is possible to ensure closed-loop constraint satisfaction in case of $x(0|k) = \mathbb{E}(x(k) \mid x(k-1))$, i.e., with the backup strategy discussed in Remark 2.2 and in particular when $e(0|k) \neq 0$.

Definition 2.8 ([51, Def. 3.1])**.** *A distribution $\mathcal{Q}$ supported in $\mathbb{R}^n$ is called central convex unimodal (CCU) if it is in the closed convex hull of the set of all uniform distributions on symmetric compact convex bodies in $\mathbb{R}^n$.*

Therefore, we make use of a property of CCU distributions called monotone unimodality [51], which, assuming that $z(0|k) = z(1|k-1)$, allows to assert that

$$\mathbb{P}(e(t|k) \in \mathcal{R}) \leq \mathbb{P}(e(0|t+k) \in \mathcal{R}), \tag{2.19}$$

cf. [77]. This yields that (2.17) can be upper bounded as follows

$$\begin{aligned} p &\leq \mathbb{P}(x(1|k) \in \bar{\mathbb{X}} \mid x(k-1)) = \mathbb{P}(z(1|k) + e(1|k) \in \mathbb{Z} \oplus \mathcal{R} \mid x(k-1)) \\ &\overset{(2.19)}{\leq} \mathbb{P}(z(1|k) + e(0|k+1) \in \mathbb{Z} \oplus \mathcal{R} \mid x(k-1)) = \mathbb{P}(x(k+1) \in \bar{\mathbb{X}} \mid x(k-1)), \end{aligned} \tag{2.20}$$

which again verifies the chance constraints in closed-loop. A similar conclusion can be made by constraining the initial state $z(0|k)$ on a line between $x(k)$ and $z(1|k-1)$, as proposed by [93, 153].

Remark 2.3. *In direct feedback schemes with a feasibility-based initialization constraint (that is, use always $z(0|k) = x(k)$ if possible and $z(0|k) = z(1|k-1)$ else), the closed-loop chance constraints are typically verified conservatively. In particular, the chance constraint* (2.7) *is unconditionally satisfied, i.e., conditioned on the initial state $x(0)$, which therefore allows for multiple subsequent constraint violations, while the chance constraint* (2.20) *is enforced conditioned on $x(k)$, making subsequent constraint violations highly improbable [78]. Hence, constraint* (2.20) *effectively enforces the chance constraint with regard to an immediate effect of a disturbance realization only, cf. [78] for a detailed discussion.*

Remark 2.4. *An alternative to direct feedback (Remark 2.2) is the so-called indirect feedback initialization proposed by [76]. The idea is to initialize the nominal dynamics* (2.13) *always with $z(0|k) = z(1|k-1)$ for $k \geq 1$. In view of this, the closed-loop error $e(k)$ evolves autonomously from the prediction dynamics, which allows us to formulate a PRS directly for the closed-loop error $e(k)$ instead of the prediction error* (2.14)*. This immediately ensures non-conservative closed-loop chance constraint satisfaction by conditioning the chance constraints on $e(0)$ instead of $x(0|k)$, i.e.,*

$$\mathbb{P}(x(1|k) \in \bar{\mathbb{X}} \mid e(0)) = \mathbb{P}(z(1|k) + e(k+1) \in \bar{\mathbb{X}} \mid e(0)) = \mathbb{P}(x(k+1) \in \bar{\mathbb{X}} \mid e(0)) \geq p,$$

where $x(0|k)$ is always initialized with the measurement $x(k)$.

Cost function In order to analytically reformulate the cost function (2.10), we assume that the terminal and stage cost functions are quadratic, so that $J = J_\mathrm{a}$ with

$$J_\mathrm{a}(x(\cdot|k), u(\cdot|k)) = \mathbb{E}\left(\|x(N|k)\|_P^2 + \sum_{t=0}^{N-1} \|x(t|k)\|_Q^2 + \|u(t|k)\|_R^2 \,\middle|\, x(0|k) \right), \tag{2.21}$$

where Q, R, P are symmetric positive definite weighting matrices and P additionally satisfies the Lyapunov inequality

$$A_\mathrm{K}^\top P A_\mathrm{K} + Q + K^\top R K - P \preceq 0.$$

Using standard arguments from linear quadratic stochastic control [57], and in particular the i.i.d. assumption of $w(k)$, the expected value in (2.21) can be analytically evaluated and the cost can be separated into a mean part $J_{\mathrm{a}}^{\mathrm{m}}$ and a variance part $J_{\mathrm{a}}^{\mathrm{v}}$, such that $J_{\mathrm{a}} = J_{\mathrm{a}}^{\mathrm{m}} + J_{\mathrm{a}}^{\mathrm{v}}$ with

$$J_{\mathrm{a}}^{\mathrm{m}}(z(\cdot|k), v(\cdot|k)) = \|z(N|k)\|_P^2 + \sum_{t=0}^{N-1} \|z(t|k)\|_Q^2 + \|v(t|k)\|_R^2$$

$$J_{\mathrm{a}}^{\mathrm{v}}(\mathrm{var}(e(\cdot|k))) = \mathrm{tr}(\mathrm{var}(e(N|k) \mid x(0|k))P) + \sum_{t=0}^{N-1} \mathrm{tr}(\mathrm{var}(e(t|k) \mid x(0|k)(Q + K^\top RK))).$$

Note that in case of an affine state/error-feedback parameterization, the feedback gain K is typically fixed, cf. [77, 93, 109, 153]. In this case, $J_{\mathrm{a}}^{\mathrm{v}}(\cdot)$ cannot be improved and can be neglected in a receding horizon implementation. However, other publications use dynamic feedback gains or require a variance prediction in order to ensure recursive feasibility, cf. [57, 59, 60], in which case the variance cost $J_a^v(\cdot)$ must be taken into account.

Remark 2.5. *The expected value in* (2.10) *is conditioned on* $x(0|k)$*, which renders this operator ambiguous. In particular, it follows from the discussion in Remark 2.2 that* $z(0|k)$ *must be either* $x(k)$ *or* $\mathbb{E}(x(k)\,|\,x(k-1))$ *to ensure recursive feasibility. Thus, the expectation operator is either conditioned on* $x(k)$*, if* $z(0|k) = x(k)$ *is feasible, or on* $x(k-1)$*, if the backup solution* $z(0|k) = z(1|k-1)$ *is adopted. This can similarly be observed in the variance cost* $J_{\mathrm{a}}^{\mathrm{v}}$*, where the variance operator is conditioned on* $x(0|k)$*, i.e., if* $x(0|k) = x(k)$*, the variance is* $var(e(0|k) \mid x(0|k)) = 0$*, while if* $x(0|k) = \mathbb{E}(x(k) \mid x(k-1))$*, the variance is* $var(e(0|k) \mid x(0|k)) \succ 0$*.*

This impacts the proof of convergence, which in general only exists for analytical frameworks with quadratic cost, e.g., the authors of [60] use the predicted nominal state/input and error variance to ensure a Lyapunov decrease condition, while [77] use only the nominal cost without the variance part, but require an additional Lipschitz argument of the value function to prove convergence. Recently, the authors of [93] have shown that by constraining $z(0|k)$ *on a line between* $x(k)$ *and* $z(1|k-1)$*, a more elegant proof of convergence exists, while no additional assumptions are required.*

Optimization problem Finally, we present a PRS-based SMPC optimization problem that we solve for all times $k \in \mathbb{N}$.

$$\begin{aligned}
\min_{z(\cdot|k), v(\cdot|k)} \quad & J_{\mathrm{a}}(z(\cdot|k), v(\cdot|k)) && && \text{(2.22a)}\\
\text{s.t.} \quad & z(t+1|k) = Az(t|k) + Bv(t|k) && \forall t \in \{0, \dots, N-1\} && \text{(2.22b)}\\
& z(t|k) \in \bar{\mathbb{X}} \ominus \mathcal{R} && \forall t \in \{0, \dots, N-1\} && \text{(2.22c)}\\
& z(N|k) \in \mathbb{Z}_{\mathrm{f}} && && \text{(2.22d)}\\
& z(0|k) \in \{x(k), z(1|k-1)\}, && && \text{(2.22e)}
\end{aligned}$$

where $\mathbb{Z}_\mathrm{f} \subseteq \bar{\mathbb{X}} \ominus \mathcal{R}$ is a terminal set that ensures that for all $z \in \mathbb{Z}_\mathrm{f}$ the constraints $z \in \mathbb{Z}$ and $Az + B\pi_\mathrm{f}(z) \in \mathbb{Z}_\mathrm{f}$ are verified. The function $\pi_\mathrm{f}(z)$ is a so-called terminal controller, which is typically chosen as $\pi_\mathrm{f}(z) = K_\mathrm{f} z$, where $K_\mathrm{f} \in \mathbb{R}^{n \times m}$ is a stabilizing gain for the matrix pair (A, B). Note that the terminal controller gain does not have to coincide with the tube-controller gain K in (2.12). However, for the sake of simplicity we select $K_\mathrm{f} = K$. Furthermore, if the MPC optimization problem is subject to input constraints of the type $v(t|k) \in \mathbb{V} \subseteq \bar{\mathbb{U}} \ominus \mathcal{R}_u$, the terminal set would require the additional condition $\pi_\mathrm{f}(z) \in \mathbb{V}$. The control input for system (2.6) follows from the tube controller (2.12) and is given by

$$u(k) = v^*(0|k) + K(x(k) - z^*(0|k)),$$

where $(z^*(0|k), v^*(0|k))$ are the first elements of the optimizer of (2.22). Then, $z^*(1|k)$ is stored, the remaining sequence is discarded and the optimization is repeated at the next time instant $k + 1$.

Remark 2.6. *Under the assumption that the MPC optimization problem* (2.22) *is feasible with* $z(0|0) = x(0)$*, it can be shown that the optimization problem remains feasible for any realization of the state, i.e., it is recursively feasible. Furthermore, the chance constrains are verified in closed-loop and the closed-loop system achieves an asymptotic average performance bound [77].*

However, a drawback in the present optimization problem is the choice of cost function, which unfortunately deteriorates the closed-loop performance compared to the linear controller $u = Kx$*. Köhler and Zeilinger [93] addressed this issue by using an interpolated initial constraint, which results in a different interpretation of the predicted states and inputs. While* (2.22a) *uses the nominal states and inputs* z *and* v*, the cost function from [93] uses the state and input mean* $\bar{x}$ *and* $\bar{u}$*, i.e.,* (2.9) *and* (2.12) *without the disturbance* w*, which follows from the interpolating initial constraint* $z(0|k) = \lambda_k x(k) + (1 - \lambda_k) z(1|k-1)$ *with* $\lambda_k \in [0, 1]$*. Therefore, no case distinction has to be considered, e.g., as required by [77], which establishes a performance bound no worse than from the linear tube controller* (2.12)*, cf. Chapter 7. The approach is closely related to SMPC with indirect feedback [76], which uses the same cost for the predicted mean.*

2.3.2 Scenario-based SMPC

In the following, we briefly outline a scenario-based SMPC, which, unlike the analytical case, requires only samples of a distribution but no analytical expression of it. Hence, we can treat arbitrary uncertainty sources.

Chance constraints In scenario-based methods, the chance constraints (2.7) are approximated by sampling $N_\mathrm{s} \in \mathbb{N}$ deterministic state predictions based on extracted uncertainty samples of $w(k)$ [27]. Note that this can also be combined with different control parameterizations, such as error-feedback [79]. In this introduction, we focus on a so-called sample-and-discard approach for stochastic programming [28], which replaces the chance constraint

(2.7) with a sample-based surrogate

$$x^{(i)}(t|k) \in \bar{\mathbb{X}} \quad \forall i \in \mathcal{I}_\mathrm{s} \quad \forall t \in \{0, \dots, N-1\}, \tag{2.23}$$

where $\mathcal{I}_\mathrm{s}$ denotes a subset of $\mathcal{I} = \{1, \dots, N_\mathrm{s}\}$ with cardinality $|\mathcal{I}_\mathrm{s}| = N_\mathrm{s} - N_\mathrm{d}$. The constraint satisfaction probability p is regulated by discarding N_d samples of $\mathcal{I}$, such that all discarded samples intentionally violate (2.23). The amount of samples to discard N_d is obtained, e.g., by applying the Chernoff bound to the binomial tail as proposed by [28, Thm. 2.1]. Note that in view of this the chance constraints only hold with a confidence of $1-\beta$ for $\beta \in (0, 1)$, where β depends on the sample size N_s. We will use this concept in Section 4.3.1, where we use scenario-based PRS [79] for indirect output-feedback SMPC.

Cost function We consider the cost function (2.10) and replace the expectation operator with the empirical expectation, which is typically done in the related literature [10, 79, 139, 150]. This implies that the cost function is only approximately minimized, i.e., $J \approx J_\mathrm{s}$ with

$$J_\mathrm{s}(x(\cdot|k), u(\cdot|k)) = \mathbb{E}_{\hat{\mu}^N}\left(V_\mathrm{f}(x(N|k)) + \sum_{t=0}^{N-1} l(x(t|k), u(t|k)) \right). \tag{2.24}$$

To formalize the empirical expectation, we consider the N-fold joint distribution measure $\hat{\mu}^N : \mathcal{B}(\mathbb{W}^N) \to [0, 1]$ of the sequence $w_N = \{w(k), \dots, w(k+N-1)\}$, which is defined as

$$\hat{\mu}^N(F) = \hat{\mu}(F_k) \times \cdots \times \hat{\mu}(F_{k+N-1}) \quad \forall F = (F_k, F_{k+1}, \dots F_{k+N-1}) \in \mathcal{B}(\mathbb{W}^N),$$

where the empirical probability measure $\hat{\mu}(F_k) = N_\mathrm{s}^{-1} \sum_{i=1}^{N_\mathrm{s}} \delta_{w^{(i)}(k)}(F_k)$ concentrates the probability mass N_s^{-1} uniformly on the $\mathcal{I}$ samples $w^{(i)}(k) \sim \mu(w(k))$ via the Dirac delta measure. In other words, the expected cost function is approximated via a sample average approximation (SAA) with N_s predicted state trajectories $x^{(i)}(t|k)$ for $t = 0, \dots, N$, i.e., (2.24) is equal to

$$J_\mathrm{s}(x(\cdot|k), u(\cdot|k)) = \frac{1}{N_\mathrm{s}} \left(\sum_{i=1}^{N_\mathrm{s}} \left[V_\mathrm{f}(x^{(i)}(N|k)) + \sum_{t=0}^{N-1} l(x^{(i)}(t|k), u(t|k)) \right] \right). \tag{2.25}$$

The minimizer $u^*(\cdot|k)$ of the SAA cost function (2.25) is with confidence $1-\beta$ also the minimizer of (2.10), which can be deduced from the empirical nature of the SAA. In fact, the SAA is asymptotically consistent [90], indicating that for $N_\mathrm{s} \to \infty$ the confidence $1-\beta \to 1$ and the minimizer of (2.25) is with probability 1 the minimizer of (2.10). However, the sample size N_s cannot be chosen arbitrarily large, since the computational complexity grows at least linearly, and often exponentially, in N_s [90].

Notably, we do not require any assumptions on the cost functions $V_\mathrm{f} : \mathbb{X} \to \mathbb{R}_{\geq 0}$ and $l : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_{\geq 0}$ except convexity, e.g., any arbitrary norm satisfies this property.

Optimization problem The resulting scenario-based SMPC optimization problem is defined as

$$\begin{aligned}
\min_{x^{(i)}(\cdot|k),u(\cdot|k)} \quad & J_{\mathrm{s}}(x(\cdot|k),u(\cdot|k)) && \text{(2.26a)}\\
\text{s.t.} \quad & x^{(i)}(t+1|k) = Ax^{(i)}(t|k) + Bu(t|k) + w^{(i)}(t|k) && \text{(2.26b)}\\
& x^{(j)}(t|k) \in \bar{\mathbb{X}} \quad \forall j \in \mathcal{I}_{\mathrm{s}} && \text{(2.26c)}\\
& x^{(i)}(0|k) = x_0 && \text{(2.26d)}
\end{aligned}$$

for all $t = 0, \ldots, N-1$ and for all $i \in \mathcal{I}$. It should be noted that recursive feasibility of the above optimization problem is not guaranteed in its present form, and in general is difficult to verify, cf. [123, 139, 150], where this aspect is usually not covered theoretically or only ensured by assumption. Notable exceptions include the work on scenario-based indirect feedback SMPC [79, 128], which ensures closed-loop chance constraints satisfaction with $1-\beta$ confidence while being recursively feasible. For details, we refer to Section 4.3 for a related approach.

Part I

Distributed Stochastic Model Predictive Control

3 Tracking of piece-wise constant references

Motivated by the practical aspect that many control problems require online adjustment of nonzero setpoints [104], such as distributed power systems [166] or water supply systems [102], in this chapter we develop a distributed stochastic MPC framework for tracking of piece-wise constant output reference points for stochastic distributed systems.

In Section 3.2, we introduce the components of the controller that are suitable for distributed optimization, while this section ends with a centralized MPC optmization problem. Section 3.3 introduces a DMPC framework, which, similar to [39], uses distributed consensus ADMM to solve the central MPC optimization problem in a distributed fashion. In addition, we use results from [92] to study the impact of inexact dual optimization on closed-loop chance constraint satisfaction and on recursive feasibility of the MPC optimization problem. In Section 3.4, we present the main result of this chapter, while Section 3.5 is devoted to two numerical examples. This chapter is based on the publication [113] [1].

Related work In [40], the authors propose a cooperative DMPC for tracking of nominal distributed systems based on the concept of distributed invariance. A distributed terminal set for tracking ensures recursive feasibility of the DMPC scheme. In [62] a sequential nominal DMPC scheme for tracking is presented that combines the MPC and steady-state target optimization problems, such that only one problem needs to be solved at each time instant. The DMPC can handle unreachable output reference points while maintaining recursive feasibility of the optimization problem. The authors use a terminal set for tracking to ensure stability. The authors of [58] use a sequential stochastic DMPC for independent systems with coupling constraints. The chance constraints are reformulated with Cantelli's inequality, which yields chance constraint satisfaction in prediction. In this chapter, we generalize our previous work on DSMPC with distributed PRS for set point regulation [109] to the tracking case, while, compared to our proposed tracking DSMPC [114], we consider explicitly inexact minimization.

[1] C. Mark and S. Liu. "Stochastic Distributed Predictive Tracking Control Under Inexact Minimization". In: *IEEE Transactions on Control of Network Systems* 8.4 (2021), pp. 1892–1904 ©2021 IEEE.

3.1 Problem description

We consider linear systems that are decomposable into $i = 1, \dots, M$ non-overlapping subsystems with local dynamics

$$x_i(k+1) = \sum_{j=1}^{M} A_{ij} x_j(k) + B_i u_i(k) + w_i(k) \tag{3.1}$$

$$y_i(k) = \sum_{j=1}^{M} C_{ij} x_{\mathcal{N}_i}(k), \tag{3.2}$$

where $x_i \in \mathbb{R}^{n_i}$ is the local state, $u_i \in \mathbb{R}^{m_i}$ the local input, $y_i \in \mathbb{R}^{p_i}$ the local output and $w_i \in \mathbb{R}^{n_i}$ the local stochastic disturbance. The global state, input, output and disturbance vectors are defined by the stacked column vector of the local quantities, i.e., $x = \mathrm{col}_{i\in\{1,\dots,M\}}(x_i) \in \mathbb{R}^n$, $u = \mathrm{col}_{i\in\{1,\dots,M\}}(u_i) \in \mathbb{R}^m$, $y = \mathrm{col}_{i\in\{1,\dots,M\}}(y_i) \in \mathbb{R}^p$ and $w = \mathrm{col}_{i\in\{1,\dots,M\}}(w_i) \in \mathbb{R}^n$ with global dimensions

$$n = \sum_{i=1}^{M} n_i, \quad m = \sum_{i=1}^{M} m_i, \quad p = \sum_{i=1}^{M} p_i.$$

To simplify the notation, we assume that the distributed system is graph representable.

Definition 3.1. *(Distributed systems on a graph) Let $\mathcal{G} = (\mathcal{M}, \mathcal{E})$ be a graph, where each node $i \in \mathcal{M} = \{1, \dots, M\}$ corresponds to a subsystem with local dynamics* (3.1) - (3.2)*, while the edges $\mathcal{E}$ represent physical couplings between the subsystems. We denote the subset of all nodes that are connected to subsystem i as the strict neighborhood $\bar{\mathcal{N}}_i = \{j \,|\, (i,j) \in \mathcal{E}\}$. By including subsystem i into $\bar{N}_i$, we obtain the neighborhood $\mathcal{N}_i = \bar{N}_i \cup \{i\}$ with state vector $x_{\mathcal{N}_i} = \mathrm{col}_{j\in\mathcal{N}_i}(x_j) \in \mathbb{R}^{|\mathcal{N}_i| n_i}$.*

Using Definition 3.1, we can reformulate the dynamics (3.1) - (3.2) for all $i \in \mathcal{M}$ as follows

$$x_i(k+1) = A_{\mathcal{N}_i} x_{\mathcal{N}_i}(k) + B_i u_i(k) + w_i(k) \tag{3.3a}$$

$$y_i(k) = C_{\mathcal{N}_i} x_{\mathcal{N}_i}(k). \tag{3.3b}$$

We assume that w_i for all $i \in \mathcal{M}$ are zero-mean i.i.d. random variables that follow a central convex unimodal distribution (Definition 2.8) with known covariance matrix $\Sigma_i^{\mathrm{w}} \succ 0$. The overall (global) system is given by

$$x(k+1) = Ax(k) + Bu(k) + w(k) \tag{3.4a}$$

$$y(k) = Cx(k), \tag{3.4b}$$

where w is zero-mean with covariance matrix $\Sigma^{\mathrm{w}} = \mathrm{blkdiag}_{i\in\mathcal{M}}(\Sigma_i^{\mathrm{w}})$. To simplify the exposition, we make the following assumption on stabilizability.

Assumption 3.1. *The pair (A, B) is stabilizable with a structured linear feedback controller*

$$u(k) = Kx(k) = \mathrm{col}_{i\in\mathcal{M}}(K_{\mathcal{N}_i} x_{\mathcal{N}_i}(k)),$$

where $K_{\mathcal{N}_i} \in \mathbb{R}^{m_i \times n_{\mathcal{N}_i}}$, such that the spectral radius $\rho(A_{\mathrm{K}}) < 1$ with $A_{\mathrm{K}} = A + BK$.

Remark 3.1. *The structured controller gains $K_{\mathcal{N}_i}$ for all $i \in \mathcal{M}$ can be computed via structured LMIs and distributed optimization, e.g., with [42, Lemma 10 and Proposition 13].*

We impose polytopic state and input chance constraints for all times $k \in \mathbb{N}$

$$\mathbb{P}\big(x_{\mathcal{N}_i}(k) \in \mathbb{X}_{\mathcal{N}_i} \coloneqq \{x_{\mathcal{N}_i} \mid H_{\mathcal{N}_i} x_{\mathcal{N}_i} \leq h_{\mathcal{N}_i}\} \mid x(0)\big) \geq p_{\mathrm{x}} \tag{3.5a}$$
$$\mathbb{P}\big(u_i(k) \in \mathbb{U}_i \quad \coloneqq \{u_i \mid L_i u_i \leq l_i\} \mid x(0)\big) \qquad \geq p_{\mathrm{u}}, \tag{3.5b}$$

where $h_{\mathcal{N}_i} \in \mathbb{R}^{r_i}_{>0}$, $l_i \in \mathbb{R}^{q_i}_{>0}$ and $p_{\mathrm{x}}, p_{\mathrm{u}} \in (0, 1)$ are the levels of chance constraint satisfaction. In this formulation, we can enforce local input constraints and neighbor-to-neighbor coupled state constraints, i.e., constraints that include states from the neighboring subsystems. The global constraint sets are given by the Cartesian product of the local sets

$$\mathbb{X} \coloneqq \prod_{i\in\mathcal{M}} \mathbb{X}_{\mathcal{N}_i} = \{x \in \mathbb{R}^n \mid Hx \leq h\} \tag{3.6}$$
$$\mathbb{U} \coloneqq \prod_{i\in\mathcal{M}} \mathbb{U}_i \;\; = \{u \in \mathbb{R}^m \mid Lu \leq l\}. \tag{3.7}$$

The goal of this chapter is to design a distributed predictive control algorithm that steers the system output (3.4b) in expectation to a reference point y_{s} in an admissible way, i.e., $\mathbb{E}(y(k)) \to y_{\mathrm{s}}$ as $k \to \infty$. In the following section, we introduce the necessary controller ingredients to pose a centralized MPC optimization problem.

3.2 Controller design

We define for each subsystem $i \in \mathcal{M}$ a distributed error-feedback controller

$$u_i(k) = v_i(0|k) + K_{\mathcal{N}_i}(x_{\mathcal{N}_i}(k) - z_{\mathcal{N}_i}(0|k)), \tag{3.8}$$

where $z_{\mathcal{N}_i}$ and v_i denote the state and input of the i-th nominal subsystem

$$z_i(t+1|k) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + B_i v_i(t|k). \tag{3.9a}$$

The nominal state and input sequences $z_{\mathcal{N}_i}(\cdot|k), v_i(\cdot|k)$ are the result of an MPC optimization problem solved at time k, while the distributed error feedback controller, i.e., the last

term in (3.8) stabilizes the neighborhood error $e_{\mathcal{N}_i} = x_{\mathcal{N}_i} - z_{\mathcal{N}_i}$. Defining the local error as $e_i = x_i - z_i$, it follows by linear superposition that the error dynamics are governed by

$$e_i(t+1|k) = A_{\mathcal{N}_i,\mathrm{K}} e_{\mathcal{N}_i}(t|k) + w_i(t|k), \tag{3.10}$$

where $A_{\mathcal{N}_i,\mathrm{K}} = A_{\mathcal{N}_i} + B_i K_{\mathcal{N}_i}$. In addition, since the disturbance sequence is i.i.d., we have that $w_i(t|k) \overset{d}{=} w_i(t+k)$, cf. Section 2.3. The corresponding global system is given by

$$z(t+1|k) = Az(t|k) + Bv(t|k) \tag{3.11a}$$

$$e(t+1|k) = A_{\mathrm{K}} e(t|k) + w(t|k). \tag{3.11b}$$

In the following, we consider the error dynamics (3.11b) to quantify the uncertainty in a distributed fashion.

3.2.1 Distributed error propagation

In SMPC with PRS-based constraint tightening, we usually use mean-variance information of the error state (3.11b) to design the PRS, cf. Section 2.3.1. Note that under the zero-mean assumption for the disturbance w, also the predicted error (3.11b) is zero mean. Therefore, it is sufficient to consider only the error covariance in the following. With this in mind, we will first introduce global covariance dynamics and point out the necessary changes required for a distributed implementation.

Covariance dynamics We define the global t-step error covariance matrix as $\Sigma^{\mathrm{e}}(t+1|k) = \mathrm{var}(e(t+1|k))$, whose dynamics are governed by

$$\begin{aligned}
&\Sigma^{\mathrm{e}}(t+1|k)\\
=&\mathbb{E}(A_{\mathrm{K}} e(t|k) w^\top(t|k) + w(t|k) w^\top(t|k) + A_{\mathrm{K}} e(t|k) w^\top(t|k) + A_{\mathrm{K}} e(t|k) e^\top(t|k) A_{\mathrm{K}}^\top)\\
=&\mathbb{E}(A_{\mathrm{K}} e(t|k) w^\top(t|k)) + \mathbb{E}(w(t|k) w^\top(t|k)) + \mathbb{E}(A_{\mathrm{K}} e(t|k) w^\top(t|k)) + \mathbb{E}(A_{\mathrm{K}} e(t|k) e^\top(t|k) A_{\mathrm{K}}^\top))\\
=&A_{\mathrm{K}} \Sigma^{\mathrm{e}}(t|k) A_{\mathrm{K}}^\top + \Sigma^{\mathrm{w}},
\end{aligned}$$

where the second equality uses linearity of the expectation and the third equality uses the assumption that $\mathbb{E}(w) = 0$. Furthermore, $\Sigma^{\mathrm{w}} = \mathrm{var}(w(t|k))$ is the covariance matrix of the noise w and $\Sigma^{\mathrm{e}}(t|k) = \mathrm{var}(e(t|k))$ the error covariance. Since A_{K} is Schur stable, $\Sigma^{\mathrm{w}} \succ 0$ and $w(t|k) \overset{d}{=} w(t+k)$ is i.i.d. for all $t, k \in \mathbb{N}$, it is guaranteed that $\Sigma^{\mathrm{e}}(t|k)$ converges to a stationary covariance matrix $\Sigma_{\mathrm{f}}^{\mathrm{e}}$, such that $\Sigma^{\mathrm{e}}(t|k) \preceq \Sigma_{\mathrm{f}}^{\mathrm{e}}$ as $t \to \infty$ for all $k \in \mathbb{N}$. Note that the global covariance matrix Σ^{e} is in general dense, which is due to the couplings between subsystems through the matrix A_{K}. This poses a major problem for the design of distributed controllers. A simple idea to enforce a distributed structure is to upper bound the global covariance matrix $\Sigma_{\mathrm{f}}^{\mathrm{e}}$ by a block-diagonal matrix, cf. [56], where we define

$$\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} = \begin{bmatrix} \hat{\Sigma}_{\mathrm{f},1}^{\mathrm{e}} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \hat{\Sigma}_{\mathrm{f},M}^{\mathrm{e}} \end{bmatrix}, \tag{3.12}$$

such that $\Sigma_{\mathrm{f}}^{\mathrm{e}} \preceq \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}}$. Each block $\hat{\Sigma}_i^{\mathrm{e}}$ corresponds to an upper-bound of the covariance matrix of the i-th subsystems error state (3.10). In the following, we present two design methods for the distributed stationary covariance matrix.

Block-diagonal stationary covariance matrix – A centralized LMI approach

In the first approach, we design (3.12) based on the Lyapunov-like equality

$$\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} = A_{\mathrm{K}} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} A_{\mathrm{K}}^{\top} + \Sigma^{\mathrm{w}}. \tag{3.13}$$

By relaxing (3.13) as an inequality, we can establish the following equivalence for positive definite matrices $\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}}$ via the Schur complement

$$\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} - A_{\mathrm{K}} \hat{\Sigma}_{\mathrm{f}} (\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}})^{-1} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} A_{\mathrm{K}}^{\top} - \Sigma^{\mathrm{w}} \succeq 0 \iff \begin{bmatrix} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} - \Sigma^{\mathrm{w}} & A_{\mathrm{K}} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} \\ \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} A_{\mathrm{K}}^{\top} & \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} \end{bmatrix} \succeq 0.$$

A block diagonal stationary covariance matrix can be found using the following optimization problem

$$\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}*} = \underset{\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} \succeq 0}{\operatorname{argmin}} \quad \|\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}}\|_{\mathrm{F}}^2 \tag{3.14a}$$

$$\text{s.t.} \qquad \begin{bmatrix} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} - \Sigma^{\mathrm{w}} & A_{\mathrm{K}} \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} \\ \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} A_{\mathrm{K}}^{\top} & \hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}} \end{bmatrix} \succeq 0, \tag{3.14b}$$

where the cost function minimizes the Frobenius norm. It then remains to distribute the corresponding neighborhood covariance matrices $\hat{\Sigma}_{\mathcal{N}_i,\mathrm{f}}^{\mathrm{e}} = \mathrm{blkdiag}_{j \in \mathcal{N}_i}(\hat{\Sigma}_{\mathrm{f},j}^{\mathrm{e}*})$ to each subsystem $i \in \mathcal{M}$.

Block-diagonal stationary covariance matrix – An iterative distributed approach

The second approach is based on the idea presented in [56, Sec. II.B]. We define block-diagonal covariance matrices $\hat{\Sigma}_{\mathcal{N}_i}^{\mathrm{e}}$ that consist of local covariance matrices $\hat{\Sigma}_j^{\mathrm{e}}$ for all $j \in \mathcal{N}_i$, while the modified local covariance update equations are given by

$$\hat{\Sigma}_i^{\mathrm{e}}(t+1) = \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}^{\mathrm{e}}(t) \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^{\top} + \Sigma_i^{\mathrm{w}} \quad \forall i \in \mathcal{M}, \tag{3.15}$$

where $\tilde{A}_{\mathcal{N}_i,\mathrm{K}} = \sqrt{|\mathcal{N}_i|} A_{\mathcal{N}_i,\mathrm{K}}$ with $|\mathcal{N}_i|$ being the cardinality of $\mathcal{N}_i$. With [56, Lemma 1] it can be shown that, if each subsystem $i \in \mathcal{M}$ updates its local covariance matrix $\hat{\Sigma}_i^{\mathrm{e}}$ according to (3.15), then the global block-diagonal covariance matrix satisfies

$$\Sigma^{\mathrm{e}}(t+1) \preceq \begin{bmatrix} \hat{\Sigma}_1^{\mathrm{e}}(t+1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \hat{\Sigma}_M^{\mathrm{e}}(t+1) \end{bmatrix} = \hat{\Sigma}^{\mathrm{e}}(t+1).$$

Since this holds for all times $t \in \mathbb{N}$, we obtain the stationary covariance matrix $\hat{\Sigma}_{\mathrm{f}}^{\mathrm{e}}$ as $t \to \infty$, which can be found in a distributed way with Algorithm 1.

Algorithm 1 Distributed stationary covariance matrix synthesis

For each subsystem $i \in \mathcal{M}$ in parallel:
Initialize $\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i}(0) = 0$ and $t = 0$
repeat
  $\hat{\Sigma}^{\mathrm{e}}_i(t+1) = \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i}(t) \tilde{A}^{\top}_{\mathcal{N}_i,\mathrm{K}} + \Sigma^{\mathrm{w}}_i$
  Communicate $\hat{\Sigma}^{\mathrm{e}}_i(t+1)$ to neighbors $j \in \mathcal{N}_i$
  Construct $\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i}(t+1) = \mathrm{blkdiag}_{j \in \mathcal{N}_i}(\hat{\Sigma}^{\mathrm{e}}_j(t+1))$
  Increment time $t \leftarrow t+1$
until convergence, i.e., $\|\hat{\Sigma}^{\mathrm{e}}_i(t+1) - \hat{\Sigma}^{\mathrm{e}}_i(t)\|_{\mathrm{F}} \leq \text{threshold} \quad \forall i \in \mathcal{M}$
Construct stationary covariance matrices $\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}} = \mathrm{blkdiag}_{j \in \mathcal{N}_i}(\hat{\Sigma}^{\mathrm{e}}_j(t)) \quad \forall i \in \mathcal{M}$

Remark 3.2. *The conservatism of Algorithm 1 depends strongly on the cardinality of the neighborhood $\mathcal{N}_i$ as well as on the coupling strength between the subsystems. If the cardinality or coupling strength is too large/strong, this may result in a very large distributed covariance matrix that ultimately leads to a conservative feasible region of the MPC optimization problem or potentially empty constraint sets. As a general rule of thumb, a centralized design should always be preferred and is usually associated with less conservatism. We refer to the following chapter, in particular to Section 4.2.6, where we perform a numerical comparison of the two design methods in the case of output feedback. These results can be trivially carried over to the state-feedback case presented in this chapter.*

3.2.2 Chance constraints via Distributed Probabilistic Reachable Sets

In the following, we address the chance constraints making use of PRS for constraint tightening, cf. Section 2.3.1, where we use the distributed stationary covariance matrix (3.12) to derive so-called distributed PRS, i.e., we derive for each subsystem $i \in \mathcal{M}$ a PRS. We propose two PRS designs based on the stationary neighborhood covariance matrix $\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}$, where the first design is called *marginal PRS* that uses the marginal distribution along each error dimension, while the second design is called *constraint-aligned PRS*, where we consider the marginal distribution along each half-space.

Marginal PRS

A marginal PRS is characterized by the marginal distributions of the zero-mean error $e_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}}$, which is readily obtained by applying Chebyshev's inequality [34, Thm. 1] along each dimension of $e_{\mathcal{N}_i}$, i.e.,

$$\mathbb{P}\left([e_{\mathcal{N}_i}]_j^2 \geq \gamma_{ij} [\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}]_{j,j} \right) \leq \frac{1}{\gamma_{ij}} \quad \forall j \in \{1, \ldots, n_{\mathcal{N}_i}\},$$

while a joint probability statement is obtained by using the union bound

$$\mathbb{P}\Bigg(\bigcup_{j=1}^{n_{\mathcal{N}_i}}\Big[[e_{\mathcal{N}_i}]_j^2 \geq \gamma_{ij}[\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}]_{j,j}\Big]\Bigg) \leq \sum_{j=1}^{n_{\mathcal{N}_i}} \frac{1}{\gamma_{ij}} = 1 - p_{\mathrm{x}}.$$

By calculating the complementary probability, we obtain the condition $1 - \sum_{j=1}^{n_{\mathcal{N}_i}} \frac{1}{\gamma_{ij}} = p_{\mathrm{x}}$, where the risk thresholds γ_{ij} can be used to weight each error dimension individually. For simplicity, we choose $\gamma_{ij} = \gamma_i$, which represents an equally weighted joint chance constrained set, i.e.,

$$1 - \sum_{j=1}^{n_{\mathcal{N}_i}} \frac{1}{\gamma_{ij}} \overset{\gamma_{ij}=\gamma_i}{=} 1 - \frac{n_{\mathcal{N}_i}}{\gamma_i} = p_{\mathrm{x}}.$$

Thus, the marginal PRS is characterized by $\gamma_i = n_{\mathcal{N}_i}/(1 - p_{\mathrm{x}})$ and is defined as

$$\mathcal{R}_{\mathcal{N}_i} = \left\{ e_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}} \;\middle|\; \big|[e_{\mathcal{N}_i}]_j\big| \leq \sqrt{\gamma_i\,[\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}]_{j,j}} \quad \forall j \in \{1,\dots,n_{\mathcal{N}_i}\} \right\}.$$

The PRS for the input error $e_i^{\mathrm{u}} = u_i - v_i = K_{\mathcal{N}_i} e_{\mathcal{N}_i}$ can be defined analogously as

$$\mathcal{R}_i^u = \left\{ e_i^{\mathrm{u}} \in \mathbb{R}^{m_i} \;\middle|\; \big|[e_i^{\mathrm{u}}]_j\big| \leq \sqrt{\gamma_i^{\mathrm{u}}[\hat{\Sigma}^{\mathrm{u}}_{i,\mathrm{f}}]_{j,j}} \quad \forall j \in \{1,\dots,m_i\} \right\},$$

where $\hat{\Sigma}^{\mathrm{u}}_{i,\mathrm{f}} = K_{\mathcal{N}_i}\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}K_{\mathcal{N}_i}^\top$ and $\gamma_i^{\mathrm{u}} = m_i/(1 - p_{\mathrm{u}})$.

Remark 3.3. *The bound γ_i holds for arbitrary probability distributions. However, if the disturbance is normally distributed, then $\gamma_i = \mathcal{X}^2_{n_{\mathcal{N}_i}}(p_{\mathrm{x}})$ yields the tightest probability bound, where $\mathcal{X}^2_{n_{\mathcal{N}_i}}(p_{\mathrm{x}})$ is the quantile function of the Chi-squared distribution at probability level p_{x} with $n_{\mathcal{N}_i}$ degrees of freedom. This similarly holds for γ_i^{u}.*

Remark 3.4. *If the chance constraints are enforced individually, a less conservative marginal parallel-space PRS can be defined with $\gamma_i = 1/(1 - p_{\mathrm{x}})$ for arbitrary distributions or $\gamma_i = \mathcal{X}^2_1(p_{\mathrm{x}})$ for the Gaussian distribution. This should not be confused with individual chance constraints, i.e., Chebyshev's inequality is two-sided and leads to JCCs in each dimension.*

The tightened state and input constraints for the marginal PRS are defined as

$$\tilde{h}_{\mathcal{N}_i} = h_{\mathcal{N}_i} - \big|H_{\mathcal{N}_i}\big|\sqrt{\mathrm{diag}\big(\gamma_i\hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}}\big)} \tag{3.16a}$$

$$\tilde{l}_i = l_i - \big|L_i\big|\sqrt{\mathrm{diag}\big(\gamma_i^{\mathrm{u}}\hat{\Sigma}^{\mathrm{u}}_{i,\mathrm{f}}\big)}, \tag{3.16b}$$

where the absolute value $|\cdot|$ and the square root of the vector of diagonal elements $\mathrm{diag}(\cdot)$ are taken element-wise.

Constraint-aligned PRS

We consider the marginal distribution of the error $e_{\mathcal{N}_i}$ in direction of the j-th half-space $\tilde{e}_j = [H_{\mathcal{N}_i}]_j e_{\mathcal{N}_i} \sim \mathcal{Q}(0, [H_{\mathcal{N}_i}]_j \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}} [H_{\mathcal{N}_i}]_j^\top)$, for which we apply Chebyshev's inequality, i.e.,

$$\mathbb{P}\left(\tilde{e}_j^\top \left([H_{\mathcal{N}_i}]_j \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}} [H_{\mathcal{N}_i}]_j^\top\right)^{-1} \tilde{e}_j \geq \gamma_{ij}\right) \leq \frac{1}{\gamma_{ij}} \quad \forall j = \{1, \dots, r_i\}.$$

Similar to the marginal PRS, we use the union bound and set $\gamma_{ij} = \gamma_i$, resulting in the bound $\gamma_i = r_i/(1 - p_{\mathrm{x}})$, while the constraint-aligned PRS is defined as

$$\mathcal{R}^{\mathrm{A}}_{\mathcal{N}_i} = \left\{ e_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}} \;\middle|\; \left|[H_{\mathcal{N}_i}]_j e_{\mathcal{N}_i}\right| \leq \sqrt{\gamma_i \, [H_{\mathcal{N}_i}]_j \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}} [H_{\mathcal{N}_i}]_j^\top} \quad \forall j \in \{1, \dots, r_i\} \right\}.$$

The tightened state and input constraints via constraint-aligned PRS are given by

$$\tilde{h}^{\mathrm{A}}_{\mathcal{N}_i} = h_{\mathcal{N}_i} - \sqrt{\mathrm{diag}\big(\gamma_i H_{\mathcal{N}_i} \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i,\mathrm{f}} H_{\mathcal{N}_i}^\top\big)} \tag{3.17a}$$

$$\tilde{l}^{\mathrm{A}}_i = l_i - \sqrt{\mathrm{diag}\big(\gamma^{\mathrm{u}}_i L_i \hat{\Sigma}^{\mathrm{u}}_{i,\mathrm{f}} L_i^\top\big)}, \tag{3.17b}$$

where the absolute value $|\cdot|$ and the square root of the vector of diagonal elements $\mathrm{diag}(\cdot)$ are taken element-wise, while γ_i and γ^{u}_i can be chosen according to Remarks 3.3 and 3.4.

Let $\bar{h}_{\mathcal{N}_i}$ and $\bar{l}_i$ be either equal to (3.16) or (3.17), then the chance constraints (3.5) are satisfied in prediction, cf. Remark 2.2, if the nominal MPC optimization problem satisfies the tightened constraints

$$\begin{aligned} \mathbb{Z}_{\mathcal{N}_i} &= \mathbb{X}_{\mathcal{N}_i} \ominus \mathcal{R}_{\mathcal{N}_i} = \{z_{\mathcal{N}_i} | H_{\mathcal{N}_i} z_{\mathcal{N}_i} \leq \bar{h}_{\mathcal{N}_i}\} \\ \mathbb{V}_{\mathcal{N}_i} &= \mathbb{U}_i \ominus \mathcal{R}^{\mathrm{u}}_i = \{v_i | L_{\mathcal{N}_i} v_i \leq \bar{l}_i\}. \end{aligned}$$

The corresponding global constraint sets are given by

$$\mathbb{Z} \coloneqq \prod_{i \in \mathcal{M}} \mathbb{Z}_{\mathcal{N}_i} = \{z \mid Hz \leq \bar{h}\} \tag{3.18a}$$

$$\mathbb{V} \coloneqq \prod_{i \in \mathcal{M}} \mathbb{V}_i = \{v \mid Lv \leq \bar{l}\}. \tag{3.18b}$$

3.2.3 Objective function

In order to define a cost function for tracking, we need to characterize the steady-states of system (3.4a). However, due to the additive noise w, the steady-states can only be formulated in expectation, i.e., w.r.t. the nominal system (3.11a). To this end, we define a

time-varying artificial tracking target $y_\mathrm{s}(k)$ that is consistent with the artificial steady-state pair $(z_\mathrm{s}(k), v_\mathrm{s}(k))$ given by

$$\begin{bmatrix} A-I & B \\ C & 0 \end{bmatrix} \begin{bmatrix} z_\mathrm{s}(k) \\ v_\mathrm{s}(k) \end{bmatrix} = \begin{bmatrix} 0 \\ y_\mathrm{s}(k) \end{bmatrix}. \tag{3.19}$$

The idea we are pursuing is based on Limón et al. [104], where the goal is to track the reachable artificial target $y_\mathrm{s}(k)$ instead of the true reference y^ref at each time instant k. This increases the feasible region of the MPC optimization problem and renders setpoint changes always feasible [104]. To steer the artificial steady-state $y_\mathrm{s}(k)$ to a constant reference y^ref in an admissible way, we define a cost function for tracking in the deviation variables $\Delta z = z - z_\mathrm{s}$ and $\Delta v = v - v_\mathrm{s}$, while we consider a quadratic cost function over a finite prediction horizon $N \in \mathbb{N}$

$$J(\Delta z, \Delta v, y_\mathrm{s}, y^\mathrm{ref}) = \sum_{t=0}^{N-1} \|z(t|k) - z_\mathrm{s}(k)\|_Q^2 + \|v(t|k) - v_\mathrm{s}(k)\|_R^2 + \|y_\mathrm{s}(k) - y^\mathrm{ref}\|_T^2, \tag{3.20}$$

where $Q \succeq 0$, $R \succ 0$ and $T \succ 0$ are block-diagonal weighting matrices for the states, inputs and the output reference. In this chapter, we use a zero-terminal constraint strategy, i.e., $z(N|k) - z_\mathrm{s}(k) = 0$, which renders a terminal cost function obsolete (see Section 3.2.4 and constraint (3.21d) for details). In addition, we omit the variance part that is usually associated with quadratic cost functions in stochastic MPC, since the tube controller gains $K_{\mathcal{N}_i}$ for all $i \in \mathcal{M}$ in (3.8) are fixed, cf. the paragraph *cost function* in Section 2.3.1.

3.2.4 MPC optimization problem

The following MPC optimization problem is solved at every time instant $k \in \mathbb{N}$.

Problem 3.2.1 (Centralized MPC tracking problem)**.**

$$\begin{aligned}
\min_{\mathcal{Z},\mathcal{V},z_\mathrm{s},v_\mathrm{s},y_\mathrm{s}} \quad & J(\Delta z, \Delta v, y_\mathrm{s}, y^\mathrm{ref}) && && \text{(3.21a)} \\
\text{s.t.} \quad & z(t+1|k) = Az(t|k) + Bv(t|k) && \forall t \in \{0, ..., N-1\} && \\
& (z(t|k), v(t|k)) \in \mathbb{Z} \times \mathbb{V} && \forall t \in \{0, ..., N-1\} && \\
& z_\mathrm{s} = Az_\mathrm{s} + Bv_\mathrm{s}, \quad (z_\mathrm{s}, v_\mathrm{s}) \in \theta(\mathbb{Z} \times \mathbb{V}) && && \text{(3.21b)} \\
& y_\mathrm{s} = Cz_\mathrm{s} && && \text{(3.21c)} \\
& z(N|k) = z_\mathrm{s} && && \text{(3.21d)} \\
& z(0|k) \in \{x(k), z(1|k-1)\}, && && \text{(3.21e)}
\end{aligned}$$

where $\mathcal{V} = \{v(0|k), \ldots, v(N-1|k)\}$ *and* $\mathcal{Z} = \{z(0|k), \ldots, z(N|k)\}$ *denote the nominal input and state sequences, respectively.*

The MPC is categorized as a direct feedback controller due to the initial constraint (3.21e), cf. Remark 2.2, where we use the standard conditional update rule, i.e., if the MPC optimization problem 3.2.1 is feasible with $z(0|k) = x(k)$, then we solve it in what is called mode 1. Otherwise, we use the shifted optimal solution from the previous time step $z(0|k) = z(1|k-1)$, which we call mode 2. Similar to [104] the parameter $\theta \in (0, 1)$ renders the resulting steady-state pair admissible and can be set arbitrarily close to 1, while the closed-loop control input for each subsystem (3.3) is given by the tube controller (3.8).

3.3 Distributed consensus optimization with ADMM

In this section, we address a distributed solution of Problem 3.2.1 with consensus ADMM [22]. Let Ξ be the global vector of decision variables

$$\Xi = \left[\begin{array}{c} z(0|k) \\ \vdots \\ z(N|k) \\ \hline v(0|k) \\ \vdots \\ v(N-1|k) \\ \hline z_{\mathrm{s}}(k) \\ v_{\mathrm{s}}(k) \\ y_{\mathrm{s}}(k) \end{array}\right].$$

In order to decompose the MPC optimization problem 3.2.1 into M subproblems, we exploit the structure of the dynamic (equality) constraints (3.9a), steady-state (equality) constraint (3.19) and the inequality constraints (3.6) - (3.7). For each subsystem $i \in \mathcal{M}$, we define a local vector of decision variables ξ_i containing the local predictions, as well as the steady-state triplet in view of subsystem i

$$\xi_i = \left[\begin{array}{c} z^i_{\mathcal{N}_i}(0|k) \\ \vdots \\ z^i_{\mathcal{N}_i}(N|k) \\ \hline v_i(0|k) \\ \vdots \\ v_i(N-1|k) \\ \hline z_{\mathcal{N}_i,\mathrm{s}}(k) \\ v_{i,\mathrm{s}}(k) \\ y_{i,\mathrm{s}}(k) \end{array}\right].$$

Hence, each subsystem $i \in \mathcal{M}$ contains the neighboring states, inputs and steady-state triplet as independent decision variables in $\xi_j, \forall j \in \mathcal{N}_i$ and in Ξ. At this point, the global

optimization problem can be decomposed into M subproblems, each of which optimizes the local vector ξ_i, while a consensus constraint coordinates the local solutions

$$E_i \Xi = \xi_i \quad \forall i \in \mathcal{M}. \tag{3.22}$$

The matrices E_i are so-called mapping operators, where each row is a unit vector with elements in $\{0, 1\}$. The communication graph $\mathcal{G}$ is thus encoded in E_i, while the local vectors ξ_i can be understood as copies of those entries in Ξ that affect subsystem i [39].

Next, the MPC optimization problem 3.2.1 is expressed in the distributed consensus form

$$\begin{aligned}
\min_{\Xi, \xi_i \; \forall i \in \mathcal{M}} \quad & \sum_{i=1}^{M} J_i(\xi_i, y_i^{\text{ref}}) && && \text{(3.23a)}\\
\text{s.t.} \quad & \xi_i = E_i \Xi && \forall i \in \mathcal{M} && \text{(3.23b)}\\
& \xi_i \in \mathcal{S}_i(z_{i,0}) && \forall i \in \mathcal{M} && \text{(3.23c)}\\
& z_{i,0} \in \{x_i(k), z_i(1|k-1)\} && \forall i \in \mathcal{M}, && \text{(3.23d)}
\end{aligned}$$

where the i-th local objective function is given by

$$J_i(\xi_i, y_i^{\text{ref}}) = \|y_{i,\text{s}}(k) - y_i^{\text{ref}}\|_{T_i}^2 + \sum_{t=0}^{N-1} \|z_i^i(t|k) - z_{i,\text{s}}^i\|_{Q_i}^2 + \|v_i(t|k) - v_{i,\text{s}}\|_{R_i}^2$$

and $\mathcal{S}_i(z_{i,0})$ denotes a convex set that enforces all local constraints

$$\mathcal{S}_i(z_{i,0}) \coloneqq \left\{ \xi_i \;\middle|\; \begin{array}{lcl}
z_i(0|k) & = & z_{i,0} \\
z_i^i(t+1|k) & = & A_{\mathcal{N}_i} z_{\mathcal{N}_i}^i(t|k) + B_i v_i(t|k) \\
z_{\mathcal{N}_i}^i(t|k) & \in & \mathbb{Z}_{\mathcal{N}_i} \\
v_i(t|k) & \in & \mathbb{V}_i \\
z_i^i(N|k) & = & z_{i,\text{s}}^i(k) \\
\forall t \in \{0, \dots, N-1\} & & \\
\hline
z_{i,\text{s}}^i(k) & = & A_{\mathcal{N}_i} z_{\mathcal{N}_i,\text{s}}^i(k) + B_i v_{i,\text{s}}(k) \\
y_{i,\text{s}}(k) & = & C_{\mathcal{N}_i} z_{\mathcal{N}_i,\text{s}}^i(k) \\
(z_{\mathcal{N}_i,\text{s}}^i(k), v_{i,\text{s}}(k)) & \in & \theta(\mathbb{Z}_{\mathcal{N}_i} \times \mathbb{V}_i)
\end{array} \right\}. \tag{3.24}$$

To solve the distributed consensus problem (3.23) via ADMM, we define the augmented Lagrangian for the consistency constraint (3.23b), i.e.,

$$\mathcal{L}_i(\Xi, \xi_i, \lambda_i, y_i^{\text{ref}}) = J_i(\xi_i, y_i^{\text{ref}}) + \lambda_i^\top (\xi_i - E_i \Xi) + \frac{\rho}{2} \|\xi_i - E_i \Xi\|_2^2, \tag{3.25}$$

where λ_i is a dual vector and $\rho \in \mathbb{R}_{>0}$ the augmentation factor to increase convexity. The distributed consensus ADMM procedure is summarized in Algorithm 2 and illustrated in Figure 3.1, where we use the notation that ξ_{ji}^+ and λ_{ji} indicate ξ_i^+ and λ_i predicted by subsystem j.

Algorithm 2 Consensus ADMM

1: For each subsystem $i \in \mathcal{M}$ in parallel:
2: Initialize $\lambda_i = 0$, $\Xi = 0$, $z_{i,0} = \{x_i(k) \vee z_i^i(1|k-1)\}$
3: **repeat**
4: $\quad \xi_i^+ = \underset{\xi_i \in \mathcal{S}_i(z_{i,0})}{\operatorname{argmin}} \mathcal{L}_i(\Xi, \xi_i, \lambda_i, y_i^{\mathrm{ref}})$
5: $\quad$ Communicate ξ_i^+ to neighbors $j \in \mathcal{N}_i$
6: $\quad$ Average $\Xi_i^+ = \dfrac{1}{|\mathcal{N}_i|} \displaystyle\sum_{j \in \mathcal{N}_i} E_{ji}^\top (\xi_{ji}^+ + \frac{1}{\rho} \lambda_{ji})$
7: $\quad$ Communicate Ξ_i^+ to neighbors $j \in \mathcal{N}_i$
8: $\quad \lambda_i^+ = \lambda_i + \rho(\xi_i^+ - E_i \Xi^+)$
9: **until** convergence

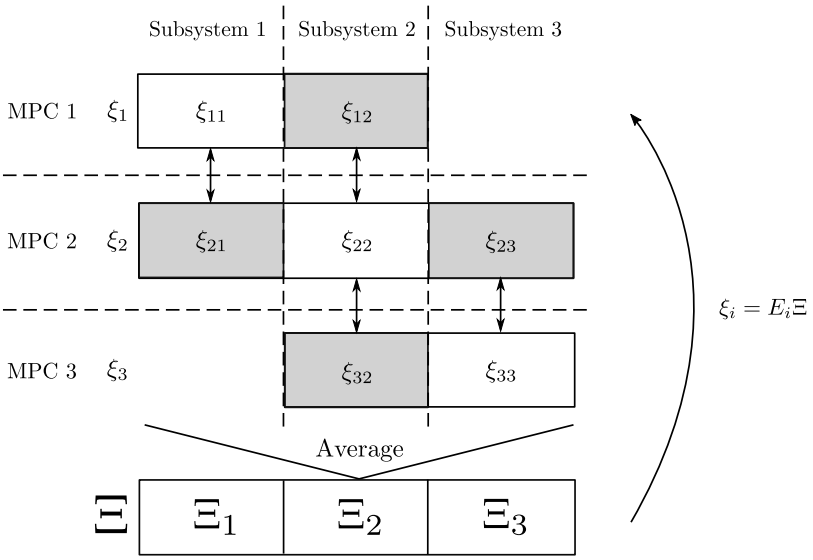

Figure 3.1: Illustration of dual consensus ADMM (Algorithm 2) for three subsystems in a chain topology. The shaded rectangles represent the coupled subsystems. For each iteration, the local solutions ξ_i, i.e., line 4 of Algorithm 2, are averaged in the global vector Ξ (Line 6). Afterwards, Ξ is decomposed into its local equivalents and distributed to each subsystem (Line 7). These steps are repeated until convergence, i.e., until $\|\xi_i - E_i \Xi\|_\infty \leq \epsilon$ for all $i \in \mathcal{M}$ (Line 9).

Remark 3.5. *The residuals $\|\xi_i - E_i \Xi\|_2^2$ converge asymptotically to zero for all $i \in \mathcal{M}$ if all functions $J_i(\xi_i, y_i^{\mathrm{ref}})$ are closed, proper and convex and the unaugmented Lagrangian (equation* (3.25) *without the last term) has a saddle point [22].*

In practice, as computing the exact optimal solution is not feasible at every time step k, Algorithm 2 is stopped when the following condition is met

$$\|\xi_i - E_i \Xi\|_\infty \leq \epsilon \quad \forall i \in \mathcal{M}, \tag{3.26}$$

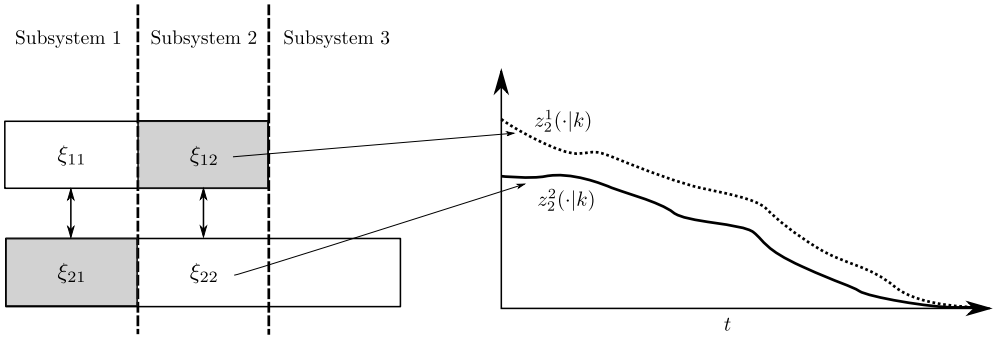

Figure 3.2: Implication of the stopping condition (3.26) for the example in Figure 3.1. The states of subsystem 2 predicted by subsystem 1 (ξ_{12}) are different from the states of subsystem 2 predicted by subsystem 2 (ξ_{22}). This inconsistency occurs in all dynamically coupled subsystems and can lead to a loss of feasibility and closed-loop chance constraint guarantees.

where $\epsilon \in \mathbb{R}_{>0}$. The stopping condition implies that the consistency constraint (3.22) is not exactly satisfied, resulting in predicted state and input trajectories that are only ϵ-feasible. The effects of this can be seen in Figure 3.2.

3.3.1 Inexact distributed optimization – Implications on feasibility

In the subsequent analysis, we quantify the prediction error of an ϵ-feasible solution by introducing uncertain dynamics and steady-state conditions that reflect the inexactness through bounded disturbances δ_i and γ_i, i.e.,

$$z_i^i(t+1|k) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + B_i v_i(t|k) + \delta_i(t) \quad \forall t \in \{0,\dots,N-1\} \tag{3.27a}$$

$$z_{i,\mathrm{s}}^i(k) = A_{\mathcal{N}_i} z_{\mathcal{N}_i,\mathrm{s}}(k) + B_i v_{i,\mathrm{s}}(k) + \gamma_i(k). \tag{3.27b}$$

Definition 3.2. *An ϵ-feasible input sequence resulting from stopping condition* (3.26) *is denoted as $v_{i,\epsilon}(\cdot|k)$, while the ϵ-exact steady input is given by $v_{i,\mathrm{s},\epsilon}(k)$. The corresponding ϵ-feasible state trajectory is denoted as $z_{i,\epsilon}(t|k) = z_i^i(t|k)$ and the ϵ-exact steady-state as $z_{i,\mathrm{s},\epsilon}(k) = z_{i,\mathrm{s}}^i(k)$.*

By inserting the dynamic constraint and steady-state condition from equation (3.24) into $z_i^i(t+1|k)$ and $z_{i,\mathrm{s}}^i(k)$, we obtain

$$\delta_i(t) = A_{\mathcal{N}_i}[z_{\mathcal{N}_i}^i(t|k) - z_{\mathcal{N}_i}(t|k)] \quad \forall t \in \{0,\dots,N\}$$
$$\gamma_i(k) = A_{\mathcal{N}_i}[z_{\mathcal{N}_i,\mathrm{s}}^i(k) - z_{\mathcal{N}_i,\mathrm{s}}(k)],$$

while the stopping condition (3.26) implies that the state trajectories and the steady-state of two neighboring subsystems i and j differ at most by

$$\left\| \begin{bmatrix} z_j^i - z_j^j \\ z_{j,\mathrm{s}}^i - z_{j,\mathrm{s}}^j \end{bmatrix} \right\|_\infty \leq 2\epsilon.$$

Therefore, for each $i \in \mathcal{M}$, we can define local compact sets $\mathbb{D}_i$ and $\mathbb{G}_i$ to represent the inexactness of the predictions

$$\mathbb{D}_i := \left\{ \delta_i \in \mathbb{R}^{n_i} \;\middle|\; \|\delta_i\|_\infty \leq 2\epsilon \Big(\sum_{j \in \mathcal{N}_i \setminus \{i\}} \|A_{ij}\|_\infty \Big) \right\} \tag{3.28}$$

$$\mathbb{G}_i := \left\{ \gamma_i \in \mathbb{R}^{n_i} \;\middle|\; \|\gamma_i\|_\infty \leq 2\epsilon \Big(\sum_{j \in \mathcal{N}_i \setminus \{i\}} \|A_{ij}\|_\infty \Big) \right\}. \tag{3.29}$$

A global form of the local dynamics and steady-state equation (3.27a) - (3.27b) is given by

$$z_\epsilon(t+1|k) = Az_\epsilon(t|k) + Bv_\epsilon(t|k) + \delta(t), \quad \delta \in \mathbb{D} := \mathbb{D}_1 \times \ldots \times \mathbb{D}_M \tag{3.30}$$

$$z_{\mathrm{s},\epsilon}(k) = Az_{\mathrm{s},\epsilon}(k) + Bv_{\mathrm{s},\epsilon}(k) + \gamma(k), \quad \gamma \in \mathbb{G} := \mathbb{G}_1 \times \ldots \times \mathbb{G}_M. \tag{3.31}$$

Feasibility of nominal predictions in inexact DMPC

So far, it has been shown that the inexact predictions can be formulated as equality constraints with bounded additive disturbances. To guarantee constraint satisfaction despite these uncertainties, we adopt the approach used in [94] and utilize robust MPC techniques to tighten the nominal constraint sets (3.18).

It is important to note that the nominal system (3.11a) receives its input from the tube controller (3.8), leading to the definition of the nominal consolidated input trajectory

$$\bar{v}_\epsilon(t|k) = v_\epsilon(t|k) + K[z(t|k) - z_\epsilon(t|k)] \quad \forall t \in \{0, \ldots, N\}, \tag{3.32}$$

which is consistent with the consolidated state and error trajectories.

Definition 3.3. *Let $\bar{z}_\epsilon(\cdot|k), \bar{e}_\epsilon(\cdot|k)$ be the nominal consolidated state and error trajectories that are consistent with the consolidated input trajectory $\bar{v}_\epsilon(\cdot|k)$ and dynamic constraints*

$$\begin{aligned} \bar{z}_\epsilon(t+1|k) &= A\bar{z}_\epsilon(t|k) + B\bar{v}_\epsilon(t|k) \\ \bar{e}_\epsilon(t+1|k) &= A_{\mathrm{K}}\bar{e}_\epsilon(t|k) + \delta(t), \quad \delta(t) \in \mathbb{D} \\ \bar{z}_\epsilon(0|k) &= z(0|k), \end{aligned}$$

where $\bar{e}_\epsilon = z_\epsilon - \bar{z}_\epsilon$. We define the consolidated steady-state as

$$\bar{z}_{\mathrm{s}} = A\bar{z}_{\mathrm{s}} + Bv_{s,\epsilon}, \tag{3.33}$$

such that the true steady-state lies in a bounded set $z_{\mathrm{s}} \in \{\bar{z}_{\mathrm{s}}\} \oplus \mathbb{G}$.

Remark 3.6. *In classical robust tube-based MPC, the additive disturbance δ describes an uncertainty in the system dynamics, while the consolidated trajectories in Definition 3.3 reflect an uncertainty in the prediction [92]. However, the treatment in both cases is equivalent, as we demonstrate in the following.*

The goal is to ensure constraint satisfaction of the consolidated trajectories and admissibility of the steady-state despite inexact minimization. Therefore, the nominal constraints (3.18) must be robustly tightened to counteract the prediction error $\bar{e}_\epsilon$, which can be achieved by using the t-step support function.

Definition 3.4 (t-step support function [41])**.** *The t-step support function is defined as*

$$\begin{aligned} \sigma_{\mathbb{D}}(a,t) = \sup_{\delta \in \mathbb{D}^t} \quad & a\zeta(t) \\ \text{s.t.} \quad & \zeta(0) = 0 \\ & \zeta(l+1) = A_{\mathrm{K}}\zeta(l) + \delta(l) \quad \forall l \in \{0,\dots,t-1\}, \end{aligned} \tag{3.34}$$

where $\mathbb{D}^t$ denotes the Cartesian product $\prod_{i=0}^{t-1} \mathbb{D}$ and $\delta(0),\dots,\delta(t-1)$ a sequence of t disturbance realizations.

By utilizing Definition 3.4, the t-step tightened constraint sets are defined as

$$\bar{\mathbb{Z}}_t = \{z \mid [H]_j z \leq [\bar{h}_t]_j \quad \forall j \in \{1,\dots,r\}\} \tag{3.35a}$$
$$\bar{\mathbb{V}}_t = \{z \mid [L]_j v \leq [\bar{l}_t]_j \quad \forall j \in \{1,\dots,q\}\}, \tag{3.35b}$$

with the right-hand sides given by

$$[\bar{h}_t]_j = [\bar{h}]_j - \sigma_{\mathbb{D}}([H]_j, t) \tag{3.36a}$$
$$[\bar{l}_t]_j = [\bar{l}]_j - \sigma_{\mathbb{D}}([L]_j K, t). \tag{3.36b}$$

Problem 3.3.1 (Tightened MPC tracking problem with inexact optimization)**.**

$$\min_{\mathcal{Z},\mathcal{V},z_{\mathrm{s}},v_{\mathrm{s}},y_{\mathrm{s}}} \quad J(\Delta z, \Delta v, y_{\mathrm{s}}, y^{\mathrm{ref}}) \tag{3.37a}$$
$$\text{s.t.} \quad z(t+1|k) = Az(t|k) + Bv(t|k) \qquad \forall t \in \{0,...,N-1\} \tag{3.37b}$$
$$(z(t|k), v(t|k)) \in \bar{\mathbb{Z}}_t \times \bar{\mathbb{V}}_t \qquad \forall t \in \{0,...,N-1\} \tag{3.37c}$$
$$z_{\mathrm{s}} = Az_{\mathrm{s}} + Bv_{\mathrm{s}}, \quad (z_{\mathrm{s}}, v_{\mathrm{s}}) \in (\bar{\mathbb{Z}}_{\mathrm{s}} \times \bar{\mathbb{V}}_s) \tag{3.37d}$$
$$y_{\mathrm{s}} = Cz_{\mathrm{s}} \tag{3.37e}$$
$$z(N|k) = z_{\mathrm{s}} \tag{3.37f}$$
$$z(0|k) \in \{x(k), z(1|k-1)\}, \tag{3.37g}$$

where $\bar{\mathbb{Z}}_{\mathrm{s}} \coloneqq \bar{\mathbb{Z}}_{N-1} \ominus A_{\mathrm{K}}^N \mathbb{D} \ominus \mathbb{G}$ and $\bar{\mathbb{V}}_s \coloneqq \bar{\mathbb{V}}_{N-1} \ominus K A_{\mathrm{K}}^{N-1} \mathbb{D}$.

The tightened MPC optimization problem 3.3.1 can be solved similarly to the original MPC optimization problem 3.2.1 with ADMM, where local minimization (step 4 in Algorithm 2) uses the tightened constraints $\bar{S}_i$ instead of the original constraints S_i for all $i \in \mathcal{M}$. The sets $\bar{S}_i$ are defined similarly to (3.24) by utilizing the tightened constraints (3.35) as opposed to (3.18).

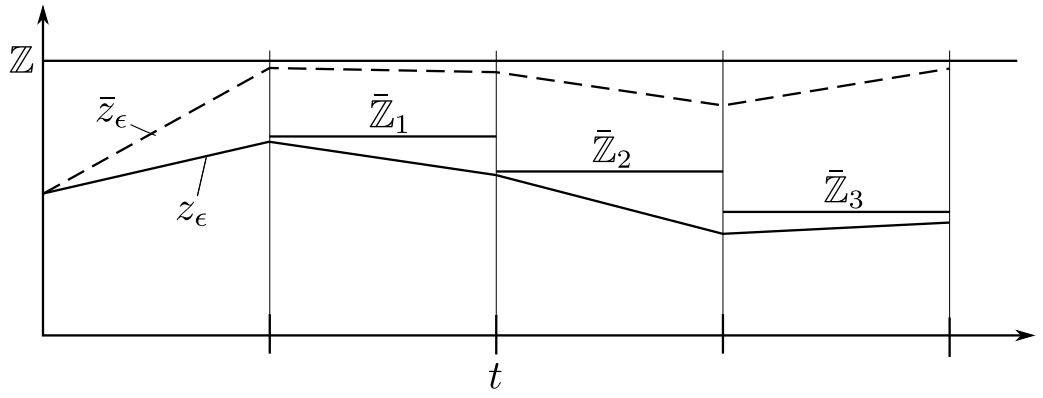

Figure 3.3: The uncertain (inexact) trajectory z_ϵ satisfies the tightened constraints $\bar{\mathbb{Z}}_t$. This implies that the uncertainty-free (consolidated) trajectory $\bar{z}_\epsilon$ verifies the original constraint $\mathbb{Z}$. Note that the uncertainty lies in the prediction, cf. Remark 3.6.

Proposition 3.1. *Let $z_\epsilon(\cdot|k)$ and $v_\epsilon(\cdot|k)$ be an ϵ-feasible solution of the MPC optimization problem 3.3.1 resulting from Algorithm 2, then the consolidated state and input trajectory $\bar{z}_\epsilon(\cdot|k), \bar{v}_\epsilon(\cdot|k)$ satisfy the state and input constraints* (3.35) *for all $i \in \mathcal{M}$.*

Proof. The proof can be found in Section 3.7. □

In Figure 3.3, we illustrate the implication of Proposition 3.1 on feasibility of the consolidated trajectory $\bar{z}_\epsilon$. The resulting nominal closed-loop system is given by

$$\bar{v}_\epsilon(0|k) = v_\epsilon(0|k) + K(z(k) - z_\epsilon(0|k)) \tag{3.38a}$$
$$\bar{z}_\epsilon(1|k) = z(k+1) = Az(k) + B\bar{v}_\epsilon(0|k). \tag{3.38b}$$

Proposition 3.1 can be employed to analyze recursive feasibility of the MPC optimization problem 3.3.1 with respect to the nominal dynamics (3.38), i.e., under application of Mode 2 of the conditional initialization (3.37g).

Lemma 3.1. *Given an approximate solution $z_\epsilon(\cdot|k), v_\epsilon(\cdot|k)$ of Problem 3.3.1 corresponding to the artificial steady-state pair $(z_{\mathrm{s},\epsilon}(k), v_{\mathrm{s},\epsilon}(k))$ at time k, the candidate sequence*

$$\begin{aligned}
\tilde{z}(0|k+1) &= z(k+1) = z_\epsilon(1|k) + \delta \quad \forall \delta \in \mathbb{D} \\
\tilde{z}(t|k+1) &= z_\epsilon(t+1|k) + A_{\mathrm{K}}^t \delta \quad \forall t \in \{0, \dots, N-1\} \\
\tilde{v}(t|k+1) &= v_\epsilon(t+1|k) + K[\tilde{z}(t|k+1) - z_\epsilon(t+1|k)] \quad \forall t \in \{0, \dots, N-2\} \\
\tilde{v}(N-1|k+1) &= v_{\mathrm{s},\epsilon}(k) + K[\tilde{z}(N-1|k+1) - z_{\mathrm{s},\epsilon}(k)] \\
\tilde{z}(N|k+1) &= A\tilde{z}(N-1|k+1) + B\tilde{v}(N-1|k+1)
\end{aligned}$$

is a feasible solution to Problem 3.3.1 at time $k+1$. Problem 3.3.1 is recursively feasible for the nominal closed-loop system (3.38).

Proof. The proof can be found in Section 3.7. □

Algorithm 3 Online algorithm

1: Measure state $x_i(k)$ and communicate it to neighbors $j \in \mathcal{N}_i$
2: Set $z_{i,0} = x_i(k)$ for all $i \in \mathcal{M}$
3: Approximately solve Problem 3.3.1 with Algorithm 2
4: **if** Infeasibility is detected **then** Set $z_{i,0} = z_{i,\epsilon}(1|k-1)$ for all $i \in \mathcal{M}$ and solve Problem 3.3.1 with Alg. 2
5: **end if**
6: Apply control input: $u_i(k) = v_{i,\epsilon}(0|k) + K_{\mathcal{N}_i}(x_{\mathcal{N}_i}(k) - z_{\mathcal{N}_i,\epsilon}(0|k))$

Remark 3.7. *In step 4 of Alg. 3 we need to check for feasibility of the optimization problem. Primal infeasibility can be detected by infeasibility flags of the local solvers in step 4 of Alg 2. Similar to [6], consensus infeasibility can be detected via the residuals* $\alpha_i = \|\xi_i^+ - E_i\Xi^+\|_\infty$, $\beta_i = \|\xi_i^+ - \xi_i\|_\infty$. *However, for the sake of simplicity, we detect consensus infeasibility whenever the algorithm has not converged within a maximum number of iterations.*

Remark 3.8. *In the event of infeasibility, we either resolve the MPC optimization problem 3.3.1 in mode 2 or directly use the shifted optimal solution from the previous time step without solving the optimization problem again.*

3.4 Theoretical analysis

Before stating the main result, we need the following definitions and assumptions.

Definition 3.5. *The set of admissible outputs is defined as* $\mathbb{Y}_\mathrm{A} := \{y_\mathrm{s} = Cz_\mathrm{s} | z_\mathrm{s} \in \bar{\mathbb{Z}}_\mathrm{s}\}$.

To study the optimality of the proposed controller, we introduce the notion of η suboptimal solutions.

Definition 3.6. *Let* J *be the cost function of the MPC optimization problem 3.3.1 and let* $V(z(k), y^\mathrm{ref}(k))$ *be the value function, i.e., the cost function* J *evaluated with the optimal input and state sequences for the output reference* y^ref. *For an* ϵ*-feasible solution* $z_\epsilon(\cdot|k), v_\epsilon(\cdot|k), z_{\mathrm{s},\epsilon}(k)$ *and* $v_{\mathrm{s},\epsilon}(k)$, *we define the suboptimality* η *w.r.t. the optimal solution as*

$$J(\Delta\bar{z}_\epsilon(\cdot|k), \Delta\bar{v}_\epsilon(\cdot|k), y_\mathrm{s}(k), y^\mathrm{ref}) \leq V(\Delta z(k), y^\mathrm{ref}) + \eta,$$

where $\Delta\bar{z}_\epsilon(\cdot|k) = \bar{z}_\epsilon(\cdot|k) - z_{\mathrm{s},\epsilon}(k)$ *and* $\Delta\bar{v}_\epsilon(\cdot|k) = \bar{v}_\epsilon(\cdot|k) - v_{\mathrm{s},\epsilon}(k)$.

Finally, we need an assumption about the boundedness of the set of feasible initial states, which is always fulfilled in case of bounded constraints and a finite prediction horizon.

Assumption 3.2. *The set of feasible initial states $\mathbb{Z}_{\mathrm{ROA}}$ for the MPC optimization problem 3.3.1 is bounded.*

In order to prove convergence of the closed-loop system under the proposed controller, we use the known fact that the value function $V(x(k), y^{\mathrm{ref}})$ of a nominal MPC problem with quadratic cost is piece-wise quadratic in the state x [11, Thm. 4]. This fact, together with Assumption 3.2 implies that there exists a Lipschitz constant L such that

$$V(x, y^{\mathrm{ref}}) \leq V(z, y^{\mathrm{ref}}) + L\|e\|_2. \tag{3.39}$$

Theorem 3.1. *Let Assumption 3.2 hold. Let system* (3.4a) *be controlled with Algorithm 3, then the closed-loop system has the following properties:*

1. *For all $x(0) = z(0) \in \mathbb{Z}_{\mathrm{ROA}}$ and any output reference y^{ref}, the MPC optimization problem 3.3.1 is feasible for $k \in \mathbb{N}$ and the chance constraints* (3.5) *are satisfied in closed-loop.*

2. *The origin $\Delta x = 0$ of the closed-loop system is practically asymptotically stable in expectation. The cost converges to the following asymptotic average*

$$\limsup_{k\to\infty} \frac{1}{k}\mathbb{E}(V(\Delta z(k), y^{\mathrm{ref}}) - V(\Delta z(0), y^{\mathrm{ref}})) \leq \eta + \frac{L}{\sqrt{\lambda_{\min}(P)}} \cdot \mathrm{tr}(\Sigma^{\mathrm{w}} P),$$

   where η is the suboptimality w.r.t. the optimal solution (Definition 3.6), L the Lipschitz constant from (3.39) *and $P \succ 0$ satisfies the Lyapunov inequality*

$$A_{\mathrm{K}}^{\top} P A_{\mathrm{K}} - P \preceq -\kappa I \tag{3.40}$$

   for some $\kappa \in \mathbb{R}_{>0}$.

3. *For any output reference y^{ref}, the output $y(k)$ of system* (3.4b) *converges in expectation to an admissible reference $\tilde{y}_{\mathrm{s}} \in \{\bar{y}_{\mathrm{s}}\} \oplus C[\mathbb{G} \oplus \mathcal{B}_{\eta}(\bar{z}_{\mathrm{s}})]$, where $\mathcal{B}_{\eta}(\bar{z}_{\mathrm{s}})$ is a ball of radius η, $\mathbb{G}$ the disturbance set from* (3.29) *and $\bar{y}_{\mathrm{s}} = C\bar{z}_{\mathrm{s}}$ the consolidated output following from* (3.33) *that minimizes the tracking cost $\|\bar{y}_{\mathrm{s}} - y^{\mathrm{ref}}\|_T^2$.*

   If in addition $y^{\mathrm{ref}} \in \mathbb{Y}_{\mathrm{A}}$, then $\bar{y}_{\mathrm{s}} = y^{\mathrm{ref}}$.

Proof. The proof can be found in Section 3.7. □

3.5 Numerical examples

In this section, we perform two numerical examples to show the basic functionality of our approach. We first consider an academic example in which we highlight the need for tightening constraints due to inexact minimization, while we then implement our controller for a general benchmark example of a four-tank process.

3.5.1 Academic example

Consider a network of $M = 50$ second-order systems in a chain topology with local dynamics

$$x_i(k+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x_i(k) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_i(k) + w_i(k) + \sum_{j \in \mathcal{N}_i \setminus \{i\}} \begin{bmatrix} 0 & 0 \\ 0.1 & 0.1 \end{bmatrix} x_j(k) \quad \forall i \in \mathcal{M}$$
$$y_i(k) = \begin{bmatrix} 1 & 0 \end{bmatrix} x_i(k) \quad \forall i \in \mathcal{M},$$

where $w_i \sim \mathcal{N}(0, 0.001I)$ is a normally distributed noise. For subsystem 1 it holds that $\mathcal{N}_1 = \{1, 2\}$, for subsystem 50 it holds that $\mathcal{N}_{50} = \{49, 50\}$ and for subsystems $i = 2, \ldots, 49$ it holds that $\mathcal{N}_i = \{i-1, i, i+1\}$. For each subsystem $i \in \mathcal{M}$ we enforce two state chance constraints, one for each local dimension $n_i = 2$, i.e.,

$$\mathbb{P}(|[x_i]_p| \leq 2) \geq 0.9 \quad \forall i \in \mathcal{M},\ p \in \{1, 2\}.$$

The system is initialized at $x_i(0) = [0, 0]^\top$ for all $i \in \mathcal{M}$ and we chose the MPC weighting matrices $Q_i = I, R_i = 1, T_i = 70$ and the ADMM augmentation factor $\rho = 5$. We perform $N_m = 100$ Monte-Carlo simulations and keep the reference output for all subsystems $i \in \mathcal{M}$ at 0, except for subsystems 1, 25 and 50.

- For $0 \leq k < 20$, we command three unreachable references $y_1^{\text{ref}} = y_{25}^{\text{ref}} = y_{50}^{\text{ref}} = 2.5$.
- For $20 \leq k < 40$ we command two reachable references $y_1^{\text{ref}} = -1$ and $y_{25}^{\text{ref}} = 0$, and one unreachable reference $y_{50}^{\text{ref}} = 2.5$.
- For $40 \leq k \leq N_s = 59$, we command two unreachable references $y_1^{\text{ref}} = 2.5$ and $y_{50}^{\text{ref}} = 2.5$, and one reachable reference $y_{25}^{\text{ref}} = -1$.

Results In Figure 3.4, the outputs of systems $1, 25$ and 50 are plotted with the corresponding reference values. Whenever an admissible reference is commanded, it can be seen that the output converges in expectation to the reference value. However, if the commanded reference is unreachable, then the output converges to an admissible reference value that minimizes the distance to the commanded one. In Table 3.1, we compare for different convergence thresholds ϵ the average and maximum number of iterations until convergence, denoted Av[it] and Max[it], as well as the absolute suboptimality $\psi = 1 - J/J^*$. The cumulative closed-loop cost is given by $J = \sum_{i=0}^{N_m} \sum_{k=0}^{N_s} \|\Delta x(k)\|_Q^2 + \|\Delta u(k)\|_R^2$, while the optimal exact cumulative cost is denoted as J^*. Furthermore, we compare for subsystem 50 the number of constraint violations ($\#C_{\text{vio}}$) and the maximum empirical constraint violation probability of the second local state, that is, $\hat{\mathbb{P}}_{\max} = \max_{0 \leq k \leq 59} c_{\text{v}}(k)$ with

$$c_{\text{v}}(k) = \frac{1}{N_{\text{m}}} \sum_{k=0}^{N_s} \mathbb{1}_{\{|[x_{50}(k)]_2| > 2\}}, \tag{3.41}$$

where $\mathbb{1}$ denotes the indicator function of the constraint set $\{|[x_{50}(k)]_2| > 2\}$. The results in Table 3.1 can be interpreted as follows. With a reduction of ϵ, the number of iterations

Table 3.1: Convergence for different ADMM parameters. ©2021 IEEE.

ϵ	Av[it]	Max[it]	ψ[%]	$\#C_{\text{vio}}$	$\hat{\mathbb{P}}_{\text{max}}$
5e-2	9.67	13	13.02	3	0.02
5e-3	18.60	24	2.24	147	0.06
5e-4	28.53	38	0.21	225	0.08
5e-5	43.13	76	0.01	253	0.09
0	-	-	0	260	0.09

necessary for convergence increases and simultaneously, the level of suboptimality decreases. Hence, we restore the centralized solution as $\epsilon \to 0$. Furthermore, as the value of ϵ decreases, the size of the disturbance set (3.28) decreases, resulting in a less conservative tightening of the constraints (3.36a) - (3.36b), which in turn leads to an increase in constraint violations. However, the chance constraints of level $p_{\text{x}} = 0.9$, i.e., $1 - \hat{\mathbb{P}}_{\text{max}} > p_{\text{x}}$, are satisfied for any ϵ.

To illustrate the importance of accounting for solver inexactness, we implemented a DSMPC without the additional tightening of constraints (3.35). This leads to an empirical constraint violation of subsystem 50, which exceeds the desired value of $1-p_{\text{x}} = 0.1$, e.g., for $\epsilon = 5{\cdot}10^{-2}$ a closed-loop constraint violation of $\hat{\mathbb{P}} = 0.15$ can be observed, while for $\epsilon = 5{\cdot}10^{-3}$ we obtain $\hat{\mathbb{P}} = 0.11$. As the accuracy increases ($\epsilon \to 0$), the constraint tightening (3.35) becomes less restrictive as the disturbance set (3.28) vanishes in the limit.

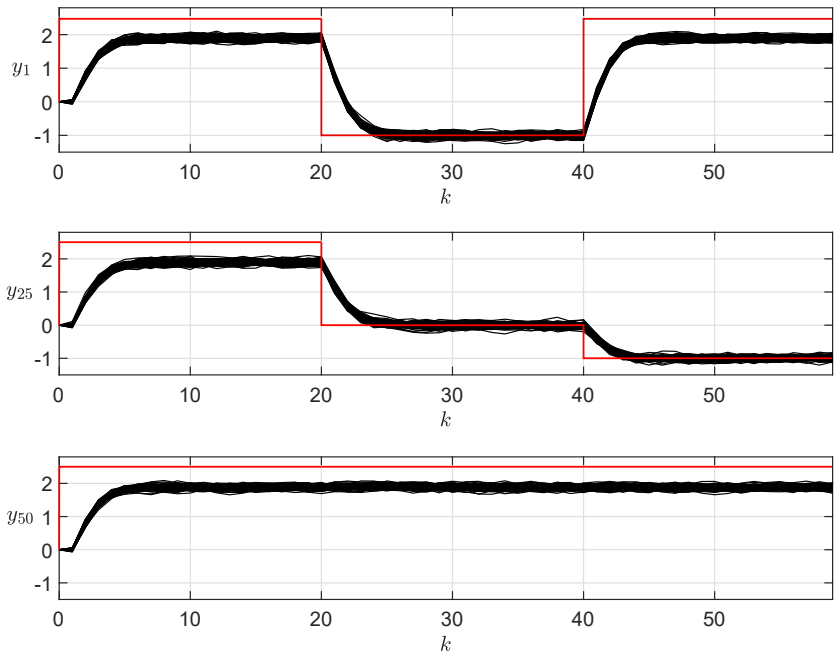

Figure 3.4: Outputs of subsystems 1, 25 and 50 in black and the corresponding commanded references in red. ©2021 IEEE.

3.5.2 Four-tank system

We consider a four-tank system initially proposed by Johansson and Nunes [86]. The goal is to track the levels of tank 1 and 3 while constraining the level for each tank near the operating point, see Figure 3.5. The continuous time dynamic of the system is given by

$$\begin{aligned}
\frac{dh_1}{dt} &= -\frac{a_1}{A_1}\sqrt{2gh_1} + \frac{a_4}{A_4}\sqrt{2gh_4} + \frac{\gamma_1 k_1}{A_1}u_1 \\
\frac{dh_2}{dt} &= -\frac{a_2}{A_2}\sqrt{2gh_2} + \frac{(1-\gamma_1)k_1}{A_2}u_1 \\
\frac{dh_3}{dt} &= -\frac{a_3}{A_3}\sqrt{2gh_3} + \frac{a_2}{A_2}\sqrt{2gh_2} + \frac{\gamma_2 k_2}{A_3}u_2 \\
\frac{dh_4}{dt} &= -\frac{a_4}{A_4}\sqrt{2gh_4} + \frac{(1-\gamma_2)k_2}{A_4}u_2,
\end{aligned}$$

where A_i and a_i are the cross-section of tank and the cross-section of the outlet of tank i, respectively. The constants k_1 and k_2 are conversion parameters that map from voltage applied to the pump to flux of the medium, while γ_1 and γ_2 denote fixed valve positions representing the fraction of water flowing into the lower tanks. The constant g is the gravitational acceleration. The above parameters are chosen according to [15].

To obtain a control-oriented model, we first define the state and input vector as $x = [h_1 \quad h_2 \quad h_3 \quad h_4]^\top$ and $u = [u_1 \quad u_2]^\top$ and linearize the continuous time dynamics around the nominal operating point $z_{\text{op}} = \begin{bmatrix}12.263 & 1.409 & 12.783 & 1.634\end{bmatrix}^\top$ and $v_{\text{op}} = \begin{bmatrix}3 & 3\end{bmatrix}^\top$.

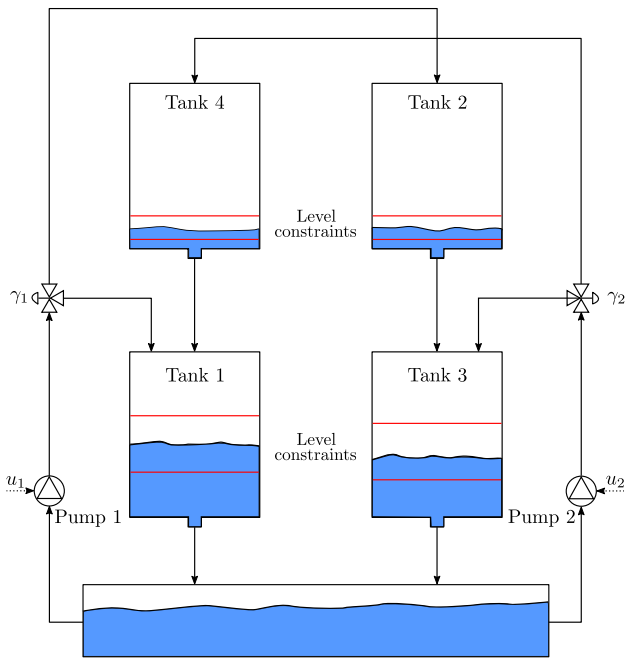

Figure 3.5: Schematic representation of the four-tank system. ©2021 IEEE.

Then, the linearized continuous dynamics is discretized with a sampling time of $T_\mathrm{s} = 5$ seconds using zero-order hold, which yields the global dynamics, input and output matrices for the dynamics (3.4a) - (3.4b)

$$A = \begin{bmatrix} 0.9229 & 0 & 0 & 0.1892 \\ 0 & 0.8469 & 0 & 0 \\ 0 & 0.1488 & 0.9463 & 0 \\ 0 & 0 & 0 & 0.8028 \end{bmatrix}, B = \begin{bmatrix} 0.4024 & 0 \\ 0.1447 & 0 \\ 0 & 0.3037 \\ 0 & 0.2135 \end{bmatrix}.$$

For tracking purposes, we set $C = \mathrm{blkdiag}([1 \quad 0], [1 \quad 0])$. The constraints on the states of the linearized system are JCCs with probability level 0.9, i.e.,

$$\begin{aligned} &\mathbb{P}(|h_2| \leq 0.3) \geq 0.9, \quad \mathbb{P}(|h_4| \leq 0.3) \geq 0.9 \\ &\mathbb{P}(|h_1| \leq 5) \geq 0.9, \quad \mathbb{P}(|h_3| \leq 5) \geq 0.9, \end{aligned}$$

while the inputs have to satisfy the following expectation constraints $|v_1| = |v_2| \leq 3$. By partitioning $x_1 = [h_1 \quad h_2]^\top$ and $x_2 = [h_3 \quad h_4]^\top$, the global system can be separated into distributed subsystems according to (3.3), where the disturbances w_1 and w_2 are assumed to be normally distributed with zero mean and covariance matrix $\Sigma_i = 0.0005I$ for $i = 1, 2$. The MPC weighting matrices are set to $Q_1 = Q_2 = 0.01I$, $R_1 = R_2 = 10$, $T_1 = T_2 = 1000$ and the prediction horizon is $N = 10$, while the augmentation factor of the ADMM algorithm is set to $\rho = 100$. Starting from the initial condition $x(0) = z_\mathrm{op}$, we perform a series of tracking objectives for y_2, while keeping $y_1^\mathrm{ref} = 0$ for all $k \in \mathbb{N}$:

- For $0 \leq k < 50$, we command $y_2^\mathrm{ref} = 0$.
- For $50 \leq k < 150$, we command $y_2^\mathrm{ref} = 0.5$.
- For $150 \leq k \leq 250$, we command $y_2^\mathrm{ref} = -0.5$.

Results In Figure 3.6, we plot the corresponding water levels of the tanks. In Table 3.2, we compare for different accuracy levels ϵ the average and maximum iterations, as well as the suboptimality w.r.t. the central optimal solution. The constraint violation $\hat{\mathbb{P}}_\mathrm{max}$ is computed for tank 4 similar to (3.41). In this scenario, the number of iterations until convergence is higher compared to the academic example in Section 3.5.1. The reason for this is the more complex interconnection between the subsystems.

Table 3.2: ADMM convergence for different parameters. ©2021 IEEE.

ϵ	Av[it]	Max[it]	ψ[%]	$\#C_\mathrm{vio}$	$\hat{\mathbb{P}}_\mathrm{max}$
5e-4	85.56	901	15.2	7	0.032
5e-5	119.74	1325	2.9	15	0.041
5e-6	288.15	2165	0.05	19	0.049
0	-	-	0	25	0.076

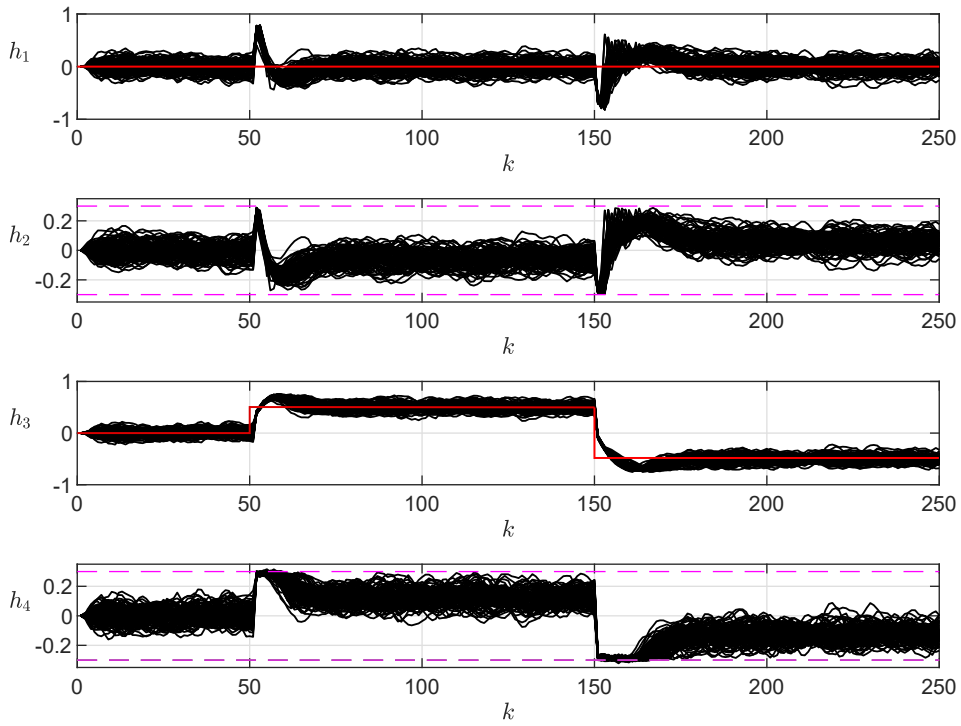

Figure 3.6: Water levels in deviation variables (black) with reference values (red) and level constraints (magenta). ©2021 IEEE.

3.6 Summary

In this chapter, we presented a cooperative DSMPC algorithm for tracking of piece-wise constant output references for distributed linear systems subject to local and coupling chance constraints. We use consensus ADMM to solve the central MPC optimization problem in a distributed fashion and explicitly incorporate the ADMM stopping condition in the MPC design to give strong closed-loop guarantees. The algorithm is proven to be recursively feasible, convergent to an asymptotic average performance bound and is able to converge to arbitrary output reference points in an admissible way. Furthermore, the PRS constraint tightening yields closed-loop chance constraints satisfaction. The numerical examples highlight the effects of different ADMM parameters regarding to the performance, chance constraint satisfaction and iterations until convergence.

3.7 Proofs

Proof of Proposition 3.1

Let $z_\epsilon(\cdot|k), v_\epsilon(\cdot|k)$ be approximate solutions of the MPC optimization problem 3.3.1. The consolidated state trajectory from Definition 3.3 and input trajectory (3.32) satisfy

$$\bar{z}_\epsilon(t|k) \in \{z_\epsilon(t|k)\} \bigoplus_{l=0}^{t-1} A_K^l \mathbb{D}$$
$$\bar{v}_\epsilon(t|k) \in \{v_\epsilon(t|k)\} \bigoplus_{l=0}^{t-1} K A_K^l \mathbb{D},$$

which implies that

$$[H]_j \bar{z}_\epsilon(t|k) \overset{(3.34)}{\leq} [H]_j z_\epsilon(t|k) + \sigma_{\mathbb{D}}([H]_j, t) \leq [\bar{h}_t]_j + \sigma_{\mathbb{D}}([H]_j, t) \overset{(3.36a)}{=} [\bar{h}]_j$$
$$[L]_j \bar{v}_\epsilon(t|k) \overset{(3.34)}{\leq} [L]_j v_\epsilon(t|k) + \sigma_{\mathbb{D}}([L]_j K, t) \leq [\bar{l}_t]_j + \sigma_{\mathbb{D}}([L]_j K, t) \overset{(3.36b)}{=} [\bar{l}]_j.$$

This concludes the proof. □

Proof of Lemma 3.1

The proof is inspired by [92] and consists of three parts.

Candidate sequences We construct the candidate sequence $\tilde{z}(t|k+1)$ by shifting $z_\epsilon(\cdot|k)$ with an additional error term δ propagated through the system dynamics to ensure satisfaction of the initial constraint, while the candidate input $\tilde{v}(\cdot|k+1)$ is constructed by shifting the input sequence $v_\epsilon(\cdot|k)$ with the error-feedback controller to compensate for the additional error term δ. First, we show that the candidate sequence satisfies the perturbed dynamics (3.30), i.e.,

$$\begin{aligned}
\tilde{z}(t|k+1) &= z_\epsilon(t+1|k) + A_{\mathrm{K}}^t \delta \\
&\overset{(3.30)}{=} A z_\epsilon(t|k) + B v_\epsilon(t|k) + \delta(t) + A_{\mathrm{K}}^t \delta \\
&= A\big[\tilde{z}(t-1|k+1) - A_{\mathrm{K}}^{t-1}\delta\big] + B\big[\tilde{v}(t-1|k+1) - K A_{\mathrm{K}}^{t-1}\delta\big] + \delta(t) + A_{\mathrm{K}}^t \delta \\
&= A\tilde{z}(t-1|k+1) + B\tilde{v}(t-1|k+1) + \delta(t),
\end{aligned}$$

where the third equality substitutes the time shifted state and input candidate sequences $z_\epsilon(t|k) = \tilde{z}(t-1|k+1) - A_{\mathrm{K}}^{t-1}\delta$ and $v_\epsilon(t|k) = \tilde{v}(t-1|k+1) - K A_{\mathrm{K}}^{t-1}\delta$.

Constraint satisfaction By the definition of the support function and due to the linear superposition we have that

$$\sigma_{\mathbb{D}}([H]_j, t+1) \geq \sigma_{\mathbb{D}}([H]_j, t) + [H]_j A_{\mathrm{K}}^t \delta \quad \forall \delta \in \mathbb{D}. \tag{3.42}$$

Therefore, the candidate sequence satisfies the following bound

$$\begin{aligned}
[H]_j \tilde{z}(t|k+1) &= [H]_j (z_\epsilon(t+1|k) + A_{\mathrm{K}}^t \delta) \\
&\leq [\bar{h}_{t+1}]_j + [H]_j A_{\mathrm{K}}^t \delta \\
&\overset{(3.36a)}{=} [\bar{h}]_j - \sigma_{\mathbb{D}}([H]_j, t+1) + [H]_j A_{\mathrm{K}}^t \delta \\
&\overset{(3.42)}{\leq} [\bar{h}]_j - \sigma_{\mathbb{D}}([H]_j, t) \leq [\bar{h}]_j \quad \forall t \in \{0, \dots, N-1\},
\end{aligned}$$

which can similarly be shown for the input sequence, so that

$$\begin{aligned}
[L]_j \tilde{v}(t|k+1) &= [L]_j v_\epsilon(t+1|k) + [L]_j K[\tilde{z}(t+1|k) - z_\epsilon(t+1|k)] \\
&\leq [\bar{l}_{t+1}]_j + [L]_j K A_{\mathrm{K}}^t d \\
&\overset{(3.36b)}{=} [\bar{l}]_j - \sigma_{\mathbb{D}}([L]_j K, t+1) + [L]_j K A_{\mathrm{K}}^t \delta \\
&\overset{(3.42)}{\leq} [\bar{l}]_j - \sigma_{\mathbb{D}}([L]_j K, t) \leq [\bar{l}]_j \quad \forall t \in \{0, \dots, N-2\}.
\end{aligned}$$

Terminal equality constraint Now we show that the terminal state of the candidate sequence strictly satisfies the state constraints, i.e.,

$$\begin{aligned}
\tilde{z}(N|k+1) &= A\tilde{z}(N-1|k+1) + B\tilde{v}(N-1|k+1) \\
&= A(z_\epsilon(N|k) + A_{\mathrm{K}}^{N-1}\delta) + B(v_\epsilon(N|k) + K[z_\epsilon(N|k) + A_{\mathrm{K}}^{N-1}\delta - z_{\mathrm{s},\epsilon}(k)]) \\
&\overset{(3.37f)}{=} A z_{\mathrm{s},\epsilon}(k) + B v_{\mathrm{s},\epsilon}(k) + A_{\mathrm{K}}^N \delta \\
&\overset{(3.31)}{=} z_{\mathrm{s},\epsilon}(k) - \gamma(k) + A_{\mathrm{K}}^N \delta \quad \forall \gamma(k) \in \mathbb{G},
\end{aligned}$$

where the third equality is due to the zero terminal constraint (3.37f) and the fourth equality follows from the steady-state condition (3.31). In addition, the terminal state satisfies

$$\tilde{z}(N|k+1) = z_{\mathrm{s},\epsilon}(k) - \gamma(k) + A_{\mathrm{K}}^N \delta \in \bar{\mathbb{Z}}_{\mathrm{s}} \oplus \mathbb{G} \oplus A_{\mathrm{K}}^N \mathbb{D} = [\bar{\mathbb{Z}}_{N-1} \ominus A_{\mathrm{K}}^N \mathbb{D} \ominus \mathbb{G}] \oplus \mathbb{G} \oplus A_{\mathrm{K}}^N \mathbb{D} \subseteq \mathbb{Z},$$

which renders $\tilde{z}(N|k+1)$ admissible. This concludes the proof. □

Proof of Theorem 3.1

The proof consists of 3 parts. First, we establish recursive feasibility and chance constraint satisfaction for the closed-loop system, while in the second part we proof convergence to the artificial steady-state and establish an asymptotic average cost bound. The third part addresses the minimization of the tracking cost function and is briefly outlined, since it uses results from [62].

Part 1

Recursive feasibility: Assume that at time k an ϵ-feasible solution $v_\epsilon(\cdot|k), z_\epsilon(\cdot|k)$ exists. Now, at time step $k+1$ we have to consider the unbounded stochastic disturbance $w(k)$ and condition the initial state of the MPC optimization problem 3.3.1 on its feasibility. For mode 2 we set $z(0|k+1) = z(1|k) = z_\epsilon(1|k) + \delta$, which, by Lemma 3.1 is guaranteed to be feasible. Thus, if $x(0) = z(0) \in \mathbb{Z}_{\mathrm{ROA}}$, then the MPC optimization problem 3.3.1 is recursively feasible for all $k \in \mathbb{N}$ under the conditional update.

Closed-loop chance constraints: By definition of the consolidated state and input trajectories and Proposition 3.1, we have that the true states and inputs of the resulting nominal closed-loop system (3.38) satisfy $z(\cdot|k) \in \mathbb{Z}$, $v(\cdot|k) \in \mathbb{V}$ for all $k \geq 0$. Furthermore, we have by Definition 2.6 that $\mathbb{P}(e(t|k) \in \mathcal{R} \mid e(0|k)) \geq p_{\mathrm{x}}$ and $\mathbb{P}(e^{\mathrm{u}}(t|k) \in \mathcal{R}^{\mathrm{u}} \mid e(0|k)) \geq p_{\mathrm{u}}$. Thus,

$$\begin{aligned}
&\forall z(t|k) \in \mathbb{Z} = \mathbb{X} \ominus \mathcal{R} \Rightarrow \mathbb{P}(x(t|k) \subseteq \mathbb{X} \mid x(0|k)) \geq p_{\mathrm{x}} \\
&\forall v(t|k) \in \mathbb{V} = \mathbb{U} \ominus \mathcal{R}^{\mathrm{u}} \Rightarrow \mathbb{P}(u(t|k) \subseteq \mathbb{U} \mid x(0|k)) \geq p_{\mathrm{u}}.
\end{aligned}$$

Since by assumption the disturbance distribution is CCU and the PRS (3.16) or (3.17) are both convex and symmetric, we can apply [77, Thm. 3] and verify closed-loop chance constraint satisfaction, cf. Section 2.3.1 for more details.

Part 2 Let J and $V(\Delta z(k), y^{\mathrm{ref}})$ be defined as in Definition 3.6. We split the expected optimal cost in the cases where mode 1 or mode 2 is applied

$$\begin{aligned}
&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}}) \\
=&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})|M^2)\mathbb{P}(M^2) \\
+&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})|M^1)\mathbb{P}(M^1).
\end{aligned} \tag{3.43}$$

For mode 2 we find

$$\begin{aligned}
&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})|M^2) = V(\Delta z(1|k), y^{\mathrm{ref}}) \\
&\leq J(\Delta\tilde{z}(\cdot|k+1), \Delta\tilde{v}(\cdot|k+1), y_{\mathrm{s}}(k), y^{\mathrm{ref}}),
\end{aligned} \tag{3.44}$$

where $\Delta\tilde{z}(\cdot|k+1) = \tilde{z}(\cdot|k+1) - z_s(k)$ and $\Delta\tilde{v}(\cdot|k+1) = \tilde{v}(\cdot|k+1) - v_s(k)$ (Lemma 3.1) are consistent with the artificial steady output $y_{\mathrm{s}}(k)$. Mode 1 evaluates to

$$\begin{aligned}
&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})|M^1) = \mathbb{E}(V(\Delta x(k+1), y^{\mathrm{ref}})|M^1) \\
&\overset{(3.39)}{\leq} V(\Delta z(1|k), y^{\mathrm{ref}}) + \mathbb{E}(L\|x(k+1) - z(1|k)\|_2|M^1) \\
&\leq J(\Delta\tilde{z}(\cdot|k+1), \Delta\tilde{v}(\cdot|k+1), y_{\mathrm{s}}(k), y^{\mathrm{ref}}) + \underbrace{L/\sqrt{\lambda_{\min}(P)}}_{c}\,\mathbb{E}(\|x(k+1) - z(1|k)\|_P|M^1),
\end{aligned}$$

where the first inequality uses the shifted optimal solution and (3.39), while the second inequality uses the fact that $\lambda_{\min}(P)\|x\|_2^2 \leq \|x\|_P^2$. Adding $c\mathbb{E}(\|x(k+1) - z(1|k)\|_P|M^2)$ to (3.44) and substituting both expressions for mode 1 and 2 in (3.43) yields

$$\begin{aligned}&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}}))\\ \leq& J(\Delta\tilde{z}(\cdot|k+1), \Delta\tilde{v}(\cdot|k+1), y_{\mathrm{s}}(k), y^{\mathrm{ref}}) + c\mathbb{E}(\|x(k+1) - z(1|k)\|_P).\end{aligned}$$

In the following, we use the suboptimality η according to Definition 3.6, which implies

$$\begin{aligned}&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})) - c\mathbb{E}(\|x(k+1) - z(1|k)\|_P)\\ &\leq J(\Delta\tilde{z}(\cdot|k+1), \Delta\tilde{v}(\cdot|k+1), y_{\mathrm{s}}(k), y^{\mathrm{ref}})\\ &\overset{\text{Lem. 3.1}}{\leq} J(\Delta\bar{z}(\cdot|k), \Delta\bar{v}(\cdot|k), y_{\mathrm{s}}(k), y^{\mathrm{ref}}) - \|\Delta z(k)\|_Q^2 - \|\Delta v(k)\|_R^2\\ &\overset{\text{Def. 3.6}}{\leq} V(\Delta z(k), y^{\mathrm{ref}}) - \|\Delta z(k)\|_Q^2 - \|\Delta v(k)\|_R^2 + \eta,\end{aligned} \tag{3.45}$$

where the first inequality is due to the principle of optimality. Furthermore, we simplify $\mathbb{E}(\|x(k+1) - z(1|k)\|_P)$ as

$$\mathbb{E}(\|x(k+1) - z(1|k)\|_P) \leq (1-\kappa)\|e(k)\|_P + \mathbb{E}(\|w(k)\|_P),$$

where $P \succ 0, \kappa \in \mathbb{R}_{>0}$ satisfy the Lyapunov inequality (3.40). After resubstitution of the above inequality into (3.45), we obtain

$$\begin{aligned}&\mathbb{E}(V(\Delta z(k+1), y^{\mathrm{ref}})) - V(\Delta z(k), y^{\mathrm{ref}})\\ \leq& - \|\Delta z(k)\|_Q^2 - \|\Delta v(k)\|_R^2 - \kappa c\|e(k)\|_P + c\mathbb{E}(\|w(k)\|_P^2) + \eta.\end{aligned}$$

Furthermore, we achieve the following asymptotic average cost bound

$$\begin{aligned}&\limsup_{k\to\infty} \frac{1}{k}\mathbb{E}(V(\Delta z(k), y^{\mathrm{ref}}) - V(\Delta z(0), y^{\mathrm{ref}}))\\ &\leq \limsup_{k\to\infty} \frac{1}{k}\sum_{i=0}^{k-1} \mathbb{E}\Big(- \|\Delta z(i)\|_Q^2 - \|\Delta v(i)\|_R^2 - \kappa c\|e(i)\|_P + c\|w(i)\|_P^2\Big) + \eta\\ &\leq \eta + c\cdot \mathrm{tr}(\Sigma_w P).\end{aligned}$$

This implies that the origin $\Delta z = 0$ is, in expectation, a practically asymptotically stable equilibrium point [70, Def. 2.15], which implies that the mean $\mathbb{E}(z(k))$ asymptotically converges to some bounded set and is then merely stable.

To formalize this, we consider the consolidated steady-state (3.33) given by $\bar{z}_{\mathrm{s}}$ and a ball $\mathcal{B}_\eta(\bar{z}_{\mathrm{s}}) \coloneqq \{z \in \mathbb{R}^n \big| \|z - \bar{z}_{\mathrm{s}}\| < \eta\}$ with radius η (due to suboptimality), see Figure 3.7. Then we have that

$$\lim_{k\to\infty} \mathrm{dist}(\mathbb{E}(z(k)), \mathbb{G} \oplus \mathcal{B}_\eta(\bar{z}_{\mathrm{s}})) = 0, \tag{3.46}$$

where the set $\mathbb{G}$ follows from $z_{\mathrm{s}} \in \{\bar{z}_{\mathrm{s}}\} \oplus \mathbb{G}$, cf. Def. 3.3. In other words, the set $\mathbb{G} \oplus \mathcal{B}_{\tilde{\eta}}(\bar{z}_{\mathrm{s}})$ accounts for the solver inexactness in the state prediction, as well as in the steady-state.

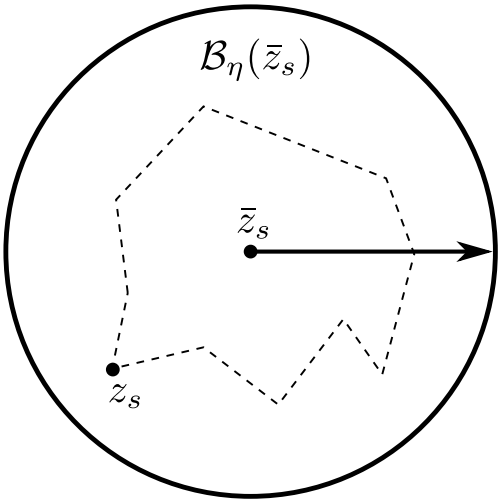

Figure 3.7: Illustration of the ball $\mathcal{B}_\eta(\bar{z}_\mathrm{s})$ centered at the consolidated steady-state $\bar{z}_\mathrm{s}(y^\mathrm{ref})$, while the ϵ-exact steady-state z_s varies in the interior (dotted trajectory).

Remark 3.9. *The set-based convergence result* (3.46) *reduces to a point-wise convergence if the ADMM threshold* $\epsilon \to 0$*, i.e., if* $\epsilon \to 0$*, then the radius* $\eta \to 0$ *and the set* $\mathbb{G}$*, defined through* (3.29)*, vanishes. Therefore, the consolidated steady-state converges to the true steady-state* $\bar{z}_s \to z_s$*, implying that* $\mathbb{E}(z(k)) \to z_s$ *in the limit.*

Part 3 The previous part established that the nominal state converges in expectation to a bounded set $\mathbb{G} \oplus \mathcal{B}_\eta(\bar{z}_\mathrm{s})$. From (3.46) and the output map $\bar{y}_\mathrm{s} = C\bar{z}_\mathrm{s}$ we have that

$$\tilde{y}_\mathrm{s} \in \{\bar{y}_\mathrm{s}\} \oplus C[\mathbb{G} \oplus \mathcal{B}_\eta(\bar{z}_\mathrm{s})], \tag{3.47}$$

where $\tilde{y}_\mathrm{s}$ is some admissible and $\bar{y}_\mathrm{s}$ the consolidated output with associated tracking cost

$$J_\mathrm{t}(\bar{y}_\mathrm{s}, y^\mathrm{ref}) = \|\bar{y}_\mathrm{s} - y^\mathrm{ref}\|_T^2. \tag{3.48}$$

Similar to [104, Thm. 1], we can show that system (3.38b) is only in steady-state if the cost (3.48) for the consolidated state is at its minimum. Then the true steady output verifies (3.47), i.e., the output lies in an area around the consolidated output. Now suppose $y^\mathrm{ref} \in \mathbb{Y}_\mathrm{A}$, then, due to admissibility of y^ref the unique minimum in (3.48) is attained at $\bar{y}_\mathrm{s} = y^\mathrm{ref}$ and the output converges in expectation to a neighborhood of y^ref. □

4 Output-feedback regulation with additive noise

In this chapter, we shift our attention to the output-feedback case, where the goal is to stabilize a predefined steady-state, e.g., the origin. This extends the state-feedback DSMPC formulation to the output-feedback case, which is still an open research question (even in central SMPC) [57]. The main difference compared to a state feedback approach is that we first need to estimate the current system state based on output measurements of the system, which is then used to initialize the DSMPC algorithm. This confronts us with the new challenge of how to deal with different sources of uncertainty, such as measurement and process noise [61]. The output-feedback case is of practical interest since the state vector of a control system is often not fully measurable, e.g., as in distributed power systems [165] or in coordination problems [141].

In Section 4.2, we present a direct output-feedback DSMPC algorithm for distributed systems with emphasis on a distributed design, such that the controller can be synthesized and operated fully distributedly. To satisfy the chance constraints in closed-loop, we extend the distributed PRS approach presented in Chapter 3 to the output-feedback case. In Section 4.3, we propose an indirect output-feedback DSMPC with scenario-based distributed PRS, which yields non-conservative closed-loop constraint satisfaction. This chapter is partly based on the publication [110] [1], i.e., Section 4.2.

Related work Due to the fact that the output-feedback case is an open research topic in central SMPC [57], the literature on distributed systems is almost non-existent. One approach is provided by [46], where the authors propose a non-iterative sequential updating DSMPC, while considering individual dynamics and ICCs for coupled constraints. The design for the controller is fully centralized, while also the operation requires a central warm-start solution. A notable central SMPC approach is proposed by [61], where individual chance constraints are analytically verified with a combined covariance matrix of the process and measurement noise via the Cantelli's inequality. In [33], the authors propose a central probabilistic tube MPC framework for bounded process and measurement noise.

[1] C. Mark and S. Liu. "A stochastic output-feedback MPC scheme for distributed systems". In: *Proc. American Control Conf. (ACC)*. extended version: arXiv:2001.10838. 2020, pp. 1937–1942 © 2020 AACC.

4.1 Problem description

We consider a network of M linear time-invariant systems on a graph (Def. 3.1), where each system $i \in \mathcal{M} = \{1, \dots, M\}$ has a state vector $x_i \in \mathbb{R}^{n_i}$, input vector $u_i \in \mathbb{R}^{m_i}$ and output vector $y_i \in \mathbb{R}^{p_i}$. The local dynamics are governed by

$$x_i(k+1) = A_{\mathcal{N}_i} x_{\mathcal{N}_i}(k) + B_i u_i(k) + w_i(k) \quad \forall k \in \mathbb{N} \tag{4.1a}$$
$$y_i(k) = C_{\mathcal{N}_i} x_{\mathcal{N}_i}(k) + d_i(k), \tag{4.1b}$$

where the distributions of the zero-mean i.i.d. process noise $w_i \in \mathbb{R}^{n_i}$ and zero-mean i.i.d. measurement noise $d_i \in \mathbb{R}^{p_i}$ are assumed to be central convex unimodal with known covariance matrices $\Sigma_i^{\mathrm{w}} \succ 0$ and $\Sigma_i^{\mathrm{d}} \succ 0$. The global system is given by

$$x(k+1) = Ax(k) + Bu(k) + w(k) \quad \forall k \in \mathbb{N} \tag{4.2a}$$
$$y(k) = Cx(k) + d(k), \tag{4.2b}$$

where $x = \mathrm{col}_{i\in\mathcal{M}}(x_i) \in \mathbb{R}^n$, $u = \mathrm{col}_{i\in\mathcal{M}}(u_i) \in \mathbb{R}^m$, $w = \mathrm{col}_{i\in\mathcal{M}}(w_i) \in \mathbb{R}^n$ and $d = \mathrm{col}_{i\in\mathcal{M}}(d_i) \in \mathbb{R}^p$. The global covariance matrices for the process and measurement noise are given by $\Sigma^{\mathrm{w}} = \mathrm{blkdiag}_{i\in\mathcal{M}}(\Sigma_i^{\mathrm{w}})$ and $\Sigma^{\mathrm{d}} = \mathrm{blkdiag}_{i\in\mathcal{M}}(\Sigma_i^{\mathrm{d}})$. Similar to Chapter 3, we impose polytopic state and input chance constraints of the form

$$\mathbb{P}\big(x_{\mathcal{N}_i}(k) \in \mathbb{X}_{\mathcal{N}_i} := \{x_{\mathcal{N}_i} \mid H_{\mathcal{N}_i} x_{\mathcal{N}_i} \le h_{\mathcal{N}_i}\} \big| x(0)\big) \ge p_{\mathrm{x}} \tag{4.3a}$$
$$\mathbb{P}\big(u_i(k) \in \mathbb{U}_i \quad := \{u_i \mid H_i^{\mathrm{u}} u_i \le h_i^{\mathrm{u}}\} \big| x(0)\big) \quad \ge p_{\mathrm{u}}, \tag{4.3b}$$

where $h_{\mathcal{N}_i} \in \mathbb{R}^{r_i}_{>0}$, $h_i^{\mathrm{u}} \in \mathbb{R}^{q_i}_{>0}$ and $p_{\mathrm{x}}, p_{\mathrm{u}} \in (0,1)$ are the levels of chance constraint satisfaction. To simplify the exposition, we impose Assumption 3.1 for the matrix pair (A, B), as well as the following assumption for distributed observability.

Assumption 4.1. *(Structured injection gain) The pair (A, C) is observable with a structured linear observer of the form*

$$\lambda(y) = Ly = \mathrm{col}_{i\in\mathcal{M}}(L_i y_i),$$

where $L_i \in \mathbb{R}^{n_i \times p_i}$, such that $\rho(A - LC) < 1$.

Remark 4.1. *The distributed injection gains L_i can be found via distributed optimization similar to the procedure outlined in Remark 3.1 by setting $(A_{\mathcal{N}_i}, B_i) = (A_{\mathcal{N}_i}^\top, C_{\mathcal{N}_i}^\top)$ for all $i \in \mathcal{M}$.*

The goal of this chapter is to develop two stabilizing DSMPC controllers for the origin of system (4.2a) based on noisy output measurements provided by (4.2b) while satisfying the closed-loop chance constraints (4.3).

4.2 Direct output-feedback DSMPC

In this section, we develop a direct output-feedback DSMPC scheme to stabilize the origin of system (4.2a). We begin by defining a distributed Luenberger observer for the distributed dynamics (4.1a)-(4.1b) that provides estimates $\hat{x}_i$ of the state x_i based solely on local output measurements y_i. The distributed observer dynamics can be expressed as follows:

$$\hat{x}_i(k+1) = A_{\mathcal{N}_i}\hat{x}_{\mathcal{N}_i}(k) + B_i u_i(k) + L_i(y_i(k) - C_{\mathcal{N}_i}\hat{x}_{\mathcal{N}_i}(k)) \quad \forall i \in \mathcal{M}.$$

To stabilize the dynamics (4.2a), we propose a distributed error-feedback control policy

$$u_i(k) = v_i(0|k) + K_{\mathcal{N}_i}(\hat{x}_{\mathcal{N}_i}(k) - z_{\mathcal{N}_i}(0|k)) \quad \forall i \in \mathcal{M}, \tag{4.4}$$

where $z_{\mathcal{N}_i}$ and v_i denote the state and input of the nominal dynamics

$$z_i(t+1|k) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + B_i v_i(t|k).$$

The nominal input sequences $v_i(\cdot|k)$ for all $i \in \mathcal{M}$ are the result of an MPC optimization problem solved at time step k, while $z_{\mathcal{N}_i}(\cdot|k)$ are the corresponding nominal state predictions. Note that, in contrast to the previous chapter, there are now two error sources acting on the system, where the state $x(k)$ can be decomposed into

$$x(k) = z(0|k) + e(k) + \tilde{x}(k),^{2} \tag{4.5}$$

where $\tilde{x}$ denotes the state estimation error and e the observation error, such that

$$\tilde{x}(k) = x(k) - \hat{x}(k) \tag{4.6a}$$
$$e(k) = \hat{x}(k) - z(0|k). \tag{4.6b}$$

The corresponding error dynamics of (4.6) are given by

$$\tilde{x}_i(t+1|k) = A_{\mathcal{N}_i,\mathrm{L}}\tilde{x}_{\mathcal{N}_i}(t|k) + w_i(t|k) - L_i d_i(t|k) \tag{4.7}$$
$$e_i(t+1|k) = A_{\mathcal{N}_i,\mathrm{K}} e_{\mathcal{N}_i}(t|k) + L_i[C_{\mathcal{N}_i}\tilde{x}_{\mathcal{N}_i}(t|k) + d_i(t|k)], \tag{4.8}$$

with $A_{\mathcal{N}_i,\mathrm{L}} = A_{\mathcal{N}_i} - L_i C_{\mathcal{N}_i}$ and $A_{\mathcal{N}_i,\mathrm{K}} = A_{\mathcal{N}_i} + B_i K_{\mathcal{N}_i}$.

By stacking the local states $\tilde{x} = \mathrm{col}_{i\in\mathcal{M}}(\tilde{x}_i)$ and $e = \mathrm{col}_{i\in\mathcal{M}}(e_i)$ and defining an extended state $\xi = \begin{bmatrix}\tilde{x}^\top & e^\top\end{bmatrix}^\top$, we can express the global extended error dynamics compactly as

$$\xi(t+1|k) = \underbrace{\begin{bmatrix} A_{\mathrm{L}} & 0 \\ LC & A_{\mathrm{K}} \end{bmatrix}}_{\Psi} \xi(t|k) + \underbrace{\begin{bmatrix} I & -L \\ 0 & L \end{bmatrix}}_{\Gamma} \omega(t|k), \tag{4.9}$$

where $\omega = \begin{bmatrix} w^\top & d^\top \end{bmatrix}^\top$, $A_{\mathrm{L}} = A - LC$ and $A_{\mathrm{K}} = A + BK$.

[2] The equation similarly holds for the local and neighborhood states, i.e., x_i and $x_{\mathcal{N}_i}$.

4.2.1 Covariance propagation of the extended state

As in Section 3.2.1, the chance constraints are reformulated using mean-variance PRS, with the covariance dynamics being computed based on the extended error (4.9), i.e.,

$$\Sigma^{\xi}(t+1|k) = \operatorname{var}(\xi(t+1|k)) = \Psi\Sigma^{\xi}(t|k)\Psi^{\top} + \Gamma\Sigma^{\omega}\Gamma^{\top}, \tag{4.10}$$

where $\Sigma^{\omega} = \operatorname{blkdiag}(\Sigma^{\mathrm{w}}, \Sigma^{\mathrm{d}})$. By Assumptions 3.1 and 4.1, we know that the matrix Ψ is Schur stable, i.e., due to the lower block triangular structure of the matrix Ψ it suffices to show that the spectral radii $\rho(A+BK)$ and $\rho(A-LC)$ are less then one (similar to the LQG problem). Since the covariance matrices Σ^{w} and Σ^{d} are positive definite by construction, convergence of $\Sigma^{\xi}(t+1|k)$ to a stationary covariance matrix $\Sigma^{\xi}_{\mathrm{f}}$ is guaranteed as $t \to \infty$ for all $k \in \mathbb{N}$. Similar to Section 3.2.1, we pursue the goal of upper bounding the stationary covariance matrix with a structured block-diagonal matrix

$$\hat{\Sigma}^{\xi} = \left[\begin{array}{c|c} \hat{\Sigma}^{\tilde{\mathrm{x}}} & 0 \\ \hline 0 & \hat{\Sigma}^{\mathrm{e}} \end{array}\right] = \left[\begin{array}{ccc|ccc} \hat{\Sigma}^{\tilde{\mathrm{x}}}_1 & \dots & 0 & & & \\ \vdots & \ddots & \vdots & & 0 & \\ 0 & \dots & \hat{\Sigma}^{\tilde{\mathrm{x}}}_M & & & \\ \hline & & & \hat{\Sigma}^{\mathrm{e}}_1 & \dots & 0 \\ & 0 & & \vdots & \ddots & \vdots \\ & & & 0 & \dots & \hat{\Sigma}^{\mathrm{e}}_M \end{array}\right], \tag{4.11}$$

where the two main diagonal blocks correspond to an upper-bound for the state estimation and observation error covariance matrix. In the following, we present two design approaches for $\hat{\Sigma}^{\xi}$.

Distributed stationary covariance matrix – A centralized LMI approach

In the first design approach, we cast (4.10) into its steady-state Lyapunov-like surrogate by taking $t \to \infty$, while constraining the matrix $\hat{\Sigma}^{\xi}_{\mathrm{f}}$ to have only block-diagonal elements akin to (4.11), i.e.,

$$\hat{\Sigma}^{\xi}_{\mathrm{f}} = \Psi\hat{\Sigma}^{\xi}_{\mathrm{f}}\Psi^{\top} + \Gamma\Sigma^{\omega}\Gamma^{\top}.$$

The idea is similar to the approach presented in Section 3.2.1, where we relax the former matrix equality into an LMI constraint of the form

$$\begin{bmatrix} \hat{\Sigma}^{\xi}_{\mathrm{f}} - \Gamma\Sigma^{\omega}\Gamma^{\top} & \Psi\hat{\Sigma}^{\xi}_{\mathrm{f}} \\ \hat{\Sigma}^{\xi}_{\mathrm{f}}\Psi^{\top} & \hat{\Sigma}^{\xi}_{\mathrm{f}} \end{bmatrix} \succeq 0,$$

while the block-diagonal stationary covariance matrix can be found by solving the following optimization problem

$$\begin{aligned}\hat{\Sigma}_{\mathrm{f}}^{\xi *} = \underset{\hat{\Sigma}_{\mathrm{f}}^{\xi} \succeq 0}{\operatorname{argmin}} \quad & \|\hat{\Sigma}_{\mathrm{f}}^{\xi}\|_{\mathrm{F}}^2 \\ \text{s.t.} \quad & \begin{bmatrix} \hat{\Sigma}_{\mathrm{f}}^{\xi} - \Gamma \Sigma^{\omega} \Gamma^{\top} & \Psi \hat{\Sigma}_{\mathrm{f}}^{\xi} \\ \hat{\Sigma}_{\mathrm{f}}^{\xi} \Psi^{\top} & \hat{\Sigma}_{\mathrm{f}}^{\xi} \end{bmatrix} \succeq 0.\end{aligned}$$

The centralized approach inherently assumes that a central coordination node is available that has access to the dynamics and noise covariance matrices of the entire plant. In the following, we present a second design approach, where we remove this assumption to obtain a fully distributed implementation.

Distributed stationary covariance matrix – An iterative distributed approach

We start by writing the individual covariance update equations for the state estimation error $\tilde{x}$ and the observation error e corresponding to the augmented dynamics (4.9)

$$\Sigma^{\tilde{\mathrm{x}}}(t+1) = A_{\mathrm{L}} \Sigma^{\tilde{\mathrm{x}}}(t) A_{\mathrm{L}}^{\top} + \Sigma^{\mathrm{w}} - L \Sigma^{\mathrm{d}} L^{\top} \tag{4.12a}$$

$$\Sigma^{e}(t+1) = A_{\mathrm{K}} \Sigma^{e}(t) A_{\mathrm{K}}^{\top} + LC \Sigma^{\tilde{\mathrm{x}}}(t) C^{\top} L^{\top} + L \Sigma^{\mathrm{d}} L. \tag{4.12b}$$

As stated earlier, the time evolution of these matrices is equivalent for all times $k \in \mathbb{N}$. Therefore, to ease the notation we drop the dependency on k. Even though the matrix pair $(\Sigma^{\tilde{\mathrm{x}}}(t), \Sigma^{e}(t))$ is assumed to be block-diagonal at time t, it is generally not the case at time $t+1$ due to couplings between neighboring systems. Therefore, as proposed in Section 3.2.1, we develop a distributed update scheme based on the block-diagonal matrices $\hat{\Sigma}^{\tilde{\mathrm{x}}}$ and $\hat{\Sigma}^{\mathrm{e}}$ as defined in (4.11), whose evolution is governed by

$$\hat{\Sigma}_i^{\tilde{\mathrm{x}}}(t+1) = \tilde{A}_{\mathcal{N}_i,\mathrm{L}} \hat{\Sigma}_{\mathcal{N}_i}^{\tilde{\mathrm{x}}}(t) \tilde{A}_{\mathcal{N}_i,\mathrm{L}}^{\top} + \Sigma_i^{\mathrm{w}} - L_i \Sigma_i^{\mathrm{d}} L_i^{\top} \qquad \forall i \in \mathcal{M} \tag{4.13a}$$

$$\hat{\Sigma}_i^{\mathrm{e}}(t+1) = \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}^{\mathrm{e}}(t) \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^{\top} + L_i \tilde{C}_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}^{\tilde{\mathrm{x}}}(t) \tilde{C}_{\mathcal{N}_i}^{\top} L_i^{\top} + L_i \Sigma_i^{\mathrm{d}} L_i \quad \forall i \in \mathcal{M}, \tag{4.13b}$$

where $\tilde{A}_{\mathcal{N}_i,\mathrm{L}} = \sqrt{|\mathcal{N}_i|} A_{\mathcal{N}_i,\mathrm{L}}$, $\tilde{A}_{\mathcal{N}_i,\mathrm{K}} = \sqrt{|\mathcal{N}_i|} A_{\mathcal{N}_i,\mathrm{K}}$, $\tilde{C}_{\mathcal{N}_i} = \sqrt{|\mathcal{N}_i|} C_{\mathcal{N}_i}$ and $|\mathcal{N}_i|$ denotes the cardinality of the neighborhood set. The following result is based on [56, Lemma 1].

Corollary 4.1. *Let $\hat{\Sigma}^{\tilde{\mathrm{x}}}$ and $\hat{\Sigma}^{\mathrm{e}}$ be defined as in (4.11) and assume that $\Sigma^{\tilde{\mathrm{x}}}(t) \preceq \hat{\Sigma}^{\tilde{\mathrm{x}}}(t)$ and $\Sigma^{e}(t) \preceq \hat{\Sigma}^{\mathrm{e}}(t)$ at time t. If $\hat{\Sigma}^{\tilde{\mathrm{x}}}(t+1)$ and $\hat{\Sigma}^{\mathrm{e}}(t+1)$ are updated according to (4.13a)-(4.13b) for all $i \in \mathcal{M}$, then also $\Sigma^{\tilde{\mathrm{x}}}(t+1) \preceq \hat{\Sigma}^{\tilde{\mathrm{x}}}(t+1)$ and $\Sigma^{e}(t+1) \preceq \hat{\Sigma}^{\mathrm{e}}(t+1)$.*

Proof. The proof can be found in Section 4.5. □

Based on Corollary 4.1, we present an iterative algorithm (Algorithm 4) that is similar to Algorithm 1 from the previous chapter.

Algorithm 4 Distributed stationary covariance matrix synthesis

For each subsystem $i \in \mathcal{M}$ in parallel:
Initialize $\hat{\Sigma}^{e}_{\mathcal{N}_i}(0) = \hat{\Sigma}^{\tilde{\mathrm{x}}}_{\mathcal{N}_i}(0) = 0$ and $t = 0$
repeat
    Update covariance matrices $(\hat{\Sigma}^{\tilde{\mathrm{x}}}_i(t+1), \hat{\Sigma}^{\mathrm{e}}_i(t+1))$ according to (4.13a), (4.13b)
    Communicate $(\hat{\Sigma}^{\tilde{\mathrm{x}}}_i(t+1), \hat{\Sigma}^{\mathrm{e}}_i(t+1))$ to neighbors $j \in \mathcal{N}_i$
    Construct $\hat{\Sigma}^{\tilde{\mathrm{x}}}_{\mathcal{N}_i}(t+1) = \mathrm{blkdiag}_{j\in\mathcal{N}_i}(\hat{\Sigma}^{\tilde{x}}_j(t+1))$, $\hat{\Sigma}^{e}_{\mathcal{N}_i}(t+1) = \mathrm{blkdiag}_{j\in\mathcal{N}_i}(\hat{\Sigma}^{e}_j(t+1))$
    Increment time $t \leftarrow t+1$
until convergence
Construct stationary covariance matrix $\hat{\Sigma}^{\xi}_{\mathcal{N}_i,\mathrm{f}} = \mathrm{blkdiag}(\hat{\Sigma}^{\tilde{\mathrm{x}}}_{\mathcal{N}_i}(t), \hat{\Sigma}^{\mathrm{e}}_{\mathcal{N}_i}(t)) \quad \forall i \in \mathcal{M}$

Remark 4.2. *As already stated in Remark 3.2, the conservatism of an iterative design akin to Algorithm 4 strongly depends on the topology of the communication graph, as well as on the coupling strength between the subsystems. In the output-feedback case, this issue is even more amplified due to multiple upper bounds that are involved in the update equations* (4.13a)-(4.13b)*. Therefore, the distributed stationary covariance matrix is expected to have much larger eigenvalues than its central counterpart, making the distributed PRS unnecessarily large and leading to conservative satisfaction of the chance constraints, cf. Section 4.2.6.*

4.2.2 Chance constraints via distributed Probabilistic Reachable Sets

In the following, we recall the state separation (4.5) and note that we intend to satisfy the chance constraints (4.3) for the state $x_{\mathcal{N}_i}$ and the input u_i. Therefore, we define a combined random variable $\delta x_{\mathcal{N}_i} = e_{\mathcal{N}_i} + \tilde{x}_{\mathcal{N}_i} = [I \quad I]\, \xi_{\mathcal{N}_i}$ whose distribution is obtained by convolution of the density functions of $e_{\mathcal{N}_i}$ and $\tilde{x}_{\mathcal{N}_i}$ [89, Thm. 14.19]. Based on the zero-mean and CCU assumptions of both random variables, it can be deduced from the property of *closure under convolution*, cf. [50], that the distribution of $\delta x_{\mathcal{N}_i}$ retains zero-mean and CCU. Note that the associated covariance matrix of $\delta x_{\mathcal{N}_i}$ is given by

$$\hat{\Sigma}_{\mathcal{N}_i,\mathrm{f}} = \begin{bmatrix} I & I \end{bmatrix} \hat{\Sigma}^{\xi}_{\mathcal{N}_i,\mathrm{f}} \begin{bmatrix} I & I \end{bmatrix}^{\top}. \tag{4.14}$$

By defining $\delta u_i = K_{\mathcal{N}_i}(\hat{x}_{\mathcal{N}_i} - z_{\mathcal{N}_i}) = [0 \quad K_{\mathcal{N}_i}]\, \xi_{\mathcal{N}_i}$, we have that $\mathbb{E}(\delta u) = 0$ and the input error covariance matrix is given by

$$\hat{\Sigma}^{\mathrm{u}}_{i,\mathrm{f}} = \begin{bmatrix} 0 & K_{\mathcal{N}_i} \end{bmatrix} \hat{\Sigma}_{\mathcal{N}_i,\mathrm{f}} \begin{bmatrix} 0 & K_{\mathcal{N}_i} \end{bmatrix}^{\top}. \tag{4.15}$$

From this point on, we can follow the lines of Section 3.2.2 and define constraint-aligned or marginal PRS for the error state $\delta x_{\mathcal{N}_i}$ and input error δu_i. For the sake of simplicity, we consider the marginal PRS in the following

$$\mathcal{R}_{\mathcal{N}_i} = \left\{ \delta x_{\mathcal{N}_i} \in \mathbb{R}^{n_{\mathcal{N}_i}} \;\middle|\; \left|[\delta x_{\mathcal{N}_i}]_j\right| \leq \sqrt{\gamma_i \, [\hat{\Sigma}_{\mathcal{N}_i,\mathrm{f}}]_{j,j}} \quad \forall j \in \{1, \ldots, n_{\mathcal{N}_i}\} \right\}, \tag{4.16}$$

where $\gamma_i = n_{\mathcal{N}_i}/(1-p_\mathrm{x})$. The input PRS is derived for the input error δu_i, i.e.,

$$\mathcal{R}_i^u = \left\{ \delta u_i \in \mathbb{R}^{m_i} \;\middle|\; \left|[\delta u_i]_j\right| \leq \sqrt{\gamma_i^\mathrm{u}[\hat{\Sigma}_{i,\mathrm{f}}^\mathrm{u}]_{j,j}} \quad \forall j \in \{1,\dots,m_i\} \right\}, \tag{4.17}$$

where $\gamma_i^\mathrm{u} = m_i/(1-p_\mathrm{u})$. Note that γ_i and γ_i^u can be chosen according to Remarks 3.3 and 3.4.

The tightened constraint sets for each subsystem $i \in \mathcal{M}$ are then simply given by the Pontryagin set differences

$$\begin{aligned} \mathbb{Z}_{\mathcal{N}_i} &= \mathbb{X}_{\mathcal{N}_i} \ominus \mathcal{R}_{\mathcal{N}_i} \\ \mathbb{V}_i &= \mathbb{U}_i \ominus \mathcal{R}_i^u. \end{aligned}$$

For details on distributed PRS constraint tightening, please refer to Section 3.2.2. Finally, a global representation of the tightened constraint sets is obtained by the Cartesian product of the local sets

$$\mathbb{Z} = \prod_{i\in\mathcal{M}} \mathbb{Z}_{\mathcal{N}_i}, \quad \mathbb{V} = \prod_{i\in\mathcal{M}} \mathbb{V}_i,$$

both of which are amendable to distributed optimization.

4.2.3 Cost function and terminal ingredients

We consider a stabilizing MPC framework with terminal cost and terminal constraints. To this end, we make the following assumption.

Assumption 4.2. *There exists a terminal cost $V_\mathrm{f}(z) = \sum_{i=1}^M V_{\mathrm{f},i}(z_i) = \sum_{i=1}^M \|z_i\|_{P_i}^2 = \|z\|_P^2$ with block diagonal matrix $P \succ 0$, a distributed terminal controller $v = Kz = \mathrm{col}_{i\in\mathcal{M}}(K_{\mathcal{N}_i} z_{\mathcal{N}_i})$ and a structured terminal set $\mathbb{Z}_\mathrm{f} \subseteq \mathbb{Z}$, such that the following conditions hold for all $z \in \mathbb{Z}_\mathrm{f}$*

$$V_\mathrm{f}((A+BK)z) \leq V_\mathrm{f}(z) - l(z, Kz) \tag{4.18a}$$

$$z \in \mathbb{Z}, \quad Kz \in \mathbb{V} \tag{4.18b}$$

$$(A+BK)z \in \mathbb{Z}_\mathrm{f}. \tag{4.18c}$$

The stage cost $l(z,v) = \sum_{i\in\mathcal{M}} l_i(z_i, v_i)$ is the sum of local stage cost functions

$$l_i(z_i, v_i) = \|z_i\|_{Q_i}^2 + \|v_i\|_{R_i}^2,$$

where $Q_i \succeq 0$, $R_i \succ 0$.

Remark 4.3. *The design of a separable terminal cost function and distributed terminal controllers can be achieved via structured LMIs, cf. [42]. Using a quadratic terminal cost function $V_\mathrm{f}(z)$ as in Assumption 4.3, the structured terminal set $\mathbb{Z}_\mathrm{f}$ is defined as the largest feasible α-level set, i.e.,*

$$\mathbb{Z}_\mathrm{f} = \{z \in \mathbb{R}^n | z^\top P z \leq \alpha\}, \quad \alpha \in \mathbb{R}_{>0},$$

where α can be obtained from a distributed linear program, see Section 5.3.4 for details.

For the MPC optimization problem, we consider a quadratic expected value cost function, which, given Assumption 4.2, can be written as follows

$$\begin{aligned}
&\mathbb{E}\left(\sum_{i=1}^{M}\left[\|x_i(N|k)\|_{P_i}^2+\sum_{t=0}^{N-1}\|x_i(t|k)\|_{Q_i}^2+\|u_i(t|k)\|_{R_i}^2\right]\,\middle|\, x(0|k)\right)\\
&=\sum_{i=1}^{M}\Bigg(\|z_i(N|k)\|_{P_i}^2+\sum_{t=0}^{N-1}\|z_i(t|k)\|_{Q_i}^2+\|v_i(t|k)\|_{R_i}^2\\
&+\operatorname{tr}(P_i\Sigma_i^{\delta\mathrm{x}}(t|k))+\sum_{t=0}^{N-1}\operatorname{tr}(Q_i\Sigma_i^{\delta\mathrm{x}}(t|k))+\operatorname{tr}(K_{\mathcal{N}_i}^{\top}R_iK_{\mathcal{N}_i}\Sigma_{\mathcal{N}_i}^{\mathrm{e}}(t|k))\Bigg),
\end{aligned}\tag{4.19}$$

where $\Sigma_i^{\delta\mathrm{x}}(t|k)=\operatorname{var}(\delta x_i(t|k)\mid x(0|k))$, $\Sigma_{\mathcal{N}_i}^{\mathrm{e}}(t|k)=\operatorname{var}(e_{\mathcal{N}_i}(t|k)\mid x(0|k))$. Note that the last line of (4.19) is independent of the MPC decision variables, since $K_{\mathcal{N}_i}$ are fixed gains for all $i\in\mathcal{M}$ and can therefore be neglected in the receding horizon cost function

$$J(z(\cdot|k),v(\cdot|k))=\sum_{i=1}^{M}\left(\|z_i(N|k)\|_{P_i}^2+\sum_{t=0}^{N-1}\|z_i(t|k)\|_{Q_i}^2+\|v_i(t|k)\|_{R_i}^2\right).\tag{4.20}$$

4.2.4 MPC optimization problem

The following optimization problem is solved via distributed optimization at every time instant $k\in\mathbb{N}$.

Problem 4.2.1 (Direct output-feedback DSMPC problem)**.**

$$\begin{aligned}
\min_{\mathcal{V},\mathcal{Z}}\quad & J(z(\cdot|k),v(\cdot|k)) && && \text{(4.21a)}\\
\text{s.t.}\quad & z(t+1|k)=Az(t|k)+Bv(t|k) && \forall t\in\{0,...,N-1\} && \text{(4.21b)}\\
& (z(t|k),v(t|k))\in\mathbb{Z}\times\mathbb{V} && \forall t\in\{0,...,N-1\} && \text{(4.21c)}\\
& z(N|k)\in\mathbb{Z}_{\mathrm{f}} && && \text{(4.21d)}\\
& z(0|k)\in\{\hat{x}(k),z(1|k-1)\}, && && \text{(4.21e)}
\end{aligned}$$

where $\mathcal{V}=\{v(0|k),\dots,v(N-1|k)\}$ *and* $\mathcal{Z}=\{z(0|k),\dots,z(N|k)\}$ *denote the nominal input and state sequences, respectively.*

Each subsystem $i\in\mathcal{M}$ takes the first element of the state and input sequences and implements them under the control law (4.4) to system (4.2).

To ensure recursive feasibility, we resort to the feasibility-based initialization strategy as proposed in Section 3.2.4, i.e., constraint (4.21e) selects the state estimate $z(0|k)=\hat{x}(k)$ whenever possible (Mode 1), and Mode 2, the backup solution $z(0|k)=z^*(1|k-1)$, when Mode 1 is infeasible.

4.2.5 Theoretical analysis

Before stating the main result, we require a technical assumption to state a Lipschitz-based convergence result.

Assumption 4.3. *The set Ξ of feasible initial states $z(0|k)$ for Problem 4.2.1 is bounded.*

Note that the previous results are presented in a centralized fashion. However, each ingredient is designed in such a way that the distributed structure is preserved, i.e., the constraints, cost function and terminal ingredients are structured. This allows for a distributed implementation as presented in Section 3.3.1, where the MPC optimization problem 4.2.1 can be solved with Algorithm 2. Since the main focus of this section is the treatment of output-feedback regulation problems, we make the following technical assumption.

Assumption 4.4. *The MPC optimization problem 4.2.1 is solved exactly by distributed optimization.*

Remark 4.4. *Assumption 4.4 allows us to prove the main result (Theorem 4.1) in a centralized fashion, i.e., the ADMM-based MPC (Algorithm 2) converges with $\epsilon = 0$. This simplifies the proof greatly, since no inexact minimization has to be considered. However, the proof can easily be extended to the inexact case ($\epsilon > 0$) by considering similar techniques as proposed in Section 3.3.1 without compromising the closed-loop guarantees.*

Theorem 4.1. *Let Assumptions 3.1, 4.1-4.4 hold. If the direct output-feedback MPC optimization problem 4.2.1 admits a feasible solution at time $k = 0$, then it is recursively feasible and the chance constraints* (4.3) *are satisfied in closed-loop for any $k \in \mathbb{N}$ with convex symmetric PRS* (4.16) - (4.17)*. In addition, conditioned on $\delta x(0) = \delta u(0) = 0$, the controller achieves the following asymptotic average cost*

$$\lim_{T\to\infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}(\|x(k)\|_Q^2 + \|u(k)\|_R^2) \leq \gamma \sqrt{\mathrm{tr}(\Gamma^\top P_\Sigma \Gamma \Omega)},$$

where $\gamma = \frac{\sqrt{2}\beta}{\sqrt{\lambda_{\min}(P_\Sigma)}} > 0$, $Q = \mathrm{blkdiag}_{i\in\mathcal{M}}(Q_i)$, $R = \mathrm{blkdiag}_{i\in\mathcal{M}}(R_i)$, β denotes a Lipschitz constant and $P_\Sigma \succ 0$ the solution of the Lyapunov inequality $\Psi^\top P_\Sigma \Psi \preceq P_\Sigma - \kappa I$ for some $\kappa \in \mathbb{R}_{>0}$.

Proof. The proof can be found in Section 4.5. □

4.2.6 Numerical example

This section is dedicated to a brief numerical example. We consider $M = 5$ subsystems in a chain as depicted in Figure 4.1. For subsystem 1 it holds that $\mathcal{N}_1 = \{1, 2\}$, for subsystem 5 it

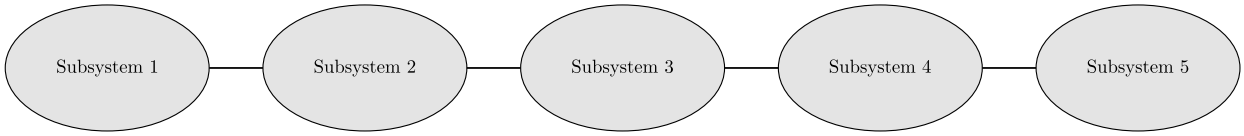

Figure 4.1: Five subsystems in a chain with a bidirectional communication graph.

holds that $\mathcal{N}_5 = \{4, 5\}$ and for the subsystems $i = \{2, 3, 4\}$ it holds that $\mathcal{N}_i = \{i-1, i, i+1\}$. We consider dynamic matrices $A_{ii} = \left[\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right], A_{ij} = \left[\begin{smallmatrix} 0.05 & 0 \\ 0 & 0.05 \end{smallmatrix}\right], \forall j \in \mathcal{N}_i \backslash \{i\}, \forall i \in \mathcal{M}$, input matrices $B_i = \left[\begin{smallmatrix} 0.5 \\ 1 \end{smallmatrix}\right], \forall i \in \mathcal{M}$ and output matrices $C_{ii} = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right], C_{ij} = \left[\begin{smallmatrix} 0 & 0 \\ 0 & 0 \end{smallmatrix}\right], \forall i \in \mathcal{M}$. Each subsystem is subject to a normally distributed process noise with $\Sigma_i^{\mathrm{w}} = 0.01I$ and a normally distributed measurement noise with $\Sigma_i^{\mathrm{d}} = 0.001$. Each subsystem has to satisfy a joint chance constraint on the second state $\mathbb{P}(|[x_i]_2| \leq 1) \geq 0.6$. The weighting matrices are set to $Q_i = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 0.1 \end{smallmatrix}\right]$, $R_i = 0.1$ and the prediction horizon is $N = 10$. The distributed controller and injection gains satisfying Assumptions 3.1 and 4.1 are computed along Remark 4.1. For simplicity, the terminal set is set to $\mathbb{Z}_{\mathrm{f}} = \{0\}$.

Comparison of covariance matrix design procedures The first case study targets the design techniques presented in Section 4.2.1 to highlight the differences between the centralized LMI and the distributed iterative approach. We compare for both methods the resulting distributed PRS volume given by

$$V = \sum_{i=1}^{M} 2\sqrt{\gamma_i [\hat{\Sigma}^e_{\mathcal{N}_i,\mathrm{f}}]_{2,2}},$$

where $\gamma_i = \mathcal{X}_1^2(p_{\mathrm{x}})$ is set according to Remark 3.3 due to normality of the noise distributions. In Table 4.1, we compare for different design procedures the resulting volume of the PRS (smaller is better). The *LMI approach* and the *iterative approach* are both distributed PRS designs, i.e., we use the block-diagonal covariance matrix as presented in Section 4.2.1, while the *central unstructured approach* is a centralized PRS design that uses directly the dense stationary covariance matrix $\Sigma_{\mathrm{f}}^{\xi}$ without block-diagonal upper bounding.

As it can be seen in Table 4.1, the central unstructured design yields the least conservative PRS, which is due to the dense covariance matrix. By imposing a distributed structure, the LMI approach yields a slightly more conservative PRS, while the iterative approach almost doubles the PRS volume compared to the LMI approach. This issue was already mentioned in Remark 4.2, which mainly results from the conservative upper bounding of the local covariance matrices during the update steps. In Figure 4.2, we quantitatively investigate

Table 4.1: Comparison of different PRS volumes regarding their design procedure.

	LMI approach	Iterative approach	Central unstructured approach
Volume	2.5620	4.8489	1.9078

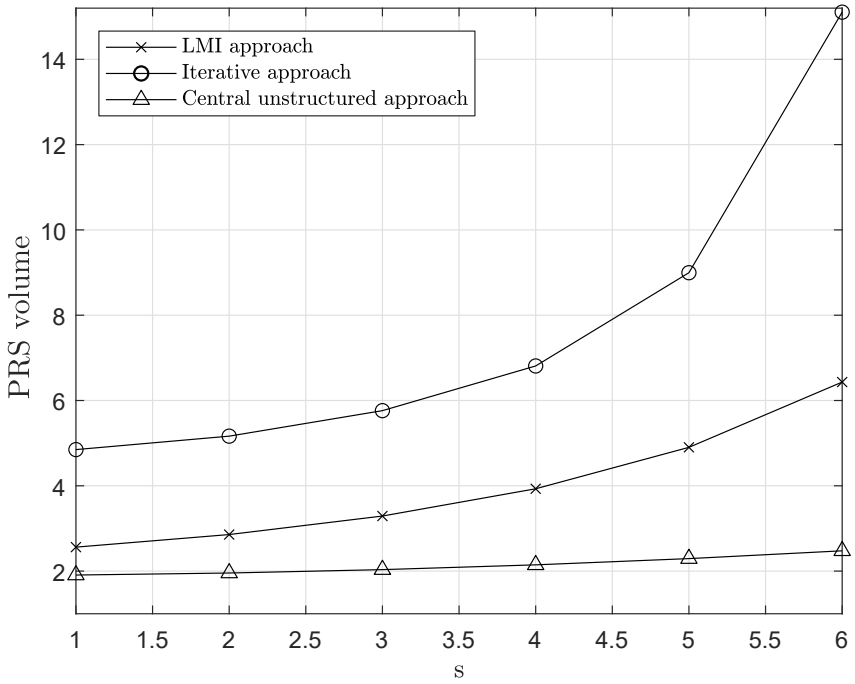

Figure 4.2: Effect of coupling strength s on PRS volume.

the effect of the coupling strength $s \in \mathbb{R}$ on the PRS volume, where we parameterize the matrices A_{ij} as sA_{ij} and compute the PRS volume for each design approach. It can be seen that the iterative approach grows quickly in size by increasing the coupling strength in comparison to the LMI approach, while the central unstructured approach is only marginally affected. This result reiterates the point that a centralized approach should be taken for the design of the distributed stationary covariance matrix whenever possible.

Performance comparison We carry out $K = 5000$ Monte-Carlo simulations of the closed-loop system for 10 time steps each starting from the initial conditions $x_1(0) = [-4, 0]^\top$, $x_5(0) = [4, 0]^\top$ and $x_2(0) = x_3(0) = x_4(0) = [0, 0]^\top$. In Table 4.2, we compare the expected closed-loop performance

$$\mathbb{E}(l(x, u)) = \sum_{j=1}^{K} \sum_{k=1}^{10} \frac{1}{10K} l(x(k), u(k))$$

and the empirical worst-case in-time constraint satisfaction rate for subsystem 1 and 5 for different controller types. We benchmark our proposed DSMPC with PRS design based on the centralized LMI approach and iterative distributed approach against a centralized MPC (C-MPC) with central unstructured PRS design, a distributed LQG (D-LQG) and central LQG (C-LQG). The controller and injection gains for the central controller setups are obtained as the solution to the linear quadratic control and estimation problem. Note

Table 4.2: Comparison between central and distributed setups.

	C-LQG	D-LQG	C-MPC	DMPC: LMI approach	DMPC: Iterative approach
$\mathbb{E}(l(x,u))$	51.34	51.37	56.31	60.65	68.00
$\mathbb{P}(\lvert[x_1]_2\rvert \leq 1)$	0%	0%	85.00%	91.32%	95.44%
$\mathbb{P}(\lvert[x_5]_2\rvert \leq 1)$	0%	0%	85.16%	90.72%	95.60%

that all MPC approaches above yield closed-loop constraint satisfaction. As expected, the central MPC implementation yields the least conservative constraint satisfaction rate of about 85%, which, however, is much higher than the prescribed level of 60%. This gap arises from the direct feedback feasibility-based initialization scheme associated with PRS-based SMPC design, cf. Remark 2.3. In case of the proposed distributed PRS, this issue is even further amplified due to the block-diagonal upper bounding procedure, cf. Figure 4.2, which translates to an even more conservative constraint satisfaction rate.

This issue gives rise to alternative SMPC schemes, such as the indirect feedback SMPC [76], which can close the gap between prescribed and observed closed-loop constraint satisfaction probability, while still maintaining the strong closed-loop chance constraint guarantees, cf. the discussion in Remark 2.4.

4.3 Indirect output-feedback DSMPC: A distributed scenario PRS approach

As we have seen in the previous section, a distributed design is associated with large conservatism regarding closed-loop chance constraint satisfaction. In the following, we extend the direct feedback formulation to an indirect one, with the goal of achieving non-conservative closed-loop chance constraint satisfaction through scenario-based distributed PRS [79]. A similar distributed scenario-based approach has been presented by [128] for the state-feedback case.

To this end, instead of solving Problem 4.2.1, we aim to solve the following indirect output-feedback problem.

Problem 4.3.1 (Indirect output-feedback DSMPC)**.**

$$\begin{aligned}
\min_{\mathcal{V},\mathcal{Z}} \quad & J(\hat{\bar{x}}(\cdot|k), \bar{u}(\cdot|k)) && \text{(4.22a)}\\
\text{s.t.} \quad & z(t+1|k) = Az(t|k) + Bv(t|k) && \text{(4.22b)}\\
& \hat{\bar{x}}(t+1|k) = A\hat{\bar{x}}(t|k) + B\bar{u}(t|k) && \text{(4.22c)}\\
& \bar{e}(t|k) = \hat{\bar{x}}(t|k) - z(t|k) && \text{(4.22d)}\\
& \bar{u}(t|k) = K\bar{e}(t|k) + v(t|k) && \text{(4.22e)}\\
& (z(t|k), v(t|k)) \in \mathbb{Z}^{\mathrm{s}}(k+t) \times \mathbb{V}^{\mathrm{s}}(k+t) && \text{(4.22f)}\\
& z(N|k) \in \mathbb{Z}_{\mathrm{f}}^{\mathrm{s}} && \text{(4.22g)}\\
& \hat{\bar{x}}(0|k) = \hat{x}(k), z(0|k) = z(k), && \text{(4.22h)}
\end{aligned}$$

for all $t = 0, \dots, N-1$*, where* $\mathcal{V} = \{v(0|k), \dots, v(N-1|k)\}$ *and* $\mathcal{Z} = \{z(0|k), \dots, z(N|k)\}$ *denote the nominal input and state sequences, respectively.* [3]

The above problem is connected to system (4.2a) through the control input (4.4), where $z(k)$ is the nominal closed-loop state and $\hat{x}(k)$ the observer state initialized with $z(0) = \hat{x}(0) = x(0)$. Therefore, similar to (4.5), we can decompose the state vector x into $x(k) = z(k) + e(k) + \tilde{x}(k)$, which is governed by the closed-loop dynamics

$$\begin{aligned}
z(k+1) &= Az(k) + Bv(k)\\
\tilde{x}(k+1) &= (A - LC)\tilde{x}(k) + w(k) - Ld(k) && \text{(4.23a)}\\
e(k+1) &= (A + BK)e(k) + LC\tilde{x}(k) + Ld(k). && \text{(4.23b)}
\end{aligned}$$

Note that (4.22b) is always initialized with the nominal closed-loop state via (4.22h), which results in linear closed-form expressions for the state estimation and observer error dynamics $\tilde{x}(k)$ and $e(k)$. Hence, both errors evolve autonomously from the chosen MPC control input. Nevertheless, feedback is introduced indirectly through the cost function (4.22a) via the predicted mean (4.22c), since we always initialize this sequence with the observer state, i.e., $\hat{\bar{x}}(0|k) = \hat{x}(k)$ through (4.22h). Finally, the constraints (4.22f) are tightened for each time step k individually via k-step PRS for the combined random variable $\delta x(k) = e(k) + \tilde{x}(k)$, which ensures maximum flexibility and allows for inclusion of correlated disturbance sequences.

4.3.1 Scenario optimization – A brief recap

In contrast to the direct output-feedback case, we aim to design the constraint sets $\mathbb{Z}^{\mathrm{s}}$ and $\mathbb{V}^{\mathrm{s}}$ in (4.22f) via scenario-based PRS. We consider techniques from scenario optimization [28] that solve chance constrained problems of the form

$$\begin{aligned}
\min_{x \in \mathbb{X} \subseteq \mathbb{R}^d} \quad & c^\top x && \text{(4.24a)}\\
\text{s.t.} \quad & \mathbb{P}(x \in \mathbb{X}_\delta) \geq p, && \text{(4.24b)}
\end{aligned}$$

[3] Variables with a bar denote the mean values, i.e., $\bar{e}(t|k) = \mathbb{E}(e(t|k)|\hat{x}(k))$.

where $\mathbb{X}_\delta$ is a convex and closed set for any realization of the random variable δ. By drawing N_s samples $\delta^{(i)}$ of δ we can approximate problem (4.24) as its sample-based surrogate

$$\min_{x \in \mathbb{X} \subseteq \mathbb{R}^d} \quad c^\top x \tag{4.25a}$$

$$\text{s.t.} \quad x \in \mathbb{X}_{\delta^{(i)}} \quad \forall i \in \mathcal{I}_\mathrm{s}, \tag{4.25b}$$

where $\mathcal{I}_\mathrm{s}$ denotes a subset of all samples with cardinality $|\mathcal{I}_\mathrm{s}| = N_\mathrm{s} - N_\mathrm{d}$, which is found by discarding N_d samples from the original set $\mathcal{I} = \{1, \ldots, N_\mathrm{s}\}$. To use a simple greedy sample removing technique, we make the following technical assumption.

Assumption 4.5 ([28])**.** *The optimal solution x^* of the sample-based problem* (4.25) *violates all N_d discarded constraints $\mathbb{X}_{\delta^{(i)}}$ with $i \in \mathcal{I} \backslash \mathcal{I}_\mathrm{s}$.*

Thus, by discarding constraints (samples), we can improve the objective function value (4.25a) while maintaining probabilistic guarantees of the optimal solution w.r.t. chance constraint satisfaction. To quantify the violation probability $\mathbb{P}(x^* \notin \mathbb{X}_\delta)$, we recall the following result from [28].

Theorem 4.2 ([28, Thm. 2.1])**.** *Let $\beta \in (0, 1)$ be a small confidence parameter, let $x \in \mathbb{X} \subseteq \mathbb{R}^d$ and choose N_s and N_d, such that*

$$\binom{N_\mathrm{d} + d - 1}{N_\mathrm{d}} \sum_{i=0}^{N_\mathrm{d}+d-1} \binom{N_\mathrm{s}}{i} (1-p)^i p^{N_\mathrm{s}-i} \leq \beta, \tag{4.26}$$

then w.r.t. the N_s-fold product measure we have that $\mathbb{P}^{N_\mathrm{s}}(\mathbb{P}(x^* \in \mathbb{X}_\delta) \geq p) \geq 1 - \beta$.

A sufficient condition for (4.26) is provided by the authors of [28], i.e.,

$$N_\mathrm{d} \leq (1-p)N_\mathrm{s} - d + 1 - \sqrt{2(1-p)N_\mathrm{s} \ln\left(\frac{((1-p)N_\mathrm{s})^{d-1}}{\beta}\right)}, \tag{4.27}$$

which gives us an approximate number of samples to remove, such that the optimal solution $x^* \in \mathbb{X}_\delta$ is verified with a probability of at least $1 - \beta$. By setting $N_\mathrm{d} = 0$ and solving for N_s, we find

$$N_\mathrm{s} \geq \frac{2}{1-p}((d-1)\ln(2) - \ln(\beta))$$

as an estimate of how many samples N_s are required to ensure that $\mathbb{P}^{N_\mathrm{s}}(\mathbb{P}(x^* \in \mathbb{X}_\delta) \geq p) \geq 1 - \beta$ without removing any samples.

4.3.2 Distributed scenario PRS

In view of the indirect feedback initialization, we know that the closed-loop error dynamics (4.23a) - (4.23b) evolve autonomously and do not depend on the MPC control input. Thus, we can sample the closed-loop error dynamics offline under the assumptions that we know the duration of the control task, i.e., the task horizon, and the distribution of w and d in (4.2a)-(4.2b) is known exactly.

Assumption 4.6. *The finite task horizon $N_{\mathrm{T}} \in \mathbb{N}$ is known a-priori.*

The design relies on a distributed error sampling algorithm, which is capable of computing the state estimation error and observer error trajectories fully distributed in a neighbor-to-neighbor fashion.

Algorithm 5 Distributed error sampling

1: For each subsystem $i \in \mathcal{M}$ in parallel:
2: Initialize $e_{\mathcal{N}_i}(0) = 0$, $\tilde{x}_{\mathcal{N}_i}(0) = 0$ and $k = 0$
3: **repeat**
4: Sample successor error states $\tilde{x}_i(k+1)$ and $e_i(k+1)$ according to (4.7), (4.8)
5: Communicate $(\tilde{x}_i(k+1), e_i(k+1))$ to neighbors $j \in \mathcal{N}_i$
6: Construct $\tilde{x}_{\mathcal{N}_i}(k+1) = \mathrm{col}_{j\in\mathcal{N}_i}(\tilde{x}(k+1))$ and $e_{\mathcal{N}_i}(k+1) = \mathrm{col}_{j\in\mathcal{N}_i}(e(k+1))$
7: Increment time $k \leftarrow k+1$
8: **until** $k = N_{\mathrm{T}}$

The result of Algorithm 5 is the j-th closed-loop state estimation and observer error trajectory $(\tilde{x}^{(j)}_{\mathcal{N}_i}(k), e^{(j)}_{\mathcal{N}_i}(k))$ for $k = 0, \dots, N_{\mathrm{T}}$. The main idea is to sample a sufficient amount of error scenarios along Algorithm 5, such that each subsystem contains N_{s} scenarios. Afterwards, each subsystem $i \in \mathcal{M}$ solves for each time instant $k = 0, \dots, N_{\mathrm{T}}$ a scenario program for the combined error sample $\delta x^{(j)}_{\mathcal{N}_i}(k) = \tilde{x}^{(j)}_{\mathcal{N}_i}(k) + e^{(j)}_{\mathcal{N}_i}(k)$, i.e.,

$$\min_{b(k)>0} \quad \|b(k)\|_1 \tag{4.28a}$$

$$\text{s.t.} \quad \|H_{\mathcal{N}_i}\delta x^{(j)}_{\mathcal{N}_i}(k)\|_\infty \leq b(k) \quad \forall j \in \mathcal{I}_{\mathrm{s}}. \tag{4.28b}$$

A simpler solution is found thanks to Assumption 4.5, which allows us to successively remove N_{d} samples $\delta x^{(j)}_{\mathcal{N}_i}(k)$ with the largest violation $\|H_{\mathcal{N}_i}\delta x^{(j)}_{\mathcal{N}_i}(k)\|_\infty$. The optimal solution to problem (4.28) is then readily given by $b^*(k) = \max_{j\in\mathcal{I}_{\mathrm{s}}} H_{\mathcal{N}_i}\delta x^{(j)}_{\mathcal{N}_i}(k)$. The result of the optimization problem formally yields a so-called k-step PRS.

Corollary 4.2. *Let N_{s} and N_{d} satisfy condition (4.26) for some probability level $p \in (0,1)$, confidence $\beta \in (0,1)$, constraint function dimension $d = r_i$, and let $b^*(k)$ be the optimal solution of problem (4.28) for $k \in \{0, \dots, N_{\mathrm{T}}\}$. With probability $1-\beta$ the set $\mathcal{R}^s_{\mathcal{N}_i}(k) = \{\delta x_{\mathcal{N}_i} \,|\, H_{\mathcal{N}_i}\delta x_{\mathcal{N}_i}(k) \leq b(k)\}$ is a k-step PRS of probability level p for the combined stochastic error part $\delta x_{\mathcal{N}_i} = \tilde{x}_{\mathcal{N}_i} + e_{\mathcal{N}_i}$ at time k initialized with $\delta x_{\mathcal{N}_i}(0) = 0$.*

Proof. The proof follows immediately by applying Theorem 4.2 to the scenario program (4.28) with optimal solution $b^*(k)$ for $k = 0, \dots, N_\mathrm{T}$. □

Analogous to the k-step distributed PRS for the state error $\delta x_{\mathcal{N}_i}(k)$, we can derive a k-step distributed PRS $\mathcal{R}^s_{\mathrm{u},i}(k)$ for the input error $e_{\mathrm{u},i}(k) = K_{\mathcal{N}_i}\tilde{x}_{\mathcal{N}_i}(k)$, which follows from the tube controller (4.4).

Remark 4.5. *If the distributions of w and d are not available, instead of sampling via Algorithm 5, we can directly use historical data of $\delta x_{\mathcal{N}_i}(k)$ recorded during closed-loop operation and apply Corollary 4.2 to obtain a scenario PRS.*

Remark 4.6. *So far, we have established a result for the joint chance constraints* (4.3)*. Note that JCCs were only necessary to ensure closed-loop constraint satisfaction in the direct feedback case, i.e., along convex symmetric PRS, see Theorem 4.1. In the sample-based indirect feedback framework, we can drop this requirement, which allows us to use non-symmetric PRS that are defined on single half-spaces and yield non-conservative chance constraint satisfaction. In view of Corollary 4.2, we can obtain a k-step PRS for the closed-loop system by simply setting $d = 1$ and deriving the value $b^*(k)$ for each half-space constraint individually. Hence, we can decompose $H_{\mathcal{N}_i}$ into its row vectors $[H_{\mathcal{N}_i}]_c$ for $c = 1, \dots, r_i$, such that $[H_{\mathcal{N}_i}]_c \delta x_{\mathcal{N}_i}(k) \leq [b_i(k)]_c$ and define the alternative optimization problem*

$$\min_{[b_i(k)]_c > 0} \quad [b_i(k)]_c \tag{4.29a}$$

$$s.t. \quad [H_{\mathcal{N}_i}]_c \delta x^{(j)}_{\mathcal{N}_i}(k) \leq [b_i(k)]_c \quad \forall j \in \mathcal{I}_\mathrm{s}, \quad \forall c \in \{1, \dots, r_i\}. \tag{4.29b}$$

This can similarly be solved thanks to Assumption 4.5 by successively removing N_d samples $\delta x^{(j)}_{\mathcal{N}_i}(k)$ with the highest values $[H_{\mathcal{N}_i}]_c \delta x^{(j)}_{\mathcal{N}_i}(k)$.

In Algorithm 6, we propose a distributed constraint tightening procedure that results in a dynamic constraint tightening along the closed-loop trajectory for $k = 0, \dots, N_\mathrm{T}$. This can

Algorithm 6 Distributed constraint tightening

1: **Input**: Chance constraints (4.3), Confidence β, Sample size N_s
2: Distributedly generate N_s samples via Algorithm 5
3: Initialize $k = 0$
4: For each subsystem $i \in \mathcal{M}$ in parallel:
5: **repeat**
6: Compute the number of samples to discard N_d via (4.27)
7: Solve optimization problem (4.28) or (4.29) and obtain scenario-based PRS $\mathcal{R}^\mathrm{s}_{\mathcal{N}_i}(k)$ and/or $\mathcal{R}^\mathrm{s}_{\mathrm{u},i}(k)$ via Corollary 4.2.
8: Tighten constraints with $\mathbb{Z}^\mathrm{s}_{\mathcal{N}_i}(k) = \mathbb{X}_{\mathcal{N}_i} \ominus \mathcal{R}^\mathrm{s}_{\mathcal{N}_i}(k)$ and $\mathbb{V}^\mathrm{s}_i(k) = \mathbb{U}_i \ominus \mathcal{R}^\mathrm{s}_{\mathrm{u},i}(k)$
9: Increment time $k \leftarrow k + 1$
10: **until** $k = N_\mathrm{T}$

be used in case of correlated disturbance sequences. The global constraint sets are given by the Cartesian products

$$\mathbb{Z}^{\mathrm{s}}(k) = \Pi_{i\in\mathcal{M}}\mathbb{Z}^{\mathrm{s}}_{\mathcal{N}_i}(k)$$
$$\mathbb{V}^{\mathrm{s}}(k) = \Pi_{i\in\mathcal{M}}\mathbb{V}^{\mathrm{s}}_{i}(k)$$

for all $k = 0, \ldots, N_{\mathrm{T}}$.

Remark 4.7. *If the noise terms $w(k)$ and $d(k)$ are i.i.d., the assumption on the task horizon (Assumption 4.6) can be dropped, since any sample $\delta x^{(j)}(k) = \tilde{x}^{(j)}(k) + e^{(j)}(k)$ for all $k \in \mathbb{N}$ is drawn from the same distribution. This optionally simplifies the design procedure* (4.28) *or* (4.29)*, since the time dependency k is no longer necessary, resulting in a PRS instead of k-step PRS.*

Objective function In general, in indirect feedback schemes, the cost function can be reformulated analytically, cf. Section 2.3.1, or sample-based, cf. Section 2.3.2. While the latter facilitates the use of arbitrary cost functions, analytical approximation schemes require a quadratic cost. To obtain a comparable result to the direct output-feedback case from the previous section, we consider a quadratic cost function in the state and input mean, cf. [76], while the receding horizon cost is defined as

$$J(\hat{\bar{x}}(\cdot|k), \bar{u}(\cdot|k)) = \|\hat{\bar{x}}(N|k)\|_P^2 + \sum_{t=0}^{N-1} \|\hat{\bar{x}}(t|k)\|_Q^2 + \|\bar{u}(t|k)\|_R^2, \tag{4.30}$$

where the matrices Q, R, P are block-diagonal matrices defined as in Assumption 4.2. Note that, similar to the direct output-feedback case, we have omitted the variance part of the cost function, cf. (4.19), as it cannot be improved by the MPC decision variables. The difference between the cost function in *direct feedback* (4.19) compared to *indirect feedback* (4.30) is that we now use the mean predictions $\hat{\bar{x}}(\cdot|k)$ and $\bar{u}(\cdot|k)$, while (4.19) uses the nominal predictions $z(\cdot|k)$ and $v(\cdot|k)$. In other words, the mean predictions include the tube feedback, while the nominal ones do not. Finally, to ensure recursive feasibility, we consider a terminal set/controller approach, where we make the following assumption.

Assumption 4.7. *Let Assumption 4.2 hold true with $\mathbb{Z} = \mathbb{Z}^{\mathrm{s}}_{\mathrm{f}}$ and $\mathbb{V} = \mathbb{V}^{\mathrm{s}}_{\mathrm{f}}$, where $\mathbb{Z}^{\mathrm{s}}_{\mathrm{f}} = \cap_{k=1}^{N_{\mathrm{T}}}\mathbb{Z}^{\mathrm{s}}(k)$ and $\mathbb{V}^{\mathrm{s}}_{\mathrm{f}} = \cap_{k=1}^{N_{\mathrm{T}}}\mathbb{V}^{\mathrm{s}}(k)$.*

4.3.3 Theoretical analysis

We can establish the following result on recursive feasibility and closed-loop chance constraint satisfaction.

Proposition 4.3.2. *Consider system* (4.2) *under control law* (4.4) *resulting from the indirect feedback MPC optimization Problem 4.3.1 satisfying Assumptions 3.1, 4.1, 4.4, 4.6 and 4.7.*

If Problem 4.3.1 admits a feasible solution with $z(0) = \hat{x}(0) = x(0)$, then it is feasible for all $k \in \{0, \dots, N_\mathrm{T} - N\}$ and with a probability of no less than $1 - \beta$ the chance constraints (4.3) *are satisfied in closed-loop.*

Proof. The proof can be found in Section 4.5. □

Remark 4.8. *Under a zero-mean i.i.d. assumption on $w(k)$ and $d(k)$, we can remove Assumption 4.6 in Proposition 4.3.2 and prove the same theoretical properties for all $k \in \mathbb{N}$. This, however, is omitted for brevity and can easily established by following the same steps of the proof (Section 4.5) with time-invariant constraint sets $\mathbb{Z}^\mathrm{s}, \mathbb{V}^\mathrm{s}, \mathbb{Z}_\mathrm{f}^\mathrm{s}$ due to PRS constraint tightening (instead of k-step PRS).*

Moreover, an expected cost decrease and convergence to an asymptotic average performance bound, similar to [76, Thm. 3, Cor. 1], can be established. Notably, thanks to the cost function (4.30) *for the mean predictions, the resulting performance bound is no worse than that from the linear controller $u = K\hat{x}$, which improves the bound from Theorem 4.1 by setting $\gamma = 1$.*

4.3.4 Numerical example continued

We continue the example from Section 4.2.6 to illustrate that the indirect feedback MPC scheme verifies the closed-loop chance constraints non-conservatively. For comparison reasons, we use a constant PRS constraint tightening by computing only one PRS over the entire task horizon $N_\mathrm{T} = 10$ (Remark 4.7), which is justified since the process and measurement noise are both i.i.d. in the considered example.

To do so, we compute a PRS at time $k = N_\mathrm{T}$ with Algorithm 6 and define $\mathbb{Z}_{\mathcal{N}_i}(k) = \mathbb{Z}_{\mathcal{N}_i}(N_\mathrm{T})$ for all $k = 0, \dots, N_\mathrm{T}$. In the following, we consider $N_\mathrm{s} = 10^6$, $\beta = 10^{-9}$ and $p_\mathrm{x} = 0.6$ for each experiment. All offline computations were carried out on standard hardware, i.e., an Intel i7-9700k CPU with 16gb of RAM, within 60 seconds.

PRS volume We design a non-symmetric scenario-based PRS with optimization problem (4.29) for the case of ICCs, which yields a volume of 0.584, while the symmetric PRS resulting from optimization problem (4.28) for the case of JCCs yields a volume of 1.218. In both cases, the scenario PRS are significantly smaller compared to the analytical counterparts, cf. Table 4.1, where the scenario PRS with JCCs is approximately 36% and the scenario PRS with ICCs nearly 70% smaller compared to the central analytical design.

Closed-loop results In view of Remark 4.6, we formulate the PRS based on the individual half-space constraints, which results in non-conservative closed-loop constraint satisfaction. To this end, we consider the same setup as in Section 4.2.6 and simulate the closed-loop system over the task horizon $N_\mathrm{T} = 10$. To facilitate comparison with the results of the

Table 4.3: Comparison between direct and indirect feedback DMPC.

	C-MPC	DMPC: LMI approach	DMPC: Iterative approach	IF-DMPC
$\mathbb{E}(l(x,u))$	56.31	60.65	68.00	48.67
$\mathbb{P}(\lvert[x_1]_2\rvert \leq 1)$	85.00%	91.32%	95.44%	60.06%
$\mathbb{P}(\lvert[x_5]_2\rvert \leq 1)$	85.16%	90.72%	95.60%	60.14%

previous section, we have extended Table 4.2 with the results from the indirect feedback approach (IF-DMPC), see Table 4.3. Figure 4.3 illustrates the mean trajectories for different DMPC approaches, averaged over 5000 Monte-Carlo runs. As it can be seen, the indirect feedback DMPC operates on average the closest to the constraint $[x_1]_2 \leq 1$, while the direct feedback approaches with feasibility-based initialization have inherently larger conservatism due to the conditioning of the constraints on the state $x(0|k)$ instead of the closed-loop error $e(0)$, cf. the discussion in Remarks 2.3 and 2.4.

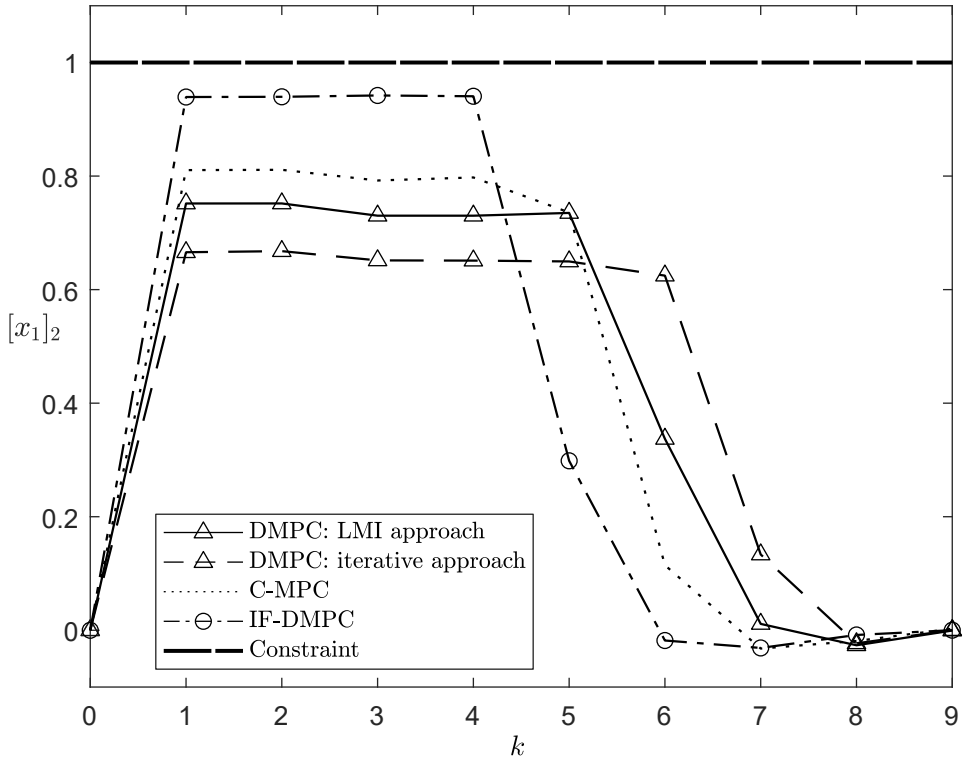

Figure 4.3: Mean trajectories of subsystem 1 averaged over 5000 Monte-Carlo runs.

4.4 Summary

In this chapter, we studied an output-feedback SMPC scheme for distributed systems using distributed PRS. First, we extended a state-feedback distributed SMPC algorithm to the output-feedback case, where we proposed two synthesis methods for distributed PRS that combine the effect of process and measurement noise via structured block-diagonal stationary covariance matrices. The central LMI-based approach yields a less conservative distributed PRS compared to the distributed iterative approach. However, the main advantage of the iterative distributed approach is that no central coordination node is required. The output-feedback DSMPC is highlighted through its fully distributed synthesis of the controller ingredients, the distributed PRS computation and the reduction to a quadratic program, which renders the optimization problem applicable to ADMM. The optimization problem is proven to be recursively feasible, convergent to an average performance bound, while the chance constraints are guaranteed for the closed-loop system. The numerical example reveals that the distributed PRS synthesis comes at the price of increased conservatism, which results in a higher empirical chance constraint satisfaction rate than necessary.

The second contribution is the extension of the direct feedback algorithm to the indirect output-feedback case, where we made use of distributed scenario PRS for constraint tightening. We propose a distributed constraint tightening algorithm that uses scenario optimization results to provide probabilistic guarantees for closed-loop constraint satisfaction, while the resulting MPC optimization problem is provably recursively feasible. In a numerical example, we illustrate the reduction in conservatism with respect to the direct feedback DSMPC case.

4.5 Proofs

Proof of Corollary 4.1

Consider the covariance matrices $\Sigma^{\tilde{\mathrm{x}}}$ and Σ^{e} and the block-diagonal covariance matrix (4.11), such that $\Sigma^{\tilde{\mathrm{x}}}(t) \preceq \hat{\Sigma}^{\tilde{\mathrm{x}}}(t)$ and $\Sigma^{\mathrm{e}}(t) \preceq \hat{\Sigma}^{\mathrm{e}}(t)$. From (4.12a) - (4.12b) we have

$$\Sigma^{\tilde{\mathrm{x}}}(t+1) = \Upsilon_1 + \Sigma^{\mathrm{w}} - L\Sigma^{\mathrm{d}}L^\top \tag{4.31}$$

$$\Sigma^{e}(t+1) = \Upsilon_2 + \Upsilon_3 + L\Sigma^{\mathrm{d}}L, \tag{4.32}$$

where, for simplicity, we denoted the coupled terms as

$$\begin{aligned}
\Upsilon_1 &\coloneqq A_{\mathrm{L}}\Sigma^{\tilde{\mathrm{x}}}(t)A_{\mathrm{L}}^\top \\
\Upsilon_2 &\coloneqq A_{\mathrm{K}}\Sigma^{e}(t)A_{\mathrm{K}}^\top \\
\Upsilon_3 &\coloneqq LC\Sigma^{\tilde{\mathrm{x}}}(t)C^\top L^\top.
\end{aligned}$$

Note that $\Sigma^{\tilde{\mathrm{x}}}(t)$ and $\Sigma^{e}(t)$ are positive semidefinite by design, which renders also $\Upsilon_1, \Upsilon_2, \Upsilon_3$ positive semidefinite. This allows us to follow the proof of [56, Lemma 1] to find that

$$\Upsilon_1 \preceq A_{\mathrm{L}}\hat{\Sigma}^{\tilde{\mathrm{x}}}(t)A_{\mathrm{L}}^{\top} \preceq \begin{bmatrix} \tilde{A}_{\mathcal{N}_1,\mathrm{L}}\hat{\Sigma}_1^{\tilde{x}}(t)\tilde{A}_{\mathcal{N}_1,\mathrm{L}}^{\top} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \tilde{A}_{\mathcal{N}_M,\mathrm{L}}\hat{\Sigma}_M^{\tilde{\mathrm{x}}}(t)\tilde{A}_{\mathcal{N}_M,\mathrm{L}}^{\top} \end{bmatrix}$$

$$\Upsilon_2 \preceq A_{\mathrm{K}}\hat{\Sigma}^{\tilde{\mathrm{x}}}(t)A_{\mathrm{K}}^{\top} \preceq \begin{bmatrix} \tilde{A}_{\mathcal{N}_1,\mathrm{K}}\hat{\Sigma}_1^{\mathrm{e}}(t)\tilde{A}_{\mathcal{N}_1,\mathrm{K}}^{\top} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \tilde{A}_{\mathcal{N}_M,\mathrm{K}}\hat{\Sigma}_M^{\mathrm{e}}(t)\tilde{A}_{\mathcal{N}_M,\mathrm{K}}^{\top} \end{bmatrix}$$

$$\Upsilon_3 \preceq LC\hat{\Sigma}^{\tilde{\mathrm{x}}}(t)C^{\top}L^{\top} \preceq \begin{bmatrix} L_1\tilde{C}_{\mathcal{N}_1}\hat{\Sigma}_1^{\tilde{x}}(t)\tilde{C}_{\mathcal{N}_1}^{\top}L_1^{\top} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & L_M\tilde{C}_{\mathcal{N}_M}\hat{\Sigma}_M^{\tilde{\mathrm{x}}}(t)\tilde{C}_{\mathcal{N}_M}^{\top}L_M^{\top} \end{bmatrix},$$

where $\tilde{A}_{\mathcal{N}_i,\mathrm{L}} = \sqrt{|\mathcal{N}_i|}A_{\mathcal{N}_i,\mathrm{L}}$, $\tilde{A}_{\mathcal{N}_i,\mathrm{K}} = \sqrt{|\mathcal{N}_i|}A_{\mathcal{N}_i,\mathrm{K}}$, $\tilde{C}_{\mathcal{N}_i} = \sqrt{|\mathcal{N}_i|}C_{\mathcal{N}_i}$ and $|\mathcal{N}_i|$. The proof concludes by recalling that the terms Σ^{w} and $L\Sigma^{\mathrm{d}}L^{\top}$ in (4.31)-(4.32) are block-diagonal by design. □

Proof of Theorem 4.1

The proof consists of four parts. First, we show recursive feasibility and predictive satisfaction of chance constraints, followed by closed-loop satisfaction of chance constraints and convergence, while the last part deals with the asymptotic average cost bound. Due to the assumption of exact feasibility, we can use the global vectors during the proof.

Part 1: Recursive feasibility Suppose that at time k a feasible solution to Problem 4.2.1 exists. Then, at time $k+1$, we need to consider the possibly suboptimal solution due to mode 2, for which we define the candidate solutions

$$\tilde{z}(t|k+1) = [z^*(1|k), ..., z^*(N|k), A_{\mathrm{K}}z^*(N|k)]$$
$$\tilde{v}(t|k+1) = [v^*(1|k), ..., v^*(N-1|k), Kz^*(N|k)].$$

From feasibility at time $k+1$ follows that $(\tilde{z}(t|k+1), \tilde{v}(t|k+1)) \in (\mathbb{Z} \times \mathbb{V})$ for $t = 0, ..., N-2$. For $t = N-1$ we have that $\tilde{z}(N-1|k+1) \in \mathbb{Z}_{\mathrm{f}}$. Thus, by Assumption 4.2, and in particular from the invariance property (4.18c), recursive feasibility follows. Predictive chance-constraint satisfaction is then a direct consequence, since for all $z \in \mathbb{Z}_{\mathrm{f}}$ the terminal constraints (4.18b) are satisfied.

Part 2: Closed-loop chance constraint satisfaction For brevity, we show the closed-loop guarantees only for the state constraints. Consider the combined error $\delta x = \tilde{x} + e$ with

$\delta x(0|0) = \delta x(0) = 0$ and assume that $\mathcal{R}$ is a convex symmetric PRS. Now, at time $k+1$, we condition the probability on feasibility of Problem 4.2.1 in Mode 1 or 2

$$\begin{aligned}&\mathbb{P}(\delta x(k+1) \in \mathcal{R})\\ =&\mathbb{P}(\delta x(k+1) \in \mathcal{R}|M^1)\mathbb{P}(M^1) + \mathbb{P}(\delta x(k+1) \in \mathcal{R}|M^2)\mathbb{P}(M^2).\end{aligned} \tag{4.33}$$

In Mode 2 we have $z(0|k+1) = z(1|k)$, so that

$$\mathbb{P}(\delta x(k+1) \in \mathcal{R}|M^2) = \mathbb{P}(e(1|k) + \tilde{x}(1|k) \in \mathcal{R}). \tag{4.34}$$

In Mode 1 we have $z(0|k+1) = \hat{x}(k+1)$ and thus $e(k+1) = 0$, such that

$$\begin{aligned}\mathbb{P}(\delta x(k+1) \in \mathcal{R}|M^1) &= \mathbb{P}(\tilde{x}(k+1) \in \mathcal{R})\\ \geq&\mathbb{P}(\tilde{x}(1|k) \in \mathcal{R}) \geq \mathbb{P}(e(1|k) + \tilde{x}(1|k) \in \mathcal{R}^X),\end{aligned}$$

where the first inequality follows from central convex unimodality and [77, Thm. 3]. The second inequality is due to [2, Thm. 1]. Substituting the latter inequality and (4.34) into (4.33) yields

$$\begin{aligned}&\mathbb{P}(\delta x(k+1) \in \mathcal{R})\\ \geq&\mathbb{P}(e(1|k) + \tilde{x}(k+1) \in \mathcal{R})\mathbb{P}(M^1) + \mathbb{P}(e(1|k) + \tilde{x}(k+1) \in \mathcal{R})\mathbb{P}(M^2)\\ =&\mathbb{P}(\delta x(1|k) \in \mathcal{R}).\end{aligned}$$

Closed-loop chance constraint satisfaction is then a direct consequence of predictive chance constraint satisfaction.

Part 3: Optimal cost decrease The convergence proof follows the same structure as in the proof of Theorem 3.1. Let $V(z(k)) = \|z^*(N|k)\|_P^2 + \sum_{t=0}^{N-1} \|z^*(t|k)\|_Q^2 + \|v^*(t|k)\|_R^2$ be the value function of Problem 4.2.1. We condition the expected cost at time $k+1$ on feasibility of Problem 4.2.1 in Mode 1 or Mode 2

$$\mathbb{E}(V(z(k+1)) = \mathbb{E}(V(z(k+1))|M^2)\mathbb{P}(M^2) + \mathbb{E}(V(z(k+1))|M^1)\mathbb{P}(M^1). \tag{4.35}$$

By optimality, the first term directly satisfies

$$\mathbb{E}(V(z(k+1))|M^2) \leq J(\tilde{z}(\cdot|k+1), \tilde{v}(\cdot|k+1)), \tag{4.36}$$

where $\tilde{z}(\cdot|k+1), \tilde{v}(\cdot|k+1)$ are the shifted state and control sequences. Next, we use the fact that the value function $V(x)$ of a nominal MPC problem is piece-wise quadratic in x [11], which, together with Assumption 4.3 implies the existence of a Lipschitz constant β, such that

$$V(z + \delta x) \leq V(z) + \beta\|\delta x\|_2. \tag{4.37}$$

This allows us to evaluate the expected value in Mode 1 as follows

$$\begin{aligned}&\mathbb{E}(V(\hat{x}(k+1))|M^1) = \mathbb{E}(V(z(0|k+1))|M^1)\\ &\overset{(4.37)}{\leq} J(\tilde{z}(\cdot|k+1),\tilde{v}(\cdot|k+1)) + \beta\mathbb{E}\big(\|x(k+1)-z(1|k)\|_2|M^1\big).\end{aligned}$$

Now we can add $\beta\mathbb{E}(\|x(k+1)-z(1|k)\|_2|M^2)$ to (4.36) and substitute both inequalities into (4.35) resulting in

$$\mathbb{E}(V(z(k+1)) \leq J(\tilde{z}(\cdot|k+1),\tilde{v}(\cdot|k+1)) + \beta\mathbb{E}\big(\|x(k+1)-z(1|k)\|_2\big).$$

The latter term can be further evaluated by considering the decomposition $x(k+1)-z(1|k) = \tilde{x}(1|k)+e(1|k) = [I \quad I]\,\xi(1|k)$, so that

$$\begin{aligned}&\beta\mathbb{E}\big(\|x(k+1)-z(1|k)\|_2\big) = \beta\mathbb{E}\big(\|[I \quad I]\,\xi(1|k)\|_2\big)\\ &\leq \sqrt{2}\beta\mathbb{E}\big(\|\,\xi(1|k)\|_2\big)\\ &\overset{(4.9)}{=} \sqrt{2}\beta\mathbb{E}\Big(\|\Psi\xi(0|k)\|_2 + \|\Gamma\omega(0|k)\|_2\Big)\\ &\leq \underbrace{\frac{\sqrt{2}\beta}{\sqrt{\lambda_{\min}(P_\Sigma)}}}_{\gamma}\Big(\|\Psi\xi(0|k)\|_{P_\Sigma} + \mathbb{E}\big(\|\Gamma\omega(0|k)\|_{P_\Sigma}\big)\Big),\end{aligned} \tag{4.38}$$

where the first inequality is due to the triangle inequality together with (4.9). The second inequality uses $\sqrt{\lambda_{\min}(P_\Sigma)}\|\xi\|_2 \leq \|\xi\|_{P_\Sigma}$, where $P_\Sigma \succ 0$ solves the Lyapunov inequality in the theorem statement for some $\kappa \in \mathbb{R}_{>0}$, so that

$$\|\Psi\xi(0|k)\|_{P_\Sigma} \leq \|\xi(0|k)\|_{P_\Sigma} - \kappa\|\xi(0|k)\|_{P_\Sigma}.$$

Therefore, (4.38) can be bounded as follows

$$\beta\mathbb{E}\big(\|x(k+1)-z(1|k)\|_2\big) \leq \gamma\Big((1-\kappa)\|\xi(0|k)\|_{P_\Sigma} + \mathbb{E}\big(\|\Gamma\omega(0|k)\|_{P_\Sigma}\big)\Big).$$

By combining the latter inequality with the nominal MPC cost decrease due to the terminal controller (4.18a), we obtain

$$\mathbb{E}\big(V\big(z(k+1)\big)\big) - V\big(z(k)\big) \leq -\|z(k)\|_Q^2 - \|v(k)\|_R^2 - \gamma\kappa\|\xi(k)\|_{P_\Sigma} + \gamma\mathbb{E}(\|\Gamma\omega(k)\|_{P_\Sigma},$$

where $z(k) = z(0|k)$ and $\Omega = \operatorname{blkdiag}(\Sigma^{\mathrm{w}}, \Sigma^{\mathrm{d}})$.

Part 4: Asymptotic average cost bound Using standard arguments from stochastic control, we obtain

$$\begin{aligned}0 &\leq \lim_{T\to\infty}\frac{1}{T}\mathbb{E}\big(V\big(z(k)\big)\big) - V\big(z(0)\big)\\ &\leq \lim_{T\to\infty}\sum_{k=0}^{T-1}\mathbb{E}\Big(-\|z(k)\|_Q^2 - \|v(k)\|_R^2 - \gamma\kappa\|\xi(k)\|_{P_\Sigma} + \gamma\mathbb{E}(\|\Gamma\omega(k)\|_{P_\Sigma}\Big)\\ &\leq \lim_{T\to\infty}\sum_{k=0}^{T-1}\gamma\mathbb{E}(\|\Gamma\omega(k)\|_{P_\Sigma} = \gamma\sqrt{\operatorname{tr}(\Gamma^\top P_\Sigma\Gamma\Omega)} = c,\end{aligned}$$

which concludes the proof. □

Proof of Proposition 4.3.2

Recursive feasibility for $k = 0, \dots, N_{\mathrm{T}} - N$ follows from standard arguments by showing feasibility of the shifted candidate sequence. Let $(z^*(\cdot|k), v^*(\cdot|k))$ be the optimal solution to Problem 4.3.1 at time k. Applying (4.4) to (4.2) yields a new state estimate $\tilde{x}(k+1)$ and nominal state $z(k+1) = z^*(1|k)$, for which we consider the candidate sequence $v(t|k+1) = v^*(t+1|k)$ for $t = 0, \dots, N-2$ appended with the terminal controller $v(N-1|k) = Kz^*(N|k)$. Since $v(t|k+1) \in \mathbb{V}^{\mathrm{s}}(k+1+t)$ for $t = 0, \dots, N-2$ and $v(N-1|k) \in \mathbb{V}^{\mathrm{s}}_{\mathrm{f}} \subseteq \mathbb{V}^{\mathrm{s}}(k+N)$ by Assumption 4.7, we have that $v(\cdot|k)$ satisfies the input constraints (4.22f) at time $k+1$. Similarly, the state constraints $z(t|k+1) = z^*(t+1|k) \in \mathbb{Z}^{\mathrm{s}}(k+1+t)$ are verified for all $t = 0, \dots, N-1$, while $z(N|k+1) = (A+BK)z^*(N|k) \in \mathbb{Z}^{\mathrm{s}}_{\mathrm{f}} \subseteq \mathbb{Z}^{\mathrm{s}}(k+N+1)$ by Assumption 4.7. This verifies the state constraints (4.22f) and terminal constraint (4.22g) at time $k+1$.

Constraint satisfaction: Let $\mathcal{R}^s_{\mathcal{N}_i}(k)$ and $\mathcal{R}^s_{\mathrm{u},i}(k)$ be state and input PRS for all $i \in \mathcal{M}$ designed along Corollary 4.2, then with no less than $1-\beta$ probability, the errors $\delta x_{\mathcal{N}_i}$ and $e_{\mathrm{u},i}$ lie within the PRS of probability level p_{x} and p_{u}, respectively. Given that $(z^*(0|k), v^*(0|k)) \in \mathbb{Z}^{\mathrm{s}}(k) \times \mathbb{V}^{\mathrm{s}}(k)$ due to recursive feasibility and $\mathbb{Z}^{\mathrm{s}}(k) = \Pi_{i\in\mathcal{M}}(\mathbb{X}_{\mathcal{N}_i} \ominus \mathcal{R}^s_{\mathcal{N}_i}(k))$ and $\mathbb{V}^{\mathrm{s}}(k) = \Pi_{i\in\mathcal{M}}(\mathbb{U}_i \ominus \mathcal{R}^s_{\mathrm{u},i}(k))$, we therefore have that with no less than $1-\beta$ probability the chance constraints (4.3) are verified in closed-loop conditioned on $x(0)$, i.e., since $\delta x(0) = 0$ by design. □

5 Regulation problem with multiplicative noise

In the previous chapters, we studied the tracking and regulation problem for distributed systems under additive stochastic uncertainty/noise. In the following, we extend the idea of dual consensus ADMM-based DSMPC to handle distributed systems with multiplicative noise, which is attracting increasing interest in the control community due to its ability to approximate complex dynamical systems [44, 68]. Since multiplicative noise affects both the predicted states and the inputs, we cannot resort to an offline constraint tightening approach as in Chapters 3 and 4, but instead consider a probabilistic framework with online constraint tightening as proposed by [59]. This increases the computational overhead of the MPC optimization problem, since in addition to the nominal states, the state variance must also be propagated in the form of LMIs, which ultimately renders the optimization problem a distributed semidefinite program (SDP).

In Section 5.2, we introduce the controller ingredients in a centralized fashion, where the emphasis lies on a distributed structure. In Section 5.3, we propose a distributed synthesis method for the distributed terminal cost function and the distributed terminal set, while Section 5.4 introduces a distributed optimization-based DSMPC. This chapter concludes with Section 5.5, where a numerical example is carried out. This chapter is based on the publication [115] [1].

Related work In [46], an output-feedback DSMPC approach for bounded additive and parametric uncertainties is proposed, where the algorithm is non-iterative and sequentially updating. Furthermore, the coupling is only imposed for the individual chance constraints. The same authors presented a non-iterative sequential state-feedback DSMPC in [48], where the coupling again affects only the individual chance constraints. In [148], a scenario-based parallel updating DSMPC for dynamically coupled systems is proposed, while each subsystem respects local joint chance constraints.

[1] C. Mark and S. Liu. "A stochastic MPC scheme for distributed systems with multiplicative uncertainty". In: *Automatica* 140 (2022), p. 110208 ©2022 Elsevier Ltd.

5.1 Problem description

We consider a network of M linear time-invariant systems, where each subsystem $i \in \mathcal{M}$ has a state vector $x_i \in \mathbb{R}^{n_i}$ and input vector $u_i \in \mathbb{R}^{m_i}$. We consider distributed systems on a graph (Def. 3.1) with local dynamics for each subsystem $i \in \mathcal{M}$ defined as

$$x_i(k+1) = A_{\mathcal{N}_i} x_{\mathcal{N}_i}(k) + B_i u_i(k) + [C_{\mathcal{N}_i} x_{\mathcal{N}_i}(k) + D_i u_i(k)] w_i(k), \tag{5.1}$$

where $A_{\mathcal{N}_i} \in \mathbb{R}^{n_i \times n_{\mathcal{N}_i}}$, $B_i \in \mathbb{R}^{n_i \times m_i}$, $C_{\mathcal{N}_i} \in \mathbb{R}^{n_i \times n_{\mathcal{N}_i}}$ and $D_i \in \mathbb{R}^{n_i \times m_i}$.

The random variable $w_i(k) \in \mathbb{R}$ is a zero mean white noise with unit variance and unbounded support, where the influence of the multiplicative noise on the nominal system, defined via $A_{\mathcal{N}_i}$ and B_i, can be weighted with the perturbation matrices $C_{\mathcal{N}_i}$ and D_i.

Assumption 5.1 (Uncorrelated disturbances)**.** $\mathbb{E}(w_i(k) w_j(t)) = 0$ *for all* $t, k \in \mathbb{N}$ *and for all* $i \neq j$.

We impose individual chance constraints for the local states and inputs

$$\mathbb{P}(H^{\mathrm{x}}_{i,r} x_i(k) \leq 1) \geq p^{\mathrm{x}}_{i,r} \quad \forall r \in \{1, \dots, n_{i,\mathrm{r}}\} \tag{5.2a}$$
$$\mathbb{P}(H^{\mathrm{u}}_{i,s} u_i(k) \leq 1) \geq p^{\mathrm{u}}_{i,s} \quad \forall s \in \{1, \dots, n_{i,\mathrm{s}}\}, \tag{5.2b}$$

where $H^{\mathrm{x}}_{i,r} \in \mathbb{R}^{n_{i,\mathrm{r}} \times n_{i,\mathrm{x}}}$ and $H^{\mathrm{u}}_{i,s} \in \mathbb{R}^{n_{i,\mathrm{s}} \times n_{i,\mathrm{u}}}$, $p^{\mathrm{x}}_{i,r}$ and $p^{\mathrm{u}}_{i,s}$ are the probability levels of constraint satisfaction for the $n_{i,\mathrm{r}}$ state and $n_{i,\mathrm{s}}$ input half-space constraints. The global dynamics are obtained by stacking the local inputs and states as $u = \mathrm{col}_{i \in \mathcal{M}}(u_i) \in \mathbb{R}^m$ and $x = \mathrm{col}_{i \in \mathcal{M}}(x_i) \in \mathbb{R}^n$, i.e.,

$$x(k+1) = Ax(k) + Bu(k) + [Cx(k) + Du(k)] w(k). \tag{5.3}$$

To simplify the exposition throughout this chapter, we make the following assumption.

Assumption 5.2 (Stabilizability)**.** *There exists a structured linear feedback controller*

$$u \coloneqq Kx = \mathrm{col}_{i \in \mathcal{M}}(K_{\mathcal{N}_i} x_{\mathcal{N}_i}),$$

where $K \in \mathbb{R}^{m \times n}$ *and* $K_{\mathcal{N}_i} \in \mathbb{R}^{m_i \times n_{\mathcal{N}_i}}$ $\forall i \in \mathcal{M}$*, such that the global system* (5.3) *is asymptotically stable in the mean-square sense [53].*

Remark 5.1.1. *Asymptotic stability in the mean-square sense implies that* $\mathbb{E}(x(k)) \to 0$ *and* $var(x(k)) \to 0$ *as* $k \to \infty$*. In Section 5.3.1, we propose an LMI-based design technique to synthesized a mean-square stabilizing controller that satisfies Assumption 5.2.*

5.2 Controller design

In the following, we present the main ingredients of the proposed tractable SMPC scheme. We begin by separating the stochastic dynamics (5.1) into its mean and variance dynamics, while the control inputs $u_i(k)$ are restricted to be affine error feedback policies that stabilize the deviation between the mean and true state. Afterwards, the chance constraints (5.2) are analytically approximated via Cantelli's inequality based on the predicted mean and covariance. Finally, we introduce the terminal ingredients in a centralized fashion, while Section 5.3 is devoted to a corresponding distributed reformulation.

5.2.1 Mean-variance dynamics

We define the predicted state mean conditioned on $x(k)$ as $z(t|k) = \mathbb{E}(x(t|k) \mid x(k))$ and consider a distributed error feedback controller of the form

$$u_i(t|k) = v_i(t|k) + K_{\mathcal{N}_i}(x_{\mathcal{N}_i}(t|k) - z_{\mathcal{N}_i}(t|k)) \quad \forall i \in \mathcal{M}, \tag{5.4}$$

where $K_{\mathcal{N}_i}$ is a structured feedback gain according to Assumption 5.2. The mean states and inputs $(z_{\mathcal{N}_i}(\cdot|k), v_i(\cdot|k))$ are obtained from an MPC optimization problem solved at time k. In view of Assumption 5.1 and $w(t|k) \overset{d}{=} w(t+k)$, it can easily be verified that the state mean $z_i(t|k)$ evolves according to the dynamics

$$z_i(t+1|k) = A_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + B_i v_i(t|k). \tag{5.5}$$

Next, we define the error variance conditioned on $x(k)$ as

$$\Sigma_i(t|k) = \operatorname{var}(x_i(t|k) - z_i(t|k) \mid x(k)) \quad \forall t \in \mathbb{N},$$

while its time evolution can be described with the variance dynamics, cf. [59, 140],

$$\begin{aligned}\Sigma_i(t+1|k) = {} & \left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right]\left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right]^\top \\ & + C_{\mathcal{N}_i,\mathrm{K}} \Sigma_{\mathcal{N}_i}(t|k) C_{\mathcal{N}_i,K}^\top + A_{\mathcal{N}_i,\mathrm{K}} \Sigma_{\mathcal{N}_i}(t|k) A_{\mathcal{N}_i,\mathrm{K}}^\top,\end{aligned} \tag{5.6}$$

where $\Sigma_{\mathcal{N}_i}(t|k) = \operatorname{var}(x_{\mathcal{N}_i}(t|k) - z_{\mathcal{N}_i}(t|k) \mid x(k))$, $A_{\mathcal{N}_i,\mathrm{K}} = A_{\mathcal{N}_i} + B_i K_{\mathcal{N}_i}$ and $C_{\mathcal{N}_i,\mathrm{K}} = C_{\mathcal{N}_i} + D_i K_{\mathcal{N}_i}$.

Remark 5.1. *Similar to the previous two chapters, we again face the problem of implementing* (5.6) *in a distributed setting, i.e., the left-hand side updates the local covariance matrix, while the right hand side requires the neighborhood covariance matrix. Thus, we enforce once again a distributed structure by introducing a positive semidefinite block-diagonal matrix* $\hat{\Sigma}(t|k) \in \mathbb{R}^{n \times n}$ *in such a way that*

$$\Sigma(t|k) \preceq \hat{\Sigma}(t|k) = \begin{bmatrix} \hat{\Sigma}_1(t|k) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \hat{\Sigma}_M(t|k) \end{bmatrix}, \tag{5.7}$$

while the local block-diagonal variance dynamics are governed by

$$\begin{aligned}\hat{\Sigma}_i(t+1|k) = &\left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right]\left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right]^\top \\ &+ \tilde{C}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(t|k) \tilde{C}_{\mathcal{N}_i,K}^\top + \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(t|k) \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^\top, \end{aligned} \tag{5.8}$$

where $\tilde{A}_{\mathcal{N}_i,\mathrm{K}} = \sqrt{|\mathcal{N}_i|} A_{\mathcal{N}_i,\mathrm{K}}$ *and* $\tilde{C}_{\mathcal{N}_i,\mathrm{K}} = \sqrt{|\mathcal{N}_i|} C_{\mathcal{N}_i,\mathrm{K}}$. *The block-diagonal neighborhood covariance matrix is given by* $\hat{\Sigma}_{\mathcal{N}_i}(t|k) = \mathrm{blkdiag}_{j\in\mathcal{N}_i}(\hat{\Sigma}_j(t|k))$.

Corollary 5.1. *Let* $\hat{\Sigma}(t+1|k)$ *be the global block-diagonal matrix consisting of* $\hat{\Sigma}_i(t+1|k)$ *for all* $i \in \mathcal{M}$ *given by* (5.8). *If* $\hat{\Sigma}(t|k) \succeq \Sigma(t|k)$, *then it holds that* $\hat{\Sigma}(t+1|k) \succeq \Sigma(t+1|k)$.

Proof. The proof can be constructed similar to Corollary 4.1 by block-diagonally upper bounding the two coupled terms involving $\tilde{A}_{\mathcal{N}_i,\mathrm{K}}$ and $\tilde{C}_{\mathcal{N}_i,\mathrm{K}}$ in the last line of (5.8). □

Covariance propagation via LMIs In the following, the inequality version of (5.8) is converted into a nonlinear matrix inequality assuming that $\hat{\Sigma}_{\mathcal{N}_i}(t|k)$ is positive definite

$$\begin{aligned}\hat{\Sigma}_i(t+1|k) &- \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(t|k) \hat{\Sigma}_{\mathcal{N}_i}^{-1}(t|k) \hat{\Sigma}_{\mathcal{N}_i}(t|k) \\ &- \tilde{C}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(t|k) \hat{\Sigma}_{\mathcal{N}_i}^{-1}(t|k) \hat{\Sigma}_{\mathcal{N}_i}(t|k) \tilde{C}_{\mathcal{N}_i,K}^\top \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^\top \\ &- \left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right] I^{-1} \left[C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k)\right]^\top \succeq 0. \end{aligned}$$

It remains to apply the Schur complement, which leads to an LMI for each subsystem $i \in \mathcal{M}$

$$\begin{bmatrix} \hat{\Sigma}_i(t+1|k) & \begin{bmatrix} \star^2 & \star & \star \end{bmatrix} \\ \begin{bmatrix} (\tilde{A}_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(t|k) + \tilde{B}_i U_{\mathcal{N}_i})^\top \\ (\tilde{C}_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(t|k) + \tilde{D}_i U_{\mathcal{N}_i})^\top \\ (C_{\mathcal{N}_i} z_{\mathcal{N}_i}(t|k) + D_i v_i(t|k))^\top \end{bmatrix} & \begin{bmatrix} \hat{\Sigma}_{\mathcal{N}_i}(t|k) & 0 & 0 \\ 0 & \hat{\Sigma}_{\mathcal{N}_i}(t|k) & 0 \\ 0 & 0 & I \end{bmatrix} \end{bmatrix} \succeq 0, \tag{5.9}$$

where we expanded the terms $\tilde{A}_{\mathcal{N}_i,\mathrm{K}}$ and $\tilde{C}_{\mathcal{N}_i,\mathrm{K}}$ to define a new variable $U_{\mathcal{N}_i} = K_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(t|k)$.

Note that the mean dynamics (5.5) and covariance dynamics (5.9) involve only local variables $z_{\mathcal{N}_i}$ and $\hat{\Sigma}_{\mathcal{N}_i}$, i.e., variables that are accessible by subsystem i through information exchange with neighbors $j \in \mathcal{N}_i$. This is necessary for a fully distributed implementation of the MPC controller, where we aim to solve the distributed SDP via distributed optimization, cf. Section 5.4. In the following sections, we reformulate the chance constraints and introduce a cost function using only the local variables, which ensures that the distributed structure of the problem is preserved.

[2] The symbol $\star$ denotes the corresponding transposed quantity.

5.2.2 Cantelli chance constraint approximation

The individual chance constraints (5.2) are implemented as probabilistic approximations via Cantelli's inequality. As reported in [59], it can be shown that the chance constraints (5.2) for the predicted states and inputs $x(t|k), u(t|k)$ are verified for all $t \in \mathbb{N}$ if we instead impose the following deterministic constraints for all $i \in \mathcal{M}$

$$H^{\mathrm{x}}_{i,r} z_i(t|k) \leq 1 - f(p^{\mathrm{x}}_{i,r}) \sqrt{H^{\mathrm{x}}_{i,r} \hat{\Sigma}_i(t|k) (H^{\mathrm{x}}_{i,r})^\top} \quad \forall r \in \{1, \dots, n_{i,\mathrm{r}}\} \tag{5.10a}$$

$$H^{\mathrm{u}}_{i,s} v_i(t|k) \leq 1 - f(p^{\mathrm{u}}_{i,s}) \sqrt{H^{u}_{i,s} \hat{\Sigma}^{\mathrm{u}}_i(t|k) (H^{\mathrm{u}}_{i,s})^\top} \quad \forall s \in \{1, \dots, n_{i,\mathrm{s}}\}, \tag{5.10b}$$

where $\hat{\Sigma}^{\mathrm{u}}_i(t|k) = K_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(t|k) K^\top_{\mathcal{N}_i}$ and $f(p) = \sqrt{p/(1-p)}$.

Remark 5.2.1. *The function $f(p) = \sqrt{p/(1-p)}$ characterizes a distributionally robust bound on the inverse cumulative distribution function (quantile function) of w at probability level p. However, this bound is quite conservative and is generally tighter when the exact quantile function is known. In this case, we can replace $f(p)$ with the exact quantile function. For instance, if w follows a standard normal distribution, $f(p) = \mathcal{N}^{-1}(p)$ gives the tightest bound, where $\mathcal{N}^{-1}(p)$ denotes the quantile function of the standard normal distribution at probability level p.*

The constraints (5.10) have unfortunately a nonlinear dependency on the covariance matrices. Therefore, as proposed by [60], we perform a simple Taylor linearization of the square root terms, i.e., for the state constraints (5.10a) we write

$$H^{\mathrm{x}}_{i,r} z_i(t|k) \leq 1 - \frac{f(p^{\mathrm{x}}_{i,r}) \sqrt{H^{\mathrm{x}}_{i,r} \hat{\Sigma}_{0,i}(t|k) (H^{\mathrm{x}}_{i,r})^\top}}{2} + \frac{f(p^{\mathrm{x}}_{i,r})}{2\sqrt{H^{\mathrm{x}}_{i,r} \hat{\Sigma}_{0,i}(t|k) (H^{\mathrm{x}}_{i,r})^\top}} H^{\mathrm{x}}_{i,r} \hat{\Sigma}_i(t|k) (H^{\mathrm{x}}_{i,r})^\top.$$

Next, we define the parameter $\epsilon^{\mathrm{x}}_{i,r} \coloneqq f(p^{\mathrm{x}}_{i,r}) \sqrt{H^{\mathrm{x}}_{i,r} \hat{\Sigma}_{0,i}(t|k) (H^{\mathrm{x}}_{i,r})^\top} \in (0, 2)$ and repeat the same linearization procedure for the input constraints (5.10b) resulting in

$$H^{\mathrm{x}}_{i,r} z_i(t|k) \leq 1 - \frac{\epsilon^{\mathrm{x}}_{i,r}}{2} - \frac{f(p^{\mathrm{x}}_{i,r})^2}{2\epsilon^{\mathrm{x}}_{i,r}} H^{\mathrm{x}}_{i,r} \hat{\Sigma}_i(t|k) (H^{\mathrm{x}}_{i,r})^\top \quad \forall r \in \{1, \dots, n_{i,\mathrm{r}}\} \tag{5.11a}$$

$$H^{\mathrm{u}}_{i,s} v_i(t|k) \leq 1 - \frac{\epsilon^{\mathrm{u}}_{i,s}}{2} - \frac{f(p^{\mathrm{u}}_{i,s})^2}{2\epsilon^{\mathrm{u}}_{i,s}} H^{\mathrm{u}}_{i,s} \hat{\Sigma}^{\mathrm{u}}_i(t|k) (H^{\mathrm{u}}_{i,s})^\top \quad \forall s \in \{1, \dots, n_{i,\mathrm{s}}\}. \tag{5.11b}$$

Remark 5.2. *The definition of ϵ^{x}_i depends on the linearization point $\sqrt{H^{\mathrm{x}}_{i,r} \hat{\Sigma}_{0,i}(t|k) (H^{\mathrm{x}}_{i,r})^\top}$. Therefore, the optimal value of ϵ^{x}_i is problem dependent, as the evolution of the covariance matrix depends on the variance dynamics* (5.6)*, which in turn depends on the initial condition $z(0|k)$. Thus, it is not possible to a-priori find a decision rule how to select ϵ^{x}_i, which was similarly reported by [59].*

5.2.3 Objective function

We consider a quadratic expected value cost function over a prediction horizon $N \in \mathbb{N}$

$$J(x(\cdot|k), u(\cdot|k)) = \mathbb{E}\bigg(\|x(N|k)\|_P^2 + \sum_{t=0}^{N-1} \Big(\|x(t|k)\|_Q^2 + \|u(t|k)\|_R^2 \Big) \bigg| x(k) \bigg), \tag{5.12}$$

where $Q \succeq 0$ and $R \succ 0$ are block-diagonal weighting matrices and P satisfies the following assumption.

Assumption 5.3. *There exists a terminal cost $V_\mathrm{f}(x) = \sum_{i\in\mathcal{M}} \|x_i\|_{P_i}^2 = \|x\|_P^2$ with block-diagonal matrix $P \succ 0$ and a distributed terminal controller $u = K_\mathrm{f} x$, such that*

$$(A + BK_\mathrm{f})^\top P (A + BK_\mathrm{f}) + (C + DK_\mathrm{f})^\top P (C + DK_\mathrm{f}) + Q + K_\mathrm{f}^\top R K_\mathrm{f} - P \preceq 0. \tag{5.13}$$

Remark 5.2.2. *The existence of a terminal cost function according to Assumption 5.3 implies that the controller $u = K_\mathrm{f} x$ is mean-square stabilizing for the global system* (5.3)*. This is the same condition that we required for the tube controller gain K by Assumption 5.2. Thus, for simplicity, we set $K_\mathrm{f} = K$.*

Similar to Section 2.3.1, we analytically reformulate the expected quadratic cost function (5.12) by decomposing it into its mean and variance components, such that $J = J_\mathrm{m} + J_\mathrm{v}$. Note that due to block-diagonality of Q, R, P and $\hat{\Sigma}$, the cost function (5.12) is fully separable, i.e., $\hat{J} = \sum_{i\in\mathcal{M}} (J_{\mathrm{m},i} + J_{\mathrm{v},i})$ with

$$\begin{aligned} J_{\mathrm{m},i}(z_i(\cdot|k), v_i(\cdot|k)) &= \|z_i(N|k)\|_{P_i}^2 + \sum_{t=0}^{N-1} \Big(\|z_i(t|k)\|_{Q_i}^2 + \|v_i(t|k)\|_{R_i}^2 \Big) \\ J_{\mathrm{v},i}(\hat{\Sigma}_{\mathcal{N}_i}(\cdot|k)) &= \mathrm{tr}(P_i \hat{\Sigma}_i(N|k)) + \sum_{t=0}^{N-1} \mathrm{tr}(\bar{Q}_i + K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i}) \hat{\Sigma}_{\mathcal{N}_i}(t|k)), \end{aligned}$$

where $\bar{Q}_i = W_i T_i^\top Q_i T_i W_i^\top$ is lifted into $\mathbb{R}^{\mathcal{N}_i \times \mathcal{N}_i}$ via lifting matrices[3] $W_i \in \{0,1\}^{\mathcal{N}_i \times n}$ and $T_i \in \{0,1\}^{n_i \times n}$. Unlike in the previous chapters, the variance cost $J_{\mathrm{v},i}$ depends now on the MPC decision variables (i.e., the predicted covariance matrices $\hat{\Sigma}_{\mathcal{N}_i}$) and thus cannot be neglected in the receding horizon implementation.

5.2.4 Terminal constraints

To ensure recursive feasibility of the MPC optimization problem and convergence of the cost, similar to Farina and Scattolini [59], we impose terminal constraints on both the predicted

[3]The lifting matrices are defined in such a way, that $x_{\mathcal{N}_i} = W_i x$ and $x_i = T_i x$

mean and the predicted variance at the end of the prediction horizon, i.e.,

$$z(N|k) \in \mathbb{Z}_\mathrm{f} \coloneqq \{z \in \mathbb{R}^n | z^\top P z \leq \alpha\} \tag{5.14}$$

$$\hat{\Sigma}(N|k) \preceq \hat{\Sigma}_\mathrm{f}. \tag{5.15}$$

The terminal set $\mathbb{Z}_\mathrm{f}$ is defined as an α sublevel set of the terminal cost function $V_\mathrm{f}(z)$ from Assumption 5.3, where P denotes the block-diagonal weighting matrix. In conjunction with the terminal controller $v_\mathrm{f} = Kz$, it can be shown that the terminal set $\mathbb{Z}_\mathrm{f}$ is positively invariant for the global nominal system $z^+ = Az + Bv_\mathrm{f}$, i.e.,

$$(A + BK)z \in \mathbb{Z}_\mathrm{f} \quad \forall z \in \mathbb{Z}_\mathrm{f}.$$

In view of (5.15), we need to compute a block-diagonal terminal covariance matrix $\hat{\Sigma}_\mathrm{f}$ that verifies the stationary condition

$$(A+BK)\hat{\Sigma}_\mathrm{f}(A+BK)^\top + (C+DK)\hat{\Sigma}_\mathrm{f}(C+DK)^\top + (C+DK)\Psi(C+DK)^\top \preceq \hat{\Sigma}_\mathrm{f}, \tag{5.16}$$

where Ψ is an arbitrary state covariance matrix that is defined in such a way that

$$\Psi \succ \psi I \succ zz^\top \succeq 0 \quad \forall z \in \mathbb{Z}_\mathrm{f} \tag{5.17}$$

for some non-negative $\psi \in \mathbb{R}_{>0}$. In addition, the chance constraints (5.11) must also be fulfilled for all $z \in \mathbb{Z}_\mathrm{f}$ under application of the terminal controller, i.e.,

$$H^\mathrm{x}_{i,r} z_i \leq 1 - \frac{\epsilon^\mathrm{x}_{i,r}}{2} - \frac{f(p^\mathrm{x}_{i,r})^2}{2\epsilon^\mathrm{x}_{i,r}} H^\mathrm{x}_{i,r} \hat{\Sigma}_{\mathrm{f},i} (H^\mathrm{x}_{i,r})^\top \quad \forall r \in \{1, \dots, n_{i,\mathrm{r}}\} \tag{5.18a}$$

$$H^\mathrm{u}_{i,s} K_{\mathcal{N}_i} z_{\mathcal{N}_i} \leq 1 - \frac{\epsilon^\mathrm{u}_{i,s}}{2} - \frac{f(p^\mathrm{u}_{i,s})^2}{2\epsilon^\mathrm{u}_{i,s}} H^\mathrm{u}_{i,s} \hat{\Sigma}^\mathrm{u}_{\mathrm{f},i} (H^\mathrm{u}_{i,s})^\top \quad \forall s \in \{1, \dots, n_{i,\mathrm{s}}\} \tag{5.18b}$$

for all $i \in \mathcal{M}$, where $\hat{\Sigma}^\mathrm{u}_{\mathrm{f},i} = K_{\mathcal{N}_i} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} K^\top_{\mathcal{N}_i}$.

Remark 5.3. *Note that in view of* (5.16) - (5.17)*, it is always possible to define a sufficiently small set* $\mathbb{Z}_\mathrm{f}$*, such that for all* $z \in \mathbb{Z}_\mathrm{f}$ *the terminal chance constraints* (5.18) *are verified. In fact, the smaller* $\mathbb{Z}_\mathrm{f}$*, the smaller* Ψ *resulting from* (5.17) *and hence the smaller* $\hat{\Sigma}_\mathrm{f}$ *resulting from* (5.16)*. If a zero terminal constraint strategy is adopted, then* $\mathbb{Z}_\mathrm{f} = \{0\}$ *results in a state covariance* $zz^\top = 0$ *in* (5.17)*. In this case,* ψ *can be selected arbitrarily close to zero to obtain a small (but positive definite) terminal covariance matrix through* (5.16)*.*

5.2.5 Central MPC optimization problem

The following MPC optimization problem is solved at every time instant $k \in \mathbb{N}$.

Problem 5.2.3 (Centralized Probabilistic MPC)**.**

$$V(x(k)) = \min_{\mathcal{Z},\mathcal{V},\mathcal{S}} \quad \sum_{i=1}^{M} \Big(J_{\mathrm{m},i}(z_i(\cdot|k), v_i(\cdot|k)) + J_{\mathrm{v},i}(\hat{\Sigma}_{\mathcal{N}_i}(\cdot|k)) \Big) \tag{5.19a}$$

$$\text{s.t.} \quad (5.5), (5.8), (5.11) \quad \forall t \in \{0, ..., N-1\} \quad \forall i \in \mathcal{M} \tag{5.19b}$$

$$z(N|k) \in \mathbb{Z}_{\mathrm{f}} \tag{5.19c}$$

$$\hat{\Sigma}(N|k) \preceq \hat{\Sigma}_{\mathrm{f}} \tag{5.19d}$$

$$(z(0|k), \hat{\Sigma}(0|k)) \in \{(x(k), 0), (z(1|k-1), \hat{\Sigma}(1|k-1))\} \tag{5.19e}$$

where $\mathcal{V} = \{v(0|k), \dots, v(N-1|k)\}$, $\mathcal{Z} = \{z(0|k), \dots, z(N|k)\}$, $\mathcal{S} = \{\hat{\Sigma}(0|k), \dots, \hat{\Sigma}(N|k)\}$ *denote the input, state and covariance sequences, respectively.*

The resulting MPC is categorized as a direct feedback controller, cf. Remark 2.3, due to the feasibility-based initial constraint (5.19e), where we consider the initial conditions $(z(0|k), \hat{\Sigma}(0|k))$ as free decision variables. We define the feedback mode (Mode 1) as $x(0|k) = \mathbb{E}(x(k) \mid x(k))$, which is selected via the condition $(z(0|k), \hat{\Sigma}(0|k)) = (x(k), 0)$ whenever possible. In case of infeasibility of Mode 1, we define the backup strategy (Mode 2) as $x(0|k) = \mathbb{E}(x(k) \mid x(k-1))$, which is enforced by the condition $(z(0|k), \hat{\Sigma}(0|k)) = (z(1|k-1), \hat{\Sigma}(1|k-1))$.

The last challenge in solving Problem 5.2.3 distributedly is the terminal set $\mathbb{Z}_{\mathrm{f}}$, where the main requirement is that $\mathbb{Z}_{\mathrm{f}}$ is decomposable into M subproblems, each of which only involves the variables $z_{\mathcal{N}_i}$. In Section 5.3.4, we propose a structured terminal set as a Cartesian product of local time-varying sets, which satisfy the decomposability property.

5.3 Distributed synthesis

In the following, we address a distributed design procedure of the controller ingredients.

5.3.1 Structured terminal cost and distributed controller

First, we develop a distributed synthesis procedure for the distributed controller and structured terminal cost as required by Assumption 5.3. More specifically, the goal is to find local quadratic functions

$$V_{\mathrm{f},i}(x_i) = x_i^\top P_i x_i \qquad \forall i \in \mathcal{M}$$
$$\gamma_i(x_{\mathcal{N}_i}) = x_{\mathcal{N}_i}^\top \Gamma_{\mathcal{N}_i} x_{\mathcal{N}_i} \quad \forall i \in \mathcal{M},$$

such that the global cost decrease condition (5.13) holds true. Similar to Conte et al. [42], we introduce indefinite relaxation functions $\gamma_i(\cdot)$ to allow the local cost $V_{\mathrm{f},i}(x_i)$ to partially

increase, as long as the expected global cost $V_{\mathrm{f}}(x)$ always decreases. These implications can be translated into the following inequalities

$$\mathbb{E}(V_{\mathrm{f},i}(x_i^+) \mid x_i) - V_{\mathrm{f},i}(x_i) + l(x_{\mathcal{N}_i}, K_{\mathcal{N}_i} x_{\mathcal{N}_i}) - \gamma_i(x_{\mathcal{N}_i}) \leq 0 \quad \forall i \in \mathcal{M} \tag{5.20a}$$

$$\sum_{i=1}^{M} \gamma_i(x_{\mathcal{N}_i}) \leq 0, \tag{5.20b}$$

where $x_i^+ = A_{\mathcal{N}_i,\mathrm{K}} x_{\mathcal{N}_i} + C_{\mathcal{N}_i,\mathrm{K}} x_{\mathcal{N}_i} w_i$ with stage cost $l(x_{\mathcal{N}_i}, K_{\mathcal{N}_i} x_{\mathcal{N}_i}) = x_{\mathcal{N}_i}^\top (\bar{Q}_i + K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i}) x_{\mathcal{N}_i}$. By replacing the expressions for $V_{\mathrm{f},i}(x_i^+)$ and $\gamma_i(x_{\mathcal{N}_i})$ in equations (5.20a) - (5.20b) with their definitions and resolving the conditional expectation, we arrive at a set of nonlinear inequalities involving the states $x_{\mathcal{N}_i}$. Since these inequalities should hold for all $x_{\mathcal{N}_i}$, we obtain a set of nonlinear matrix inequalities

$$A_{\mathcal{N}_i,\mathrm{K}}^\top P_i A_{\mathcal{N}_i,\mathrm{K}} + C_{\mathcal{N}_i,K}^\top P_i C_{\mathcal{N}_i,\mathrm{K}} - \bar{P}_i \preceq -(\bar{Q}_i + K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i}) + \Gamma_{\mathcal{N}_i} \quad \forall i \in \mathcal{M} \tag{5.21a}$$

$$\sum_{i=1}^{M} W_i^\top \Gamma_{\mathcal{N}_i} W_i \preceq 0, \tag{5.21b}$$

where $\bar{P}_i = W_i T_i^\top P_i T_i W_i^\top$ is lifted into $\mathbb{R}^{\mathcal{N}_i \times \mathcal{N}_i}$. Condition (5.21a) is structured by design, i.e., it is fully distributedly solvable, while (5.21b) connects all subsystems with a system-wide coupling constraint. In the following, we provide an LMI approximation of the nonlinear matrix inequalities (5.21).

Corollary 5.2. *Conditions* (5.21a)- (5.21b) *are equivalent to the following set of LMIs*

$$\begin{bmatrix} \bar{E}_i + F_{\mathcal{N}_i} & \begin{bmatrix} \star & \star & \star & \star \end{bmatrix} \\ \begin{bmatrix} A_{\mathcal{N}_i} E_{\mathcal{N}_i} + B_i Y_{\mathcal{N}_i} \\ C_{\mathcal{N}_i} E_{\mathcal{N}_i} + D_i Y_{\mathcal{N}_i} \\ \bar{Q}_i^{1/2} E_{\mathcal{N}_i} \\ R_i^{1/2} Y_{\mathcal{N}_i} \end{bmatrix} & \begin{bmatrix} E_i & 0 & 0 & 0 \\ 0 & E_i & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} \end{bmatrix} \succeq 0 \quad \forall i \in \mathcal{M} \tag{5.22a}$$

$$\sum_{i=1}^{M} W_i^\top F_{\mathcal{N}_i} W_i \preceq 0, \tag{5.22b}$$

where $E_i = P_i^{-1}$, $\bar{E}_i = W_i T_i^\top P_i^{-1} T_i W_i^\top$, $E_{\mathcal{N}_i} = W_i E W_i^\top$, $F_{\mathcal{N}_i} = E_{\mathcal{N}_i} \Gamma_{\mathcal{N}_i} E_{\mathcal{N}_i}$ *and* $Y_{\mathcal{N}_i} = K_{\mathcal{N}_i} E_{\mathcal{N}_i}$.

Proof. The proof follows directly from [42, Lemma 10] by considering the additional term $C_{\mathcal{N}_i,K}^\top P_i C_{\mathcal{N}_i,\mathrm{K}}$ in condition (5.21a). □

Remark 5.3.1. *The matrix* $F_{\mathcal{N}_i}$ *in* (5.22b) *is block-sparse and is the last hindrance for the distributed synthesis procedure. Therefore, Conte et al. [42] propose to use block-diagonal upper bounding matrices* $F_{\mathcal{N}_i} \preceq S_{\mathcal{N}_i}$ *together with a neighbor-to-neighbor coupling constraint* $\sum_{j \in \mathcal{N}_i} T_j W_j^\top S_{\mathcal{N}_j} W_j T_j^\top \preceq 0$ *for all* $i \in \mathcal{M}$ *to replace the system-wide coupling constraint* (5.22b). *This, however, introduces additional conservatism as we essentially upper bound a dense matrix with a block-diagonal matrix, i.e., similar to Remark 5.1.*

5.3.2 Distributed terminal covariance matrix

Next, we develop an LMI-based design procedure for the terminal covariance matrix as required by (5.15). Consider the block-diagonal covariance matrix $\hat{\Sigma}_{\mathrm{f}}$ from (5.15) and let $\hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} = W_i \hat{\Sigma}_{\mathrm{f}} W_i^\top$ be the block-diagonal neighborhood covariance matrix. To derive the distributed terminal covariance matrix, we substitute $\hat{\Sigma}_{\mathrm{f},i}, \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i}$ together with the terminal controller $v_i = K_{\mathcal{N}_i} z_{\mathcal{N}_i}$ into (5.8), resulting in the Lyapunov-like stationary covariance equation

$$\hat{\Sigma}_{\mathrm{f},i} = \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^\top + \tilde{C}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} \tilde{C}_{\mathcal{N}_i,K}^\top + \tilde{C}_{\mathcal{N}_i,\mathrm{K}} \Psi_{\mathcal{N}_i} \tilde{C}_{\mathcal{N}_i,K}^\top \quad \forall i \in \mathcal{M}, \tag{5.23}$$

where $\Psi_{\mathcal{N}_i} \in \mathbb{R}^{\mathcal{N}_i \times \mathcal{N}_i}$ is an arbitrary state covariance matrix. The nonlinear matrix equality (5.23) can be represented as an LMI by transforming the equality in (5.23) into an inequality and defining $\Psi_{\mathcal{N}_i} = \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i}$. For the latter equality to hold, it is necessary that $\hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} \succeq \psi I$ is satisfied, with $\psi = \min\{\psi_1, \dots, \psi_M\}$ additionally verifying the global condition (5.17). Finally, by defining new matrix variables $U_{\mathcal{N}_i} = K_{\mathcal{N}_i} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i}$ and $\bar{\psi}_i$, the inequality version of (5.23) can be cast as the following pair of LMIs for all $i \in \mathcal{M}$ via Schur complements

$$\begin{bmatrix} \hat{\Sigma}_{\mathrm{f},i} & \begin{bmatrix} \star & \star \end{bmatrix} \\ \begin{bmatrix} (\tilde{A}_{\mathcal{N}_i} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} + \tilde{B}_i U_{\mathcal{N}_i})^\top \\ (\tilde{C}_{\mathcal{N}_i} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} + \tilde{D}_i U_{\mathcal{N}_i})^\top \end{bmatrix} & \begin{bmatrix} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} & 0 \\ 0 & \frac{1}{2} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} \end{bmatrix} \end{bmatrix} \succeq 0 \tag{5.24a}$$

$$\begin{bmatrix} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} & I \\ I & \bar{\psi}_i I \end{bmatrix} \succeq 0. \tag{5.24b}$$

5.3.3 A unique terminal controller

It is important to note that both LMIs (5.22) and (5.24) rely on a structured feedback matrix $K_{\mathcal{N}_i}$ through the matrix variables $Y_{\mathcal{N}_i}$ and $U_{\mathcal{N}_i}$. To eliminate this ambiguity, a unique LMI problem must be posed to solve both LMIs simultaneously. A trivial idea would be to introduce an additional uniqueness constraint $U_{\mathcal{N}_i} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i}^{-1} = Y_{\mathcal{N}_i} E_{\mathcal{N}_i}^{-1}$, which unfortunately destroys the convexity of the problem, as already stated by Farina and Scattolini [59]. A simple (but conservative) way of circumventing the non-convexity issue is to set $\hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} = E_{\mathcal{N}_i}$ and $U_{\mathcal{N}_i} = Y_{\mathcal{N}_i}$, which we will consider hereafter.

Proposition 5.1. *Let $\hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} = E_{\mathcal{N}_i}$, $U_{\mathcal{N}_i} = Y_{\mathcal{N}_i}$ $\forall i \in \mathcal{M}$. If the following optimization problem admits a feasible solution*

$$\begin{aligned} \max \quad & \sum_{i=1}^{M} \log(\det(E_i)) \\ \text{s.t.} \quad & (5.22\text{a}), (5.22\text{b}), (5.24\text{a}), (5.24\text{b}) \quad \forall i \in \mathcal{M}, \end{aligned}$$

then the terminal weighting matrices P_i for all $i \in \mathcal{M}$ are unique and the volume the 1-level set of $V_{\mathrm{f}}(x) = \sum_{i \in \mathcal{M}} \|x_i\|_{P_i} = \|x\|_P$ is maximized.

Proof. For $E_{\mathcal{N}_i} = \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i}, U_{\mathcal{N}_i} = Y_{\mathcal{N}_i}, \forall i \in \mathcal{M}$, the LMIs (5.22a), (5.22b), (5.24a), (5.24b) are convex in $E_{\mathcal{N}_i}$ and $Y_{\mathcal{N}_i}$. Therefore, the minimizer is unique, while the objective function $\sum_{i=1}^{M} \log(\det(E_i))$ is convex and maximizes the volume of the 1-level set of $V_\mathrm{f}(x)$ [20]. □

Remark 5.3.2. *Proposition 5.1 yields a separable terminal cost function with weights P_i, terminal controller gains $K_{\mathcal{N}_i}$, relaxation matrices $\Gamma_{\mathcal{N}_i}$ and terminal covariance matrices $\hat{\Sigma}_{\mathrm{f},i} = E_i$ for all $i \in \mathcal{M}$ that satisfy Assumption 5.3. Infeasibility of the optimization problem in Proposition 5.1 implies that there exists no distributed stabilizing terminal controller for system* (5.1)*. In this case, we can set $\hat{\Sigma}_{\mathrm{f},i}, P_i, \Gamma_{\mathcal{N}_i}$ to zero for all $i \in \mathcal{M}$ and resort to a zero terminal constraint strategy. Furthermore, in view of Remark 5.2.2, we then have to compute a structured stabilizing tube controller K for the error system* (5.4)*, e.g., via LMIs* (5.22)*.*

5.3.4 Structured terminal sets

In the following, we propose a structured global terminal set $\hat{\mathbb{Z}}_\mathrm{f}$ that replaces the global terminal set $\mathbb{Z}_\mathrm{f}$ in (5.19c). The main idea relies on the concept of distributed invariance [42], which is based on time-varying terminal sets.

Definition 5.1 (Time-varying terminal sets)**.** *Let $\mathbb{Z}_\mathrm{f}$ be the global terminal set from* (5.14) *and define α, such that for all $z \in \mathbb{Z}_\mathrm{f}$ the constraints* (5.16) - (5.18) *are verified. Define local time-varying terminal sets as*

$$\mathbb{Z}_{\mathrm{f},i}(\alpha_i(k)) \coloneqq \{z_i \in \mathbb{R}^{n_i} | z_i^\top P_i z_i \leq \alpha_i(k)\} \quad \forall i \in \mathcal{M}, \tag{5.25}$$

where $\alpha_i(k)$ is given by the set dynamics

$$\alpha_i(k+1) = \alpha_i(k) + z_{\mathcal{N}_i}^\top(k) \Gamma_{\mathcal{N}_i} z_{\mathcal{N}_i}(k) \quad \forall i \in \mathcal{M}$$

with $\sum_{i=1}^{M} \alpha_i(0) \leq \alpha$ and $\alpha_i(0) \geq 0 \quad \forall i \in \mathcal{M}$.

In order to prove recursive feasibility under time-varying terminal sets, we recall the following two lemmas from [42].

Lemma 5.3.3 (Local invariance [42, Lem. 8])**.** *Let $\mathbb{Z}_{\mathrm{f},i}(\alpha_i(k))$ for all $i \in \mathcal{M}$ be local terminal sets as in Definition 5.1, then*

$$\begin{aligned} &z_i(k) \in \mathbb{Z}_{\mathrm{f},i}(\alpha_i(k)) \implies (A_{\mathcal{N}_i} + B_i K_{\mathcal{N}_i}) z_{\mathcal{N}_i}(k) \in \mathbb{Z}_{\mathrm{f},i}(\alpha_i(k+1)) \quad \forall i \in \mathcal{M} \\ &\alpha_i(k+1) \geq 0 \quad \forall i \in \mathcal{M}. \end{aligned}$$

Lemma 5.3.4 (Global invariance [42, Lem. 9])**.** *Let $\mathbb{Z}_\mathrm{f}$ be a global terminal set and $\mathbb{Z}_{\mathrm{f},i}(\alpha_i(k))$ for all $i \in \mathcal{M}$ as in Definition 5.1, then*

$$\hat{\mathbb{Z}}_\mathrm{f}(\alpha_1(k), \ldots, \alpha_M(k)) \coloneqq \prod_{i=1}^{M} \mathbb{Z}_{\mathrm{f},i}(\alpha_i(k)) \subseteq \mathbb{Z}_\mathrm{f} \quad \forall k \in \mathbb{N}.$$

Next, we compute the scaling factor α (Definition 5.1), such that for all $z \in \hat{\mathbb{Z}}_{\mathrm{f}}$ the terminal constraints (5.17) and (5.18) are satisfied. We propose the following distributed linear program, which is an extension of the optimization problem from [42, Sec. 4.2].

Problem 5.3.5. *(Distributed terminal set)*

$$\alpha = \max_{\hat{\alpha}>0} \quad \hat{\alpha} \tag{5.26a}$$

$$\text{s.t.} \quad \|P_i^{-\frac{1}{2}}(H_{i,r}^{\mathrm{x}})^\top\|^2\hat{\alpha} \leq (\tilde{h}_{i,r}^{\mathrm{x}})^2 \qquad \forall i \in \mathcal{M} \quad \forall r \in \{1,\dots,n_{i,\mathrm{r}}\} \tag{5.26b}$$

$$\|P_{\mathcal{N}_i}^{-\frac{1}{2}}K_{\mathcal{N}_i}^\top(H_{i,s}^{\mathrm{u}})^\top\|^2\hat{\alpha} \leq (\tilde{h}_{i,s}^{\mathrm{u}})^2 \quad \forall i \in \mathcal{M} \quad \forall s \in \{1,\dots,n_{i,\mathrm{s}}\} \tag{5.26c}$$

$$\|P_i^{-1}\|\hat{\alpha} \leq \psi_i \qquad \forall i \in \mathcal{M}, \tag{5.26d}$$

where

$$\tilde{h}_{i,r}^{\mathrm{x}} = 1 - \frac{\epsilon_{i,r}^{\mathrm{x}}}{2} - \frac{f(p_{i,r}^{\mathrm{x}})^2}{2\epsilon_{i,r}^{\mathrm{x}}} H_{i,r}^{\mathrm{x}} P_i^{-1}(H_{i,r}^{\mathrm{x}})^\top > 0$$

$$\tilde{h}_{i,s}^{\mathrm{u}} = 1 - \frac{\epsilon_{i,s}^{\mathrm{u}}}{2} - \frac{f(p_{i,s}^{\mathrm{u}})^2}{2\epsilon_{i,s}^{\mathrm{u}}} H_{i,s}^{\mathrm{u}} K_{\mathcal{N}_i} P_{\mathcal{N}_i}^{-1} K_{\mathcal{N}_i}^\top (H_{i,s}^{\mathrm{u}})^\top > 0$$

denote the right-hand side of (5.18) *with the terminal covariance matrices from Proposition 5.1, i.e.,* $\hat{\Sigma}_{\mathrm{f},i} = P_i^{-1}$ *and* $\hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} = P_{\mathcal{N}_i}^{-1}$.

Lemma 5.3.6. *Let Assumption 5.3 hold. The solution of Problem 5.3.5 defines the largest feasible level set* $\mathbb{Z}_{\mathrm{f}} = \{z \in \mathbb{R}^n | z^\top P z \leq \alpha\}$.

Proof. The proof can be found in Section 5.7. □

Once such a global level set $\mathbb{Z}_{\mathrm{f}}$ with size α is derived, the local terminal sets from Definition 5.1 can be initialized according to $\sum_{i=1}^{M} \alpha_i(0) \leq \alpha$.

Remark 5.4. *Note that* $\tilde{h}_{i,r}^{\mathrm{x}}$ *and* $\tilde{h}_{i,s}^{\mathrm{u}}$ *in Problem 5.3.5 must be positive scalars for all halfspace constraints. If these conditions are violated, we can vary the linearization parameters* $\epsilon_{i,r}^{\mathrm{x}}$ *and* $\epsilon_{i,s}^{\mathrm{u}}$ *or the weighting matrices* Q *and* R*, since they directly influence* P.

5.4 Distributed Optimization for DSMPC

At this point, the central MPC optimization problem 5.2.3 can be written entirely by means of distributed ingredients, that is, we replace the terminal set $\mathbb{Z}_{\mathrm{f}}$ with the structured terminal set from Definition 5.1 and the terminal covariance matrix $\hat{\Sigma}_{\mathrm{f}}$ with the block-diagonal matrix P^{-1}. In what follows, we briefly outline a standard distributed consensus ADMM that is conceptually similar to Chapter 3.

Let Ξ contain all global predictions of the input, state and covariance sequences. Let ξ_i consist of the state and covariance sequence of the neighboring subsystems as predicted by subsystem i, i.e., $z^i_{\mathcal{N}_i}(\cdot|k)$ and $\hat{\Sigma}^i_{\mathcal{N}_i}(\cdot|k)$, and the predicted input $v_i(\cdot|k)$ over the prediction horizon N. To coordinate the local solutions, we implement a consensus constraint $G_i\,\Xi\;=\;\xi_i\quad\forall i\in\mathcal{M}$, for which we formulate the augmented Lagrangian

$$\mathcal{L}_i(\xi_i,\Xi,\lambda_i)=J_i(\xi_i)+\lambda_i^\top(\xi_i-G_i\Xi)+\frac{\rho}{2}\|\xi_i-G_i\Xi\|_2^2\quad\forall i\in\mathcal{M},$$

where λ_i is a Lagrange multiplier and $\rho\in\mathbb{R}_{>0}$ a positive augmentation factor. The augmented Lagrangian depends only on local variables, which implies that the MPC optimization problem 5.2.3 is decomposable into M local optimization problems.

Problem 5.4.1. *(Local MPC optimization problem)*

$$\begin{aligned}
\xi_i^+=\underset{\xi_i}{\operatorname{argmin}}\quad & \mathcal{L}_i(\xi_i,\Xi,\lambda_i) && \text{(5.27a)}\\
\text{s.t.}\quad & (5.5),(5.8),(5.11)\quad\forall t\in\{0,...,N-1\} && \text{(5.27b)}\\
& z_i(N|k)\in\mathbb{Z}_{\mathrm{f},i}(\alpha_i(k)) && \text{(5.27c)}\\
& \Sigma_i(N|k)\le\hat{\Sigma}_{\mathrm{f},i}=P_i^{-1} && \text{(5.27d)}\\
& (z_i(0|k),\Sigma_i(0|k))=(z_{i,0},\Sigma_{i,0}) && \text{(5.27e)}
\end{aligned}$$

for all $i\in\mathcal{M}$, $r=1,\dots,n_{i,\mathrm{r}}$ *and* $s=1,\dots,n_{i,\mathrm{s}}$.

We introduce the following notation: ξ^+_{ji} and λ_{ji} indicate ξ_i^+ and λ_i predicted by subsystem j. For practical reasons, a simple stopping criterion for the ADMM algorithm is implemented

$$\|G_i\Xi-\xi_i\|_\infty\le\epsilon_\mathrm{c},\tag{5.28}$$

where $\epsilon_\mathrm{c}\in\mathbb{R}_{\ge0}$. The basic ADMM steps are given in Algorithm 7, while the online DSMPC steps are summarized in Algorithm 8. Note that Algorithm 7 is executed at every time instant $k\in\mathbb{N}$.

Algorithm 7 Consensus ADMM

1: For each subsystem $i\in\mathcal{M}$ in parallel:
2: Initialize $\lambda_i=0$, $\xi_i=0$ and $(z_{i,0},\Sigma_{i,0})$ according to (Mode 1) or (Mode 2)
3: **repeat**
4: Solve MPC Problem 5.4.1 and obtain ξ_i^+
5: Communicate ξ_i^+ to neighbors $j\in\mathcal{N}_i$
6: Average $\Xi_i^+=\dfrac{1}{|\mathcal{N}_i|}\displaystyle\sum_{j\in\mathcal{N}_i}G_{ji}^\top(\xi_{ji}^+\dfrac{1}{\rho}\lambda_{ji})$
7: Communicate ξ_i^+ to neighbors $j\in\mathcal{N}_i$
8: $\lambda_i^+=\lambda_i+\rho(\xi_i^+-G_i\Xi^+)$
9: **until** (5.28) is satisfied

Algorithm 8 Online DSMPC

1: Measure local states $x_i(k)$ for all $i \in \mathcal{M}$ and share with neighbors
2: Set $(z_{i,0}, \Sigma_{i,0}) = (x_i(k), 0)$ for all $i \in \mathcal{M}$ and solve Problem 5.2.3 via Alg. 7
3: **if** infeasibility is detected **then**
4: Set $(z_{i,0}, \Sigma_{i,0}) = (z_i^i(1|k-1), \hat{\Sigma}_i^i(1|k-1))$ for all $i \in \mathcal{M}$ and solve Problem 5.2.3 via Alg. 7
5: **end if**
6: Each subsystem $i \in \mathcal{M}$ applies the optimal control input

$$u_i(k) = v_i(0|k) + K_{\mathcal{N}_i}(x_{\mathcal{N}_i}(k) - z_{\mathcal{N}_i}^{i*}(0|k))$$

7: Each subsystem $i \in \mathcal{M}$ updates the local terminal set with

$$\alpha_i(k+1) = \alpha_i(k) + (z_{\mathcal{N}_i}^{i*})^\top(N|k)\Gamma_{\mathcal{N}_i} z_{\mathcal{N}_i}^{i*}(N|k)$$

8: $k \leftarrow k+1$ and go to step 1

The following theorem is the main result of this chapter and provides guarantees for recursive feasibility of the MPC optimization problem, predictive chance constraint satisfaction, and point-wise convergence of the states.

Theorem 5.4.2. *If at time $k = 0$ Problem 5.2.3 admits a feasible solution via Algorithm 8, then it is recursively feasible, $\mathbb{E}(||x(k)||_Q^2) \to 0$ as $k \to \infty$ and the chance constraints* (5.2) *are satisfied for all times $k \in \mathbb{N}$.*

Proof. The proof can be found in Section 5.7. □

5.5 Numerical example

We demonstrate our approach on a numerical example of a chain of coupled double integrators, where for subsystem 1 it holds that $\mathcal{N}_1 = \{1, 2\}$, for subsystem M it holds that $\mathcal{N}_M = \{M-1, M\}$ and for subsystems $i \in \mathcal{M}\backslash\{1, M\}$ it holds that $\mathcal{N}_i = \{i-1, i, i+1\}$.

Problem setting In the first experiment we choose $M = 5$ and consider the following dynamics matrices to define system (5.1) for all $i \in \mathcal{M}$

$$A_{ii} = \begin{bmatrix} 1 & 1 \\ 0 & 0.9 \end{bmatrix}, A_{ij} = \begin{bmatrix} 0.1 & 0 \\ 0.1 & 0.1 \end{bmatrix}, C_{ii} = \begin{bmatrix} 0.01 & 0 \\ 0.02 & 0.03 \end{bmatrix}, C_{ij} = \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix} \quad \forall j \neq i,$$

while the input matrices are given by

$$B_i = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, D_i = \begin{bmatrix} 0 \\ 0.001 \end{bmatrix}.$$

For subsystems $i \in \{1, M\}$ we impose a chance constraint $\mathbb{P}([x_i(k)]_2 \geq -1) \geq 0.7$. The multiplicative noise is normally distributed with $w \sim \mathcal{N}(0,1)$, the weighting matrices are set to $Q_i = \mathrm{diag}(5,1), R_i = 0.3$ for all $i \in \mathcal{M}$ and the prediction horizon is $N = 10$. The initial conditions are $x_1(0) = x_M(0) = [3 \quad 0]^\top$ and $x_i(0) = [1 \quad 0]^\top$ for all $i \in \{2,3,4\}$. The constraint linearization parameters are set to $\epsilon^{\mathrm{x}_i} = 0.1$, the probability bound is $f(p) = \mathcal{N}^{-1}(0.7) = 0.5244$ (Remark 5.2.1) and the ADMM augmentation factor is $\rho = 10$.

Results In the following, we carry out $N_{\mathrm{mc}} = 1000$ Monte-Carlo simulations to evaluate the empirical closed-loop chance constraint satisfaction and optimality for different parameterizations. Figure 5.1 shows the density plot of the closed-loop trajectories of subsystem 1. In Table 5.1, we compare for different values of ϵ_{c} the average and maximum number of iterations of Algorithm 8, the average cumulative closed-loop cost

$$\mathrm{av}[J] = N_{\mathrm{mc}}^{-1} \sum_{q=1}^{N_{\mathrm{mc}}} \sum_{k=0}^{13} \|x(k)\|_Q^2 + \|u(k)\|_R^2$$

and the worst-case empirical in-time constraint satisfaction of subsystem 1, i.e.,

$$c_{\mathrm{wc}} = \min_{k \in \{0,\dots,13\}} 1 - N_{\mathrm{mc}}^{-1} c_{\mathrm{v}}(k)$$

with $c_{\mathrm{v}}(k) = \sum_{q=1}^{N_{\mathrm{mc}}} \mathbb{1}_{\{H_{1,1}^{\mathrm{x}} x_1(k) > 1\}}$. Furthermore, we can see the total number of constraint violations for subsystem 1, i.e., $c_{\mathrm{total}} = \sum_{q=1}^{N_{\mathrm{mc}}} \sum_{k=0}^{14} c_{\mathrm{v}}(k)$. For comparison, we computed two central solutions, where we set up the central SMPC scheme from [59] according to:

i) The distributed design procedure from this chapter.

ii) The centralized design procedure from Farina and Scattolini [59].

By reducing the accuracy level ϵ_{c}, the average number of iterations increases, which results from the stopping condition (5.28). Therefore, ϵ_{c} influences the optimality of the solution, as well as the conservatism of the chance constraint satisfaction, i.e., for $\epsilon_{\mathrm{c}} \to 0$, we restore the central solution, cf. Chapter 3. Furthermore, it can be seen that for different values ϵ_{c}, the average cost and the number of cumulative constraint violations vary only slightly

Table 5.1: Impact of ϵ_{c} on the performance.

Controller	ϵ_{c}	av[it]	max[it]	av[J]	c_{total}	c_{wc}
DSMPC	$5 \cdot 10^{-3}$	25.03	36	112.68	439	78.90%
DSMPC	$5 \cdot 10^{-4}$	37.68	55	113.13	668	77.70%
DSMPC	$5 \cdot 10^{-5}$	49.73	75	113.64	684	77.50%
i)	0	–	–	113.67	734	77.50%
ii)	0	–	–	113.06	746	74.10%

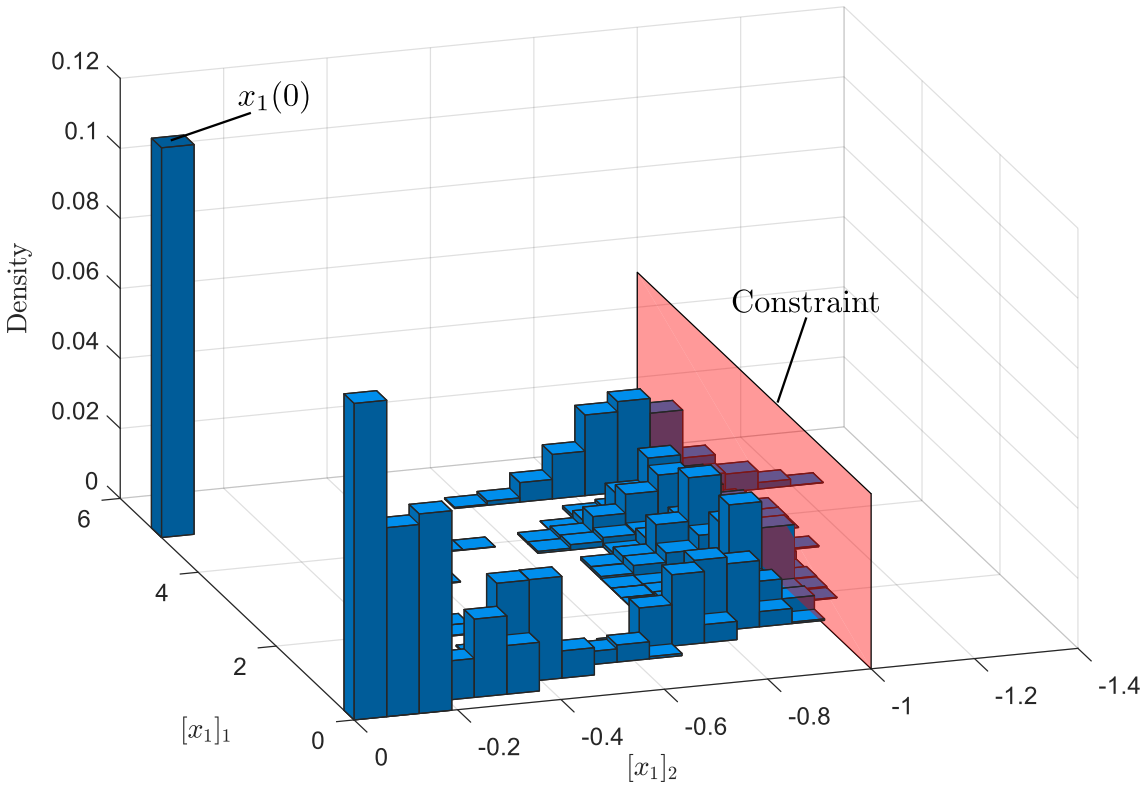

Figure 5.1: Density plot of the closed-loop trajectories of subsystem 1. The red surface denotes the constraint $[x_1]_2 \geq -1$.

compared to the central case i). In each scenario, the chance constraints of level $p_i^{\mathrm{x}} \geq 0.7$ are empirically verified. Note that due to the block-diagonal upper bounding via (5.7), the constraint satisfaction is more conservative compared to the centralized case.

Remark 5.5. *To cope with inexact solutions, similar techniques as in Chapter 3 can be used. However, to ensure that the consolidated mean and covariance trajectories, defined similarly to Definition 3.3, verify the constraints* (5.11)*, a robust constraint tightening must be applied to account for the inexact mean and covariance predictions resulting from* (5.28) *for positive* ϵ_{c}*. This, however, goes beyond the scope of the chapter.*

Remark 5.6. *The augmentation factor* ρ *should be selected in appropriate scale to the cost function* (5.19a)*. If* ρ *is too small, the primary objective is the minimization of the cost function* (5.19a)*. As a consequence, the number of iterations until convergence increases. If* ρ *is too large, the primary objective is the fulfillment of the consensus constraint. Hence, the minimization of the MPC cost becomes less important and the MPC optimization problem 5.2.3 gets solved suboptimally.*

In the following, we quantitatively investigate the effect of the number of subsystems M on the online computational demand. In Figure 5.2, we can see for $M \in \{5, 10, 15\}$ the number of iterations averaged over 100 Monte-Carlo simulations. We consider the same network topology, dynamics, constraints and initial values as before. It can be seen that the number of iterations required for convergence grows slightly as the number of subsystems M increases. This is due to the fact that more subsystems result in a larger global cost (5.19a) and therefore more iterations are required to minimize this cost.

To summarize: The number of subsystems M affects the number of iterations only marginally, which verifies the scalability of our approach. In order to further reduce the computational demand, we can sacrifice the optimality of the solution by increasing ϵ_c.

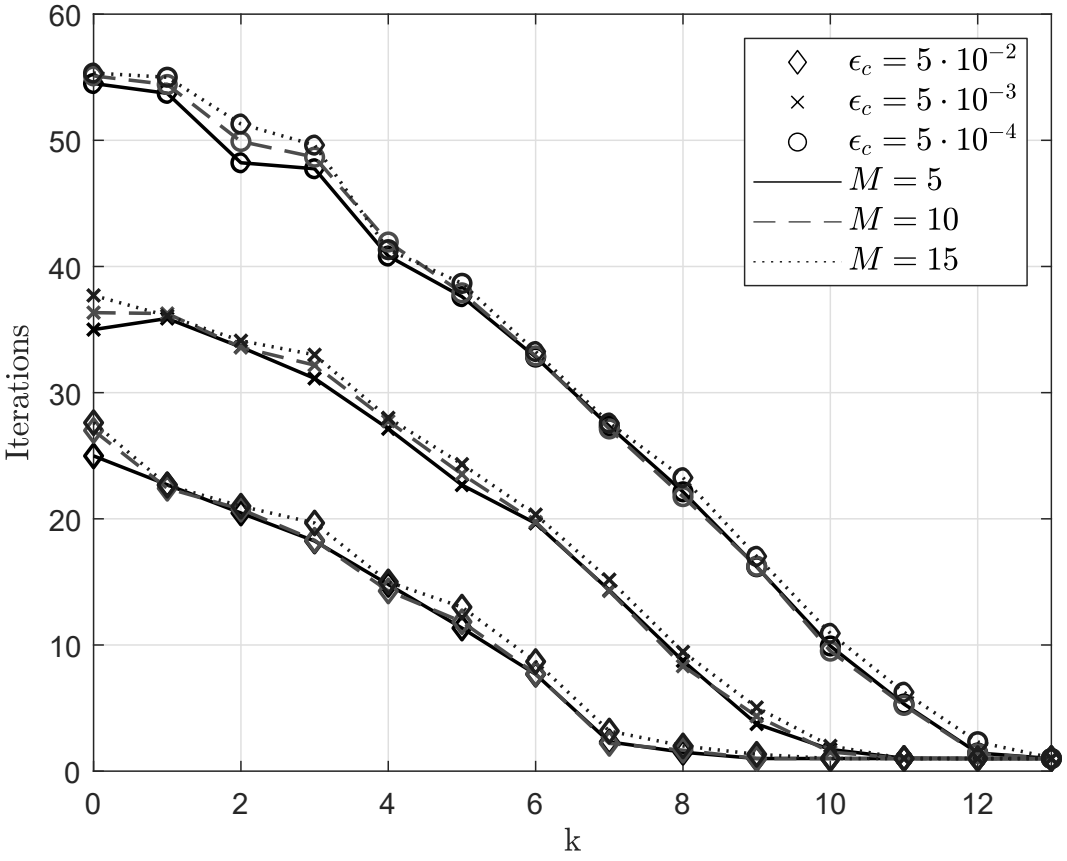

Figure 5.2: Quantitative impact of the number of subsystems M and exactness ϵ_c on the average number of iterations of Algorithm 7.

5.6 Summary

In this chapter, we have presented a DSMPC algorithm for distributed systems with unbounded multiplicative uncertainty. The chance constraints are approximated with Cantelli's inequality, while the cost function is analytically reformulated by resorting to mean-variance predictive dynamics. Each component is defined so that the controller can be synthesized and operated in a fully distributed manner. The distributed design guarantees recursive feasibility, point-wise convergence of the states and predictive chance constraint satisfaction. We presented an ADMM-based distributed MPC algorithm to solve the proposed distributed SDP, while the properties of the controller were demonstrated on a numerical example.

5.7 Proofs

Proof of Lemma 5.3.6

Constraints (5.26b) and (5.26c) are reformulations of (5.18a) and (5.18b) under usage of the support function of the 1-level set of the elliptical terminal region, see [42] for details. Constraint (5.26d) enforces (5.17), which will be shown in the following. Recall that $\psi = \min\{\psi_1, \dots, \psi_M\}$ and P is block-diagonal, thus we have the equivalence for all $\alpha \geq 0$

$$\|P_i^{-1}\|\alpha \leq \psi_i \quad \forall i \in \mathcal{M} \iff \|P^{-1}\|\alpha \leq \psi.$$

It remains to show the equivalence of the latter and (5.17). Substitution of $z = P^{-\frac{1}{2}}\tilde{z}$ into the terminal set (5.25) yields

$$\forall z \in \mathbb{Z}_\mathrm{f} : z^\top P z \leq \alpha \iff \tilde{z}^\top \tilde{z} \leq \alpha \tag{5.29}$$

and by substitution into (5.17) that

$$zz^\top \preceq \psi I \iff P^{-\frac{1}{2}}\tilde{z}\tilde{z}^\top P^{-\frac{1}{2}} \preceq \psi I.$$

Taking the norm on both sides yields

$$\|P^{-\frac{1}{2}}\tilde{z}\tilde{z}^\top P^{-\frac{1}{2}}\| \leq \|P^{-1}\|\|\tilde{z}\tilde{z}^\top\| \overset{(5.29)}{\leq} \|P^{-1}\|\alpha \leq \psi,$$

where the second inequality is due to (5.29) and the rank one matrix $\tilde{z}\tilde{z}^\top$, which implies that $\|\tilde{z}\tilde{z}^\top\| = \tilde{z}^\top\tilde{z}$. Since all constraints are convex, maximization of $\hat{\alpha}$ yields the largest feasible level set $\mathbb{Z}_\mathrm{f}$. □

Proof of Theorem 5.4.2

The proof is inspired by [59] and consists of two parts.

Recursive feasibility Assume that at time k a feasible solution to Problem 5.2.3 exists. At time $k+1$ we consider for each subsystem $i \in \mathcal{M}$ the shifted optimal solution $v_i(t|k+1) = v_i^*(t+1|k)$ for all $t = 0, \dots, N-2$ appended with the terminal controller $v_i(N-1|k+1) = K_{\mathcal{N}_i} z_{\mathcal{N}_i}^*(N|k)$, while the corresponding state and variance sequences are given by $z_i(t|k+1) = z_i^*(t+1|k)$ and $\Sigma_i(t|k+1) = \Sigma_i^*(t+1|k)$ for all $t = 0, \dots, N-1$.

From feasibility at time k follows that the state and input constraints (5.11a), (5.11b) are verified for any pair $(z_i(t|k+1), \Sigma_i(t|k+1))$ and $(v_i(t|k+1), \Sigma_i(t|k+1))$ for all $t = 0, \dots, N-1$.

At time $t = N$, given the terminal constraint (5.27c) and the invariance property of the local terminal set (Lemma 5.3.3), we have that

$$z_i(N|k+1) = A_{\mathcal{N}_i,\mathrm{K}} z_{\mathcal{N}_i}(N|k) \in \mathbb{Z}_{\mathrm{f},i}(\alpha_i(k+1)) \quad \forall i \in \mathcal{M},$$

while from Lemma 5.3.4 we can deduce that $z(N|k+1) \in \hat{\mathbb{Z}}_{\mathrm{f}}(\alpha_1(k+1), \dots, \alpha_M(k+1)) \subseteq \mathbb{Z}_{\mathrm{f}}$, which verifies the global terminal constraint (5.14) at time $k+1$. In view of the shifted candidate solution, we further have

$$\begin{aligned}
&\hat{\Sigma}_i(N|k+1) = \hat{\Sigma}_i(N+1|k) \\
&\overset{(5.8)}{=} \tilde{C}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(N|k) \tilde{C}_{\mathcal{N}_i,K}^\top + \tilde{A}_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}(N|k) \tilde{A}_{\mathcal{N}_i,\mathrm{K}}^\top + \tilde{C}_{\mathcal{N}_i,\mathrm{K}} z_{\mathcal{N}_i}(N|k) (z_{\mathcal{N}_i}(N|k))^\top \tilde{C}_{\mathcal{N}_i,K}^\top \\
&\overset{(5.23)}{\preceq} A_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} A_{\mathcal{N}_i,\mathrm{K}}^\top + C_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} C_{\mathcal{N}_i,K}^\top + C_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathrm{f},\mathcal{N}_i} C_{\mathcal{N}_i,K}^\top \\
&\overset{(5.27d)}{\preceq} \hat{\Sigma}_{\mathrm{f},i},
\end{aligned}$$

which verifies the terminal constraint (5.27d) at time $k+1$. Hence, both global terminal constraints (5.14) and (5.15) are verified at time $k+1$, which implies satisfaction of the chance constraints (5.2) for all times $k \in \mathbb{N}$.

Convergence Next, we prove point-wise convergence of the state trajectories. At time step $k+1$, we have to consider the possible shifted optimal solution due to (5.19e), i.e., Mode 2. Consider the value function $V(k) = V_{\mathrm{m}}(k) + V_{\mathrm{v}}(k)$ of the global system, where the following inequality holds due to optimality

$$V(k+1) \leq J_{\mathrm{m}}(z(\cdot|k+1), v(\cdot|k+1)) + J_{\mathrm{v}}(\hat{\Sigma}(\cdot|k+1)),$$

while the suboptimal mean and variance cost result from the shifted candidate solution due to recursive feasibility. The suboptimal mean cost is given by

$$\begin{aligned}
&J_{\mathrm{m}}(z(\cdot|k+1), v(\cdot|k+1)) \\
&= V_{\mathrm{m}}(k) - \sum_{i=1}^{M} \Big(||z_i(0|k)||_{Q_i}^2 + ||v_i^*(0|k)||_{R_i}^2 - ||z_i^*(N|k)||_{Q_i}^2 - ||K_{\mathcal{N}_i} z_{\mathcal{N}_i}^*(N|k)||_{R_i}^2 \\
&\qquad + ||z_i^*(N|k)||_{P_i}^2 - ||A_{\mathcal{N}_i,\mathrm{K}} z_{\mathcal{N}_i}^*(N|k)||_{P_i}^2 \Big) \\
&\overset{(5.21a)}{\leq} V_{\mathrm{m}}(k) - \sum_{i=1}^{M} \Big(||z_i(0|k)||_{Q_i}^2 + ||v_i^*(0|k)||_{R_i}^2 + ||z_{\mathcal{N}_i}^*(N|k)||_{\tilde{W}}^2 - \|z_{\mathcal{N}_i}^*(N|k)\|_{\Gamma_{\mathcal{N}_i}}^2 \Big), \quad (5.30)
\end{aligned}$$

where $\tilde{W} = C_{\mathcal{N}_i,K}^\top P_i C_{\mathcal{N}_i,\mathrm{K}}$. Note that $||z_i^*(N|k)||_{Q_i}^2 = ||z_{\mathcal{N}_i}^*(N|k)||_{\bar{Q}_i}^2$ and $||z_i^*(N|k)||_{P_i}^2 = ||z_{\mathcal{N}_i}^*(N|k)||_{\bar{P}_i}^2$. The suboptimal variance cost J_v is given by

$$\begin{aligned}
J_\mathrm{v}(\hat{\Sigma}(\cdot|k+1)) = V_\mathrm{v}(k) - \sum_{i=1}^{M} \Bigg(& \mathrm{tr}(Q_i \hat{\Sigma}_i(0|k)) + \mathrm{tr}(K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(0|k)) - \mathrm{tr}(Q_i \hat{\Sigma}_i^*(N|k)) \\
& - \mathrm{tr}(K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}^*(N|k)) + \mathrm{tr}\Big[P_i \hat{\Sigma}_i^*(N|k) - P_i A_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}^*(N|k) A_{\mathcal{N}_i,\mathrm{K}}^\top \\
& - P_i (C_{\mathcal{N}_i,\mathrm{K}} \hat{\Sigma}_{\mathcal{N}_i}^*(N|k) C_{\mathcal{N}_i,\mathrm{K}}^\top) - P_i (C_{\mathcal{N}_i,\mathrm{K}} z_{\mathcal{N}_i}^*(N|k) z_{\mathcal{N}_i}^{*,\top}(N|k) C_{\mathcal{N}_i,\mathrm{K}}^\top) \Big] \Bigg) \\
\overset{(5.21a)}{\leq} V_\mathrm{v}(k) - \sum_{i=1}^{M} \Bigg(& \mathrm{tr}(Q_i \hat{\Sigma}_i(0|k)) + \mathrm{tr}(K_{\mathcal{N}_i}^\top R_i K_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}(0|k)) \\
& - ||z_{\mathcal{N}_i}^*(N|k)||_{\tilde{W}}^2 - \mathrm{tr}(\Gamma_{\mathcal{N}_i} \hat{\Sigma}_{\mathcal{N}_i}^*(N|k)) \big] \Bigg),
\end{aligned} \tag{5.31}$$

where we used $\mathrm{tr}(Q_i \hat{\Sigma}_i^*(N|k)) = \mathrm{tr}(\bar{Q}_i \hat{\Sigma}_{\mathcal{N}_i}^*(N|k))$, $\mathrm{tr}(P_i \hat{\Sigma}_i^*(N|k)) = \mathrm{tr}(\bar{P}_i \hat{\Sigma}_{\mathcal{N}_i}^*(N|k))$ and the cyclic invariance property of the trace to factor out $\hat{\Sigma}_{\mathcal{N}_i}^*(N|k)$. Furthermore, note that $||z_{\mathcal{N}_i}^*(N|k)||_{\tilde{W}}^2 = \mathrm{tr}(P_i C_{\mathcal{N}_i,\mathrm{K}} z_{\mathcal{N}_i}^*(N|k)(z_{\mathcal{N}_i}^*(N|k))^\top C_{\mathcal{N}_i,\mathrm{K}}^\top)$. After combining (5.30) and (5.31), we obtain

$$\begin{aligned}
V(k+1) &\leq V(k) - \sum_{i=1}^{M} \Big(\mathbb{E}\left(||x_i(k)||_{Q_i}^2 + ||u_i(k)||_{R_i}^2 \right) - \mathbb{E}\left(\|x_{\mathcal{N}_i}\|_{\Gamma_{\mathcal{N}_i}}^2 \right) \Big) \\
&\overset{(5.20b)}{\leq} V(k) - \mathbb{E}(||x(k)||_Q^2 + \|u(k)\|_R^2).
\end{aligned}$$

Using standard arguments we conclude that $\mathbb{E}\left(||x(k)||_Q^2\right) \to 0$ as $k \to \infty$. □

Part II

Distributionally Robust Model Predictive Control

6 Wasserstein Distributionally Robust Model Predictive Control

In this chapter, we present two DR-MPC frameworks with Wasserstein ambiguity sets for linear systems subject to additive stochastic uncertainty. In Section 6.2, we present a novel scenario-based DR-MPC with indirect feedback that allows the use of correlated stochastic processes, which has not yet been considered in the related literature, cf. Table 1.3. We investigate for linear and nonlinear tube controllers whether chance constraints can be ensured in a distributionally robust manner, and moreover, whether a distributionally robust performance bound can be established. In Section 6.3, we strengthen the distributional assumption by requiring that the random variables are zero-mean and i.i.d. so that we can obtain an analytical DR-MPC scheme. In addition, we use an indirect feedback initialization that decouples the closed-loop error from the MPC predictions, which allows us to define so-called distributionally robust PRS. In this way, we obtain an offline constraint tightening mechanism that conditions the chance constraints via indirect feedback on the closed-loop error, which leads to satisfaction of the chance constraints in closed-loop. For related work, please refer to Section 1.2.2. This chapter is based on the publications [112] [1] and [111] [2].

6.1 Problem description

We consider discrete linear time-invariant systems of the form

$$x(k+1) = Ax(k) + Bu(k) + w(k) \quad \forall k \in \mathbb{N}, \tag{6.1}$$

where $x(k) \in \mathbb{X} \subseteq \mathbb{R}^n$, $u(k) \in \mathbb{U} \subseteq \mathbb{R}^m$ and $w(k) \in \mathbb{W} \subseteq \mathbb{R}^n$ denote the state, input and disturbance vectors, while $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ are matrices of conformal dimension. We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ for a finite sequence $\bar{w} : \Omega \to \mathbb{W}^{N_\mathrm{T}}$ of random variables $w : \Omega \to \mathbb{W}$, i.e., $\bar{w} = \{w(k)\}_{k=0}^{N_\mathrm{T}-1}$, where $N_\mathrm{T} \in \mathbb{N}$ denotes a large but finite task horizon.

The sequence $\bar{w}$ has a joint probability measure $\mu_{\bar{\mathrm{w}}} : \mathbb{W}^{N_\mathrm{T}} \to [0,1]$, defined such that

[1] C. Mark and S. Liu. "Data-driven distributionally robust model predictive control: An indirect feedback approach". In: *arXiv preprint arXiv:2109.09558* (2021). Submitted to International Journal of Robust and Nonlinear Control (2022).

[2] C. Mark and S. Liu. "Stochastic MPC with Distributionally Robust Chance Constraints". In: *Proc. 21st IFAC World Congress.* extended version: arXiv:2005.00313. 2020, pp. 7136–7141 ©2020 the authors.

$\mu_{\bar{\mathrm{w}}}(F) = \mathbb{P}(\omega \in \Omega : \bar{w}(\omega) \in F)$ for all $F \in \mathcal{B}(\mathbb{W}^{N_\mathrm{T}})$, where $\mathcal{B}(\mathbb{W}^{N_\mathrm{T}})$ is the Borel σ-algebra on $\mathbb{W}^{N_\mathrm{T}}$. We say that $\bar{w} \sim \mu_{\bar{\mathrm{w}}}$, where $\mu_{\bar{\mathrm{w}}}$ is the true distribution (pushforward measure).

For the sake of simplicity, we assume that the pair (A, B) is controllable and perfect state measurement is available at each time instant $k \in \mathbb{N}$. The system dynamics are subject to $r \in \mathbb{N}$ individual state chance constraints

$$\mathbb{P}(h_i^\top x(k) \leq 1) \geq p_\mathrm{x}^i \quad \forall i \in \{1, \ldots, r\}, \tag{6.2}$$

as well as hard input constraints

$$u(k) \in \bar{\mathbb{U}}, \tag{6.3}$$

where $\bar{\mathbb{U}} \subseteq \mathbb{U}$ is a convex set that contains the origin. We consider a cost function $J : \mathbb{X}^{N_\mathrm{T}+1} \times \mathbb{U}^{N_\mathrm{T}} \to \mathbb{R}_{\geq 0}$ and aim to solve the following finite-horizon stochastic optimal control problem

$$\begin{aligned}
\min_{x,u} \quad & \mathbb{E}_{\mu_{\bar{\mathrm{w}}}}\Big(J\big(x(0, \ldots, N_\mathrm{T}), u(0, \ldots, N_\mathrm{T} - 1)\big)\Big) && \text{(6.4a)}\\
\text{s.t.} \quad & x(k+1) = Ax(k) + Bu(k) + w(k) && \text{(6.4b)}\\
& \mathbb{P}(h_i^\top x(k) - 1 \leq 0) \geq p_\mathrm{x}^i \quad \forall i \in \{1, \ldots, r\} && \text{(6.4c)}\\
& u(k) \in \mathbb{U} && \text{(6.4d)}\\
& \bar{w} = \{w(0), \ldots, w(N_\mathrm{T} - 1)\} \sim \mu_{\bar{\mathrm{w}}} && \text{(6.4e)}\\
& x(0) = x_0 && \text{(6.4f)}
\end{aligned}$$

for all $k = 0, \ldots, N_\mathrm{T} - 1$.

In SMPC, one typically assumes that the distribution $\mu_{\bar{\mathrm{w}}}$ is known, e.g., as we have done in part one of this thesis. However, from a practical point of view, this is quite limiting, since the statistics of the underlying random variables in any real-world application usually have to be estimated from finite data. Therefore, we cannot use the true distribution to reformulate the FH-SOCP (6.4), which in light of this contains multiple sources of intractability:

(i) The expectation in (6.4a) is taken w.r.t. the true probability distribution $\mu_{\bar{\mathrm{w}}}$.

(ii) The chance constraints (6.4c) are evaluated under the true probability measure $\mathbb{P}$.

(iii) Optimizing over general control inputs u in the presence of possibly unbounded disturbances w yields an infinite-dimensional optimization problem.

6.1.1 Distributionally Robust Optimization

In the following, we introduce concepts from DRO to reformulate the FH-SOCP, such that the intractability sources (i) and (ii) can be cast into tractable surrogates. We follow a data-driven approach and assume the existence of a (possibly small) amount of data.

Assumption 6.1 (Distributional assumptions)**.**

1. *The true probability distribution $\mu_{\bar{\mathrm{w}}}$ is light-tailed.*
2. *There exists a set $\hat{\mathcal{W}} = \{\hat{\bar{w}}_j\}_{j=1}^{\bar{N}}$ that consists of $\bar{N} \in \mathbb{N}$ independent and identically distributed disturbance trajectories $\hat{\bar{w}}_j = \{\hat{w}_j(0), \ldots, \hat{w}_j(N_\mathrm{T}-1)\} \sim \mu_{\bar{\mathrm{w}}}$.*

Remark 6.1. *The second part of Assumption 6.1 essentially gives us at each time instant $k \in [0, N_\mathrm{T}-1]$ a collection of $\bar{N}$ i.i.d. samples. The sequence $\bar{w} = \{w(0), \ldots, w(N_\mathrm{T}-1)\}$ can also be viewed as a discrete-time stochastic process, where we distinguish between the Markovian setting, i.e., stage-wise independent distributions or the more general setting of stage-wise dependent distributions, e.g., due to time correlation.*

First, we would like to point out that the chance constraints (6.4c) *can be verified in both cases, as we evaluate them at each time step k individually. However, with respect to the cost function* (6.4a)*, the situation is much more delicate, since we have to consider a multi-stage optimization problem, where in the most general setting the distribution of $w(k)$ at stage k depends on the entire history $\{w(0), \ldots, w(k-1)\}$. This leads to nested optimization problems [138] that are difficult to solve numerically. However, by assuming that the stochastic process is stage-wise independent (Markov property), a much simpler optimization problem can be derived since the stochastic process is independent of the historical data. In the related DRO literature, this concept is also called rectangularity [138].*

To avoid the issue of nested optimization problems, we make a second distributional assumption about stage-wise independent distributions, which is necessary for some (not all) of our theoretical results.

Assumption 6.2. *The distribution $\mu_{\bar{\mathrm{w}}}$ is stage-wise independent, i.e., the joint distribution $\mu_{\bar{\mathrm{w}}}$ can be written as a product of marginal distributions so that*

$$\mu_{\bar{\mathrm{w}}} = \mu(w(0)) \times \ldots \times \mu(w(N_\mathrm{T}-1)),$$

where $\mu(w(k))$ denotes the distribution of $w(k)$ at time $k \in [0, \ldots, N_\mathrm{T}-1]$.

Sample average approximation A straight forward approach to solve problem (6.4) is to evaluate (6.4a) and (6.4c) with the empirical distribution of $\mu_{\bar{\mathrm{w}}}$, i.e.,

$$\hat{\mu}_{\bar{\mathrm{w}}} = \hat{\mu}(w(0)) \times \ldots \times \hat{\mu}(w(N_\mathrm{T}-1)),^{3}$$

where $\hat{\mu}(w(k)) = N_\mathrm{s}^{-1} \sum_{j=1}^{N_\mathrm{s}} \delta_{\hat{w}_j(k)}$ concentrates the probability mass uniformly on the N_s i.i.d. samples $\hat{w}_j(k) \in \hat{\mathcal{W}}$ via the Dirac delta measure $\delta_{\hat{w}_j(k)}$. In other words, we approximate

[3] The *true distribution* $\mu_{\bar{\mathrm{w}}}$ of a sequence $\bar{w} = \{w(0), \ldots, w(N_\mathrm{T}-1)\}$ can always be approximated as an empirical product distribution $\hat{\mu}_{\bar{\mathrm{w}}}$ consisting of empirical marginal distributions $\hat{\mu}(w(k))$ for $k \in [0, N_\mathrm{T}-1]$. This should not be confused with stage-wise independency condition from Assumption 6.2, where the *true distribution* can be written as a product of marginal distributions $\mu(w(k))$.

the FH-SOCP with a sample average approximation, which we denote as the SAA-OCP. The result of the SAA-OCP provides an optimal input sequence $\hat{u}^*(\cdot)$ that minimizes the in-sample performance, i.e., the expected cost in terms of $\hat{\mu}_{\bar{\mathrm{w}}}$, while the chance constraints are only empirically verified. Applying the input $\hat{u}^*(\cdot)$ to the dynamics (6.1) introduces new disturbances $w(k)$ into the system that were not part of the decision making process. This may result in poor out-of-sample performance (expected cost w.r.t. the true distribution $\mu_{\bar{\mathrm{w}}}$) and a potential violation of the chance constraints (6.4c).

Remark 6.2. *The optimizer of the SAA-OCP converges almost surely to the optimizer of* (6.4) *when* N_s *tends to infinity, whereas for small* N_s *the SAA control input* $\hat{u}^*(\cdot)$ *performs poorly when applied to the real system* (6.1)*. Unfortunately, the sample size cannot be chosen arbitrarily large, since the sample complexity of the SAA-OCP increases at least linearly in the sample size* N_s*, which ultimately boils down to a trade-off between accuracy and computational effort [90]. This is our main motivation to study distributionally robust SOCPs that allow us to derive meaningful control inputs from a small sample size* N_s *such that the system states satisfy the chance constraints with high probability.*

Wasserstein DRO With the goal of robustifying the FH-SOCP (6.4) against distributional uncertainties, we introduce a discrepancy-based ambiguity set using the Wasserstein metric [126]. For the sake of introduction, consider an arbitrary measure space Ξ and the space of all Borel probability measures $\mathcal{P}(\Xi)$ with q-th finite moment for $q \in [1, \infty]$ that contains all distributions ν supported on Ξ.

Definition 6.1 (Wasserstein metric [126])**.** *Let* $q \in [1, \infty]$*. The* q*-Wasserstein metric* $d_\mathrm{W}^\mathrm{q}(\nu_1, \nu_2) : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \to \mathbb{R}_{\geq 0}$ *is defined as*

$$d_\mathrm{W}^\mathrm{q}(\nu_1, \nu_2) \coloneqq \inf_{\gamma \in \mathcal{H}(\nu_1, \nu_2)} \left(\int_{\Xi \times \Xi} \|\xi_1 - \xi_2\|_q \gamma(d\xi_1, d\xi_2) \right),$$

where $\mathcal{H}(\nu_1, \nu_2)$ *is the set of all joint distributions of* ξ_1 *and* ξ_2 *with marginal distributions* $\nu_1 \in \mathcal{M}(\Xi)$ *and* $\nu_2 \in \mathcal{M}(\Xi)$*, respectively.*

The Wasserstein metric measures distances between probability distributions by solving an optimal mass transport problem, where the shortest distance is characterized by the optimal transport plan γ.

Definition 6.2. *The Wasserstein ambiguity set centered at the distribution* ν *with radius* $\epsilon \in \mathbb{R}_{\geq 0}$ *is given by*

$$\mathbb{B}_\epsilon(\nu) \coloneqq \{\nu' \in \mathcal{M}(\Xi) \mid d_\mathrm{W}^\mathrm{q}(\nu, \nu') \leq \epsilon\}.$$

In a data-driven framework, one typically centers the Wasserstein ball at the empirical distribution $\hat{\mu}_{\bar{\mathrm{w}}}$.

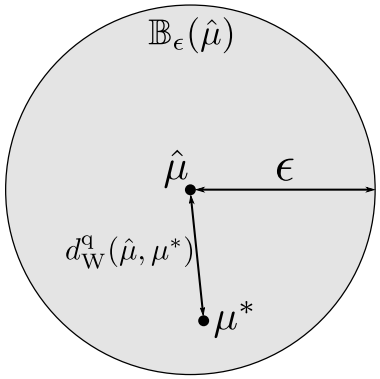

Figure 6.1: Illustration of a Wasserstein ball of radius ϵ centered at the empirical distribution $\hat{\mu}$. The distance from $\hat{\mu}$ to any distribution within radius ϵ is measured in the q-Wasserstein distance, while the figure illustrates the distance to the true distribution μ^*. The parameter ϵ regulates the amount of robustification, i.e., how many distributions should be considered within the distributionally robust optimization problem.

In view of Definition 6.2, we extend the definition of a Wasserstein ball to product distributions, such that

$$\mathbb{B}_\epsilon(\hat{\mu}_{\bar{\mathrm{w}}}) \coloneqq \mathbb{B}_\epsilon(\hat{\mu}(w(0))) \times \ldots \times \mathbb{B}_\epsilon(\hat{\mu}(w(N_{\mathrm{T}}-1))).$$

The last point we want to cover in this section is the choice of the Wasserstein radius ϵ. To this end, we recall that the Wasserstein ball itself is a random object as it is constructed from data, which in turn makes it impossible to give robust guarantees of the type $\mu_{\bar{\mathrm{w}}} \in \mathbb{B}_\epsilon(\hat{\mu}_{\bar{\mathrm{w}}})$. However, due to the first part of Assumption 6.1, we can resort to the concentration inequality result [126, Thm. 3.4], based on which one can find an optimal Wasserstein radius ϵ that satisfies the following assumption, see also [126, Thm. 3.5].

Assumption 6.3. *For a given confidence level $\beta \in (0,1)$ and sample size N_{s} there exists a Wasserstein radius $\epsilon(\beta, N_{\mathrm{s}})$, such that*

$$\mathbb{P}^{N_{\mathrm{s}}}(\mu_{\bar{\mathrm{w}}} \in \mathbb{B}_\epsilon(\hat{\mu}_{\bar{\mathrm{w}}})) \geq 1-\beta,^{4}$$

where $\mu_{\bar{\mathrm{w}}}$ and $\hat{\mu}_{\bar{\mathrm{w}}}$ are the true and empirical distributions of $\bar{w}$.

Remark 6.3. *The concentration inequality result provides a strong theoretical guarantee [126, Thm. 3.4], which is unfortunately of limited use in practice, since the Wasserstein radius still depends on constants of the unknown distribution. In practice, one typically uses tools from machine learning, such as K-fold cross validation [129], to calibrate the Wasserstein ambiguity set for a certain sample size N_{s} given all available samples $\bar{N}$ from Assumption 6.1 [126].*

[4] In case that $\mu_{\bar{\mathrm{w}}}$ is a product distribution (Assumption 6.2), the notation $\mathbb{P}^{N_{\mathrm{s}}}(\mu_{\bar{\mathrm{w}}} \in \mathbb{B}_\epsilon(\hat{\mu}_{\bar{\mathrm{w}}})) \geq 1-\beta$ denotes $\mathbb{P}^{N_{\mathrm{s}}}(\mu(w(k)) \in \mathbb{B}_\epsilon(\hat{\mu}(w(k)))) \geq 1-\beta$ for all $k \in [0, \ldots, N_{\mathrm{T}}-1]$. $\mathbb{P}^{N_{\mathrm{s}}}$ denotes the N_{s}-fold product probability measure $\mathbb{P}^{N_{\mathrm{s}}} = \mathbb{P} \times \ldots \times \mathbb{P}$, which results from the N_{s} i.i.d. samples.

In the following, we present the main result from Mohajerin Esfahani and Kuhn [126], which establishes a convex reduction of an infinite dimensional worst-case expectation problem.

Proposition 6.1.1 ([126]). *Let $\{\hat{\xi}_j\}_{j=1}^{N_\mathrm{s}}$ be a set of i.i.d. samples of ξ and assume that $f(\xi)$ is proper, convex and lower semicontinuous, $\xi \in \Xi = \mathbb{R}^n$ and let $q \in [1, \infty]$, $\epsilon > 0$. Then it holds that*

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(\xi))} \mathbb{E}_\nu(f(\xi)) = \inf_{\lambda \geq 0} \lambda\epsilon + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} \sup_{\xi \in \mathbb{R}^n} (f(\xi) - \lambda \|\xi - \hat{\xi}_j\|_q),$$

where $\hat{\mu}(\xi) = N_\mathrm{s}^{-1} \sum_{j=1}^{N_\mathrm{s}} \delta_{\hat{\xi}_j}$ is the empirical distribution and λ the Wasserstein penalty.

Proof. The proof follows immediately from [126, Thm. 4.2], which relies on marginalizing and dualizing the Wasserstein constraint $\nu \in \mathbb{B}_\epsilon(\hat{\mu}(\xi))$. □

Using the former definition, we can robustify the stochastic control problem (6.4) to sampling errors introduced by a small sample size N_s. In the subsequent section, we present an indirect feedback DR-MPC formulation that uses concepts from Wasserstein DRO, where we distinguish between the cases of i.i.d. and non-i.i.d. disturbance sequences that lead to different guarantees for the closed-loop system.

6.2 Scenario-based indirect feedback DR-MPC

In the following, we address the intractability source (iii) of the FH-SOCP (6.4) by solving the optimization problem over a shortened prediction horizon $N \in \mathbb{N}$, where $N \ll N_\mathrm{T}$. We pursue an indirect feedback tube-based approach [76] and split the dynamics (6.1) into a nominal and error part, such that the true state satisfies $x(k) = z(k) + e(k)$. Analogously, we separate the input $u(k)$ into a nominal part $v(k)$ and an error part $e_\mathrm{u}(k) = \pi(e(k))$, where $\pi(\cdot)$ is a so-called tube controller, so that $u(k) = v(k) + \pi(e(k))$. The resulting decoupled closed-loop nominal and error dynamics are then given by

$$z(k+1) = Az(k) + Bv(k) \tag{6.5a}$$
$$e(k+1) = Ae(k) + B\pi(e(k)) + w(k) \tag{6.5b}$$

with initial conditions $z(0) = x(0)$ and $e(0) = 0$. To make predictions at time k, we define the t-step predictive dynamics

$$x(t+1|k) = Ax(t|k) + Bu(t|k) + w(t|k) \tag{6.6a}$$
$$z(t+1|k) = Az(t|k) + Bv(t|k) \tag{6.6b}$$
$$e(t+1|k) = Ae(t|k) + Be_\mathrm{u}(t|k) + w(t|k) \tag{6.6c}$$
$$e_\mathrm{u}(t|k) = \pi(e(t|k)) \tag{6.6d}$$
$$u(t|k) = v(t|k) + e_\mathrm{u}(t|k), \tag{6.6e}$$

which are coupled to the closed-loop dynamics with $x(0|k) = x(k)$, $z(0|k) = z(k)$, $e(0|k) = e(k)$. Note that in general the closed-loop error $e(k) \neq 0$ for $k \geq 1$, which, unlike in reset-based (direct feedback) SMPC schemes renders the closed-loop error dynamics (6.5b) valid even under the MPC control input, cf. Remark 2.4. The predictive disturbance sequence $W(k) = \{w(k), w(1|k) \dots, w(N-1|k)\}$ is obtained by conditioning $\bar{w}$ (defined in (6.4e)) on all past disturbances, such that

$$\mu_{W(k)} = \mathbb{P}\left(\{w(k), w(1|k), \dots, w(N-1|k)\} \middle| \{w(0), \dots, w(k-1)\}\right), \tag{6.7}$$

while the predictive error sequence $E(k) = \{e(k), e(1|k) \dots, e(N|k)\}$ is obtained from (6.6c) and (6.7), i.e.,

$$\mu_{E(k)} = \mathbb{P}\left(\{e(k), e(1|k) \dots, e(N|k)\} \middle| \{e(0), \dots, e(k-1)\}\right). \tag{6.8}$$

Figure 6.2 illustrates the predictive distributions for the disturbance and error trajectories and highlights the difference to the closed-loop realization.

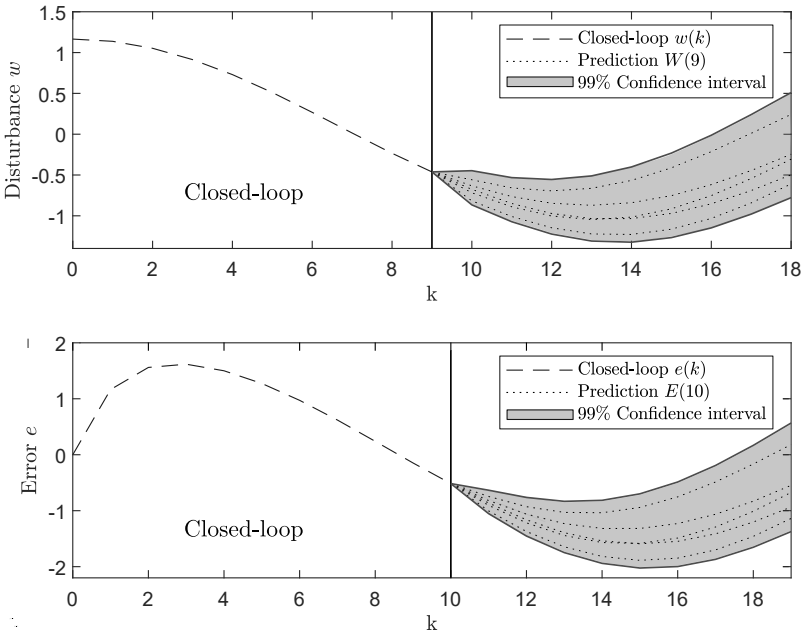

Figure 6.2: Illustration of realized and predictive disturbance and error trajectories. (**Top**) The realized disturbance trajectory for $k = 0, \dots, 9$, while for $k = 9, \dots, 18$ various predictive sequences $W(9) = [w(k), w(1|k), \dots, w(9|k)]^\top$ are drawn from the conditional distribution given $[w(0), \dots, w(9)]^\top$. (**Bottom**) The realized error trajectory for $k = 0, \dots, 10$, while for $k = 10, \dots, 19$ we plot the predicted error sequences that correspond to the disturbance sequences $W(9)$.

6.2.1 Objective function

In the following, we approximate the cost function (6.4a) over a shortened prediction horizon and maximize the expected value over the previously introduced Wasserstein ball. At this point, we invoke Assumption 6.2 to circumvent the issue of nested optimization problems, compare Remark 6.1. We define the distributionally robust MPC cost function as

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_{W(k)})} \mathbb{E}_\nu \big(l_1(x(\cdot|k)) + l_2(u(\cdot|k)) \big), \tag{6.9}$$

where $l_1 : \mathbb{X}^N \to \mathbb{R}_{\geq 0}$ and $l_2 : \mathbb{U}^N \to \mathbb{R}_{\geq 0}$ denote Borel measurable functions. The Wasserstein ball in (6.9) is centered at the empirical predictive distribution of $W(k)$, which indicates that the expectation problem is solved w.r.t. w. Furthermore, since the expectation operator involves both, x and u, this implies that the mapping $w \mapsto (x, u)$ must be jointly linear in w. This property, however, depends on the choice of the tube controller, which can be seen by separating the states and inputs as

$$x(t|k) = z(t|k) + e(t|k) \tag{6.10a}$$

$$u(t|k) = v(t|k) + e_{\mathrm{u}}(t|k), \tag{6.10b}$$

where the predicted error and input error e and e_{u} are given by (6.6c) - (6.6d). The remainder of this section examines for linear and nonlinear tube controllers whether the condition of joint linearity holds, and shows the implications when this condition is violated.

Linear tube controllers

First, we study linear tube controllers of the form $\pi(e) = Ke$, where $K \in \mathbb{R}^{m \times n}$ is a stabilizing controller gain for the matrix pair (A, B). Due to linearity, we can explicitly write the prediction and input error (6.6c) - (6.6d) as

$$e(t|k) = A_{\mathrm{K}}^t e(0|k) + \sum_{i=0}^{t-1} A_{\mathrm{K}}^{t-1-i} w(i|k) \tag{6.11a}$$

$$e_{\mathrm{u}}(t|k) = K A_{\mathrm{K}}^t e(0|k) + \sum_{i=0}^{t-1} K A_{\mathrm{K}}^{t-1-i} w(i|k), \tag{6.11b}$$

where $A_{\mathrm{K}} = A + BK$. Hence, (6.11a)-(6.11b) are both affine functions, which, together with (6.10a) - (6.10b) renders the mapping from $w \mapsto (x, u)$ jointly linear. Based on this observation, we can pose the following result.

Lemma 6.1. *Suppose that Assumptions 6.1 and 6.2 hold. Let the tube controller be a linear mapping $\pi(e) = Ke$ and let the functions l_1 and l_2 be proper, convex and Lipschitz continuous. Specifically, let $L_\phi > 0$ be the Lipschitz constant w.r.t. the q-norm of the mapping $(x, u) \mapsto l_1(x) + l_2(u)$ and assume that $\mathbb{W} = \mathbb{R}^n$. Then, for any $\epsilon > 0$, the distributionally robust cost function* (6.9) *is equal to*

$$J_{\mathrm{DR}}(x(\cdot|k), u(\cdot|k)) \coloneqq L_\phi \epsilon + \mathbb{E}_{\hat{\mu}_{W(k)}} \left(l_1(x(\cdot|k)) + l_2(u(\cdot|k)) \right). \tag{6.12}$$

Proof. The proof can be found in Section 6.5. □

Remark 6.4. *The class of proper, convex and Lipschitz functions includes many cost functions of practical interest, e.g., all norms verify these conditions. For example, if the cost functions l_1 and l_2 are defined as q-norms, the Lipschitz constant L_ϕ of the mapping $(x,u) \mapsto l_1(x) + l_2(u)$ with respect to the q-norm is equal to 1, while L_ϕ must obey the equivalence of norms when the Lipschitz continuity is measured by the r-norm, where $r \neq q$.*

Remark 6.5. *Lemma 6.1 indicates that the distributionally robust cost function (6.9) is equivalent to a sample average approximation plus an additional Wasserstein regularization term $L_\phi \epsilon$ that is independent of any decision variables. Thus, we can neglect this regularization term in an MPC implementation, while still obtaining the same minimizer. This was already mentioned in [126, Remark 6.7] for the case of single stage optimization problems. Note that in other control parameterization the Lipschitz constant $L_\phi = \sup_{\theta\in\Theta}(\|\theta\|_{q,*})$ (see proof of Lemma 6.1) might depend on the decision variables, e.g., as in [45, Lemma A.2]. In this case, the equivalence result presented in Lemma 6.1 still holds, but the term $L_\phi\epsilon$ cannot be neglected in the MPC implementation.*

Nonlinear tube controllers

A nonlinear tube controller $\pi(\cdot)$ does not allow for an explicit representation of the error and input error akin to (6.11a) - (6.11b), i.e.,

$$\begin{aligned} e(t|k) &= A^t e(0|k) + \sum_{i=0}^{t-1} A^{t-1-i}[Be_{\mathrm{u}}(i|k) + w(k+i)] \\ e_{\mathrm{u}}(t|k) &= \pi(e(t|k)), \end{aligned}$$

which implies that the map $w \mapsto (x,u)$ is not jointly linear. In fact, $w \mapsto u$ is nonlinear and thus, it is not possible to represent the worst-case expectation (6.9) together with Proposition 6.1.1 as a tractable convex optimization problem. However, we can still approximate the expected value in (6.9) empirically by generating i.i.d. error and input error samples with (6.6c) - (6.6d)

$$\hat{e}_j(t+1|k) = A\hat{e}_j(t|k) + B\pi(\hat{e}_j(t|k)) + \hat{w}_j(k+t), \quad \hat{w}_j(k+t) \in \hat{\mathcal{W}} \tag{6.13a}$$

$$\hat{e}_{u,j}(t|k) = \pi(\hat{e}_j(t|k)), \tag{6.13b}$$

where $\hat{e}_j(0|k) = e(k)$, while the empirical distributions for e and e_{u} are easily constructed from $j = 1, \ldots, N_{\mathrm{s}}$ sample trajectories

$$\begin{aligned} \hat{\mu}_{E(k)} &= \delta_{e(k)} \times \hat{\mu}(e(1|k)) \times \ldots, \hat{\mu}(e(N-1|k)) \\ \hat{\mu}_{E_u(k)} &= \delta_{e_{\mathrm{u}}(k)} \times \hat{\mu}(e_{\mathrm{u}}(1|k)) \times \ldots, \hat{\mu}(e_{\mathrm{u}}(N-1|k)). \end{aligned} \tag{6.14}$$

At time $t = 0$ we consider the Dirac distribution, since $e(0|k) = e(k)$ and $e_{\mathrm{u}}(0|k) = e_{\mathrm{u}}(k)$ are already realizations of the stochastic process. Finally, the SAA MPC cost function can be defined as

$$J_{\mathrm{SAA}}(x(\cdot|k), u(\cdot|k)) \coloneqq \mathbb{E}_{\hat{\mu}_{E(k)}}\big(l_1(x(\cdot|k))\big) + \mathbb{E}_{\hat{\mu}_{E_u(k)}}\big(l_2(u(\cdot|k))\big), \tag{6.15}$$

where x and u are given by (6.10a) - (6.10b).

Remark 6.6. *Note that we cannot give distributionally robust performance guarantees for nonlinear tube controllers, which, in view of the SAA MPC cost function* (6.15) *allows us to work with correlated disturbance sequences. Therefore, we can use Assumption 6.1 without the additional stage-wise independency condition (Assumption 6.2) as required by Lemma 6.1.*

6.2.2 DR-CVaR state constraints

Chance constraints of type (6.4c) are in the literature also known as Value-at-Risk (VaR) constraints. One of the main problems associated with VaR constraints is that their feasible set is generally not-convex, except in some generic cases, e.g., when the distribution of random variables is log-concave, in which case the feasible set happens to be convex [132]. This is an undesirable property for an MPC optimization problem, as it increases the online complexity of the algorithm and hedges the risk of finding local minima. To alleviate this problem, we formulate the chance constraints (6.2) for the predicted states

$$\mathbb{P}(h_i^\top x(t|k) \leq 1) \geq p_{\mathrm{x}}^i \quad \forall i \in \{1, \dots, r\}, \tag{6.16}$$

which we then relax by using a coherent risk measure called the conditional Value-at-Risk, serving as a convex relaxation of VaR [132].

Definition 6.3. *The conditional Value-at-Risk (CVaR) of a random variable $\xi \sim \mu$ at risk level $p \in (0, 1)$ is defined as*

$$\mathrm{CVaR}_p^\mu(\xi) \coloneqq \inf_{\tau \in \mathbb{R}} \big(-(1-p)\tau + \mathbb{E}_\mu\big((\xi + \tau)_+\big)\big),$$

which penalizes the average loss above the p-th quantile, i.e., above the Value-at-Risk.

From the above definition, it can be seen that we again need the true (unknown) distribution to evaluate the CVaR. This issue is addressed by maximizing the CVaR over all distributions contained in the Wasserstein ambiguity set, resulting in the distributionally robust CVaR.

Based on the earlier definitions, we can relax the chance constraints (6.16) as distributionally robust CVaR constraints, where we define for each half-space $i = 1, \dots, r$ the t-step ahead predicted constraint function as

$$\gamma_i(x(t|k)) \coloneqq h_i^\top x(t|k) - 1.$$

Using the fact that $x(t|k) = z(t|k) + e(t|k)$, we can impose the CVaR constraints in view of the empirical error distribution (6.14) for $t = 0, \dots, N-1$, resulting in

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \mathrm{CVaR}^{\nu}_{p^i_{\mathrm{x}}}\big(h_i^\top (z(t|k) + e(t|k)) - 1\big) = \sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \mathrm{CVaR}^{\nu}_{p^i_{\mathrm{x}}}\big(\gamma_i(x(t|k))\big) \leq 0. \tag{6.17}$$

We then define the distributionally robust CVaR constraint set at time k as

$$\mathbb{X}_{\mathrm{CVaR}} \coloneqq \left\{ \begin{bmatrix} x(0|k) \\ \vdots \\ x(N-1|k) \end{bmatrix} \in \mathbb{R}^{nN} \;\middle|\; \begin{array}{c} \sup\limits_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \mathrm{CVaR}^{\nu}_{p^i_{\mathrm{x}}}\big(\gamma_i(x(t|k))\big) \leq 0 \\ \forall i = \{1, \dots, r\}, \forall t \in \{0, \dots, N-1\} \end{array} \right\},$$

which is intractable in its present form, i.e., the evaluation of the constraint $x(\cdot|k) \in \mathbb{X}_{\mathrm{CVaR}}$ involves multiple infinite-dimensional optimization problems, one for each CVaR constraint. Therefore, we propose the following convex approximation based on Proposition 6.1.1.

Lemma 6.2. *Suppose that Assumption 6.1 holds true. Let $p^i_{\mathrm{x}} \in (0,1)$ and define $p, q \geq 1$, such that $1/p + 1/q = 1$, then*

$$\mathbb{Z} \coloneqq \left\{ \begin{bmatrix} z(0|k) \\ \vdots \\ z(N-1|k) \end{bmatrix} \in \mathbb{R}^{nN} \;\middle|\; \begin{array}{l} \exists \tau_{i,t} \in \mathbb{R}, \lambda_{i,t} \in \mathbb{R}_{\geq 0}, s_{i,j,t} \in \mathbb{R}_{\geq 0} \text{ s.t.} \\ -(1 - p^i_{\mathrm{x}})\tau_{i,t} + \epsilon \lambda_{i,t} + \frac{1}{N_{\mathrm{s}}} \sum_{j=1}^{N_{\mathrm{s}}} s_{i,j,t} \leq 0 \\ \big(\gamma_i(z(t|k) + \hat{e}_j(t|k)) + \tau_{i,t}\big)_+ \leq s_{i,j,t} \\ \|h_i^\top\|_p \leq \lambda_{i,t} \\ \forall j \in \{1, \dots, N_{\mathrm{s}}\}\, \forall i \in \{1, \dots, r\} \\ \forall t \in \{0, \dots, N-1\} \end{array} \right\} \subseteq \mathbb{X}_{\mathrm{CVaR}}, \tag{6.18}$$

where $\hat{e}_j(t|k)$ results from (6.13a).

Proof. The proof can be found in Section 6.5. □

Remark 6.7. *Note that the distributionally robust CVaR constraint penalizes the worst-case expected constraint violation above the p^i_{x}-th quantile of $\gamma_i(x)$ (compare Definition 6.3). Hence,* (6.17) *is a sufficient condition for the chance constraints* (6.16) *to hold [164].*

Remark 6.8. *The choice of p and q-norms as required in Lemma 6.2 can impose different robustness goals, e.g., by setting $q = \infty$, we measure distances in the sample space with $\|e - \hat{e}_j\|_\infty$, which results in a robustification of the Wasserstein penalty λ in the 1-norm, whereas $q = 2$ leads to a robustification of λ in the 2-norm. The relationship between q and p is formalized with the duality condition $1/p + 1/q = 1$.*

Remark 6.9. *For linear tube controllers, we already know from* (6.11a)-(6.11b) *that the error and input error can be written as affine functions of the disturbance w. Thus, we can substitute* (6.11a) *into* (6.18) *and center the Wasserstein ball at the t-step empirical disturbance distribution $\hat{\mu}(w(t|k))$ instead of $\hat{\mu}(e(t|k))$. Then we can simultaneously ensure distributionally robust performance (Lemma 6.1) and distributionally robust CVaR constraint satisfaction (Lemma 6.2).*

6.2.3 Input constraints

We originally indicated that we were imposing hard input constraints, which unfortunately contradicts with some of our main results. Below we discuss nominal and hard input constraints, both of which can be used in our framework but lead to different performance and constraint satisfaction guarantees.

Nominal input constraints

Nominal input constraints are only imposed on $v(t|k)$ acting on (6.6e), i.e., we enforce constraints of the type

$$v(t|k) \in \mathbb{V} = \mathbb{U} \quad \forall t \in \{0, \dots, N-1\}. \tag{6.19}$$

In view of the tube controller $\pi(e)$, we are able to use a linear controller $\pi(e) = Ke$, which ensures that the true input is kept close to the nominal predictions. In view of this, we can use Lemma 6.1 to give distributionally robust performance guarantees, while Lemma 6.2 additionally guarantees distributionally robust state constraint satisfaction. On the downside, the actual input (see (6.22) in the following section) acting to the system (6.1) can deviate greatly from the mean value since the disturbance is unbounded.

Remark 6.10. *Another possibility is to derive distributionally robust chance constraints for the input that at least probabilistically bounds the deviation between $u(t|k)$ and $v(t|k)$.*

Hard input constraints

In the nominal case, we required that the nominal input $v(\cdot|k)$ verifies the hard constraints (6.19), while in the following we require that this condition holds for the predicted tube input $u(t|k)$ given by (6.6e), i.e.,

$$u(t|k) \in \mathbb{U} \quad \forall t \in \{0, \dots, N-1\}. \tag{6.20}$$

To ensure this, we limit the control authority of the tube controller $\pi(\cdot)$, e.g., via a saturated LQR [82], where we make the following assumption.

Assumption 6.4. *The tube controller π satisfies $\pi(e) \in \mathcal{E}_{\mathrm{u}} \subset \mathbb{U}\, \forall e \in \mathbb{R}^n$.*

Similar to tube-based robust MPC [118], we tighten the original input constraints with $\mathbb{V} = \mathbb{U} \ominus \mathcal{E}_{\mathrm{u}}$, where $\mathbb{V}$ denotes the tightened nominal input constraint set. Note that a saturated LQR belongs to the class of nonlinear tube controllers. Thus, if we want to ensure hard input constraints (Assumption 6.4), we cannot achieve distributionally robust performance along with the reformulation used in Lemma 6.1.1. However, distributionally robust state chance constraints can still be verified thanks to Lemma 6.2.

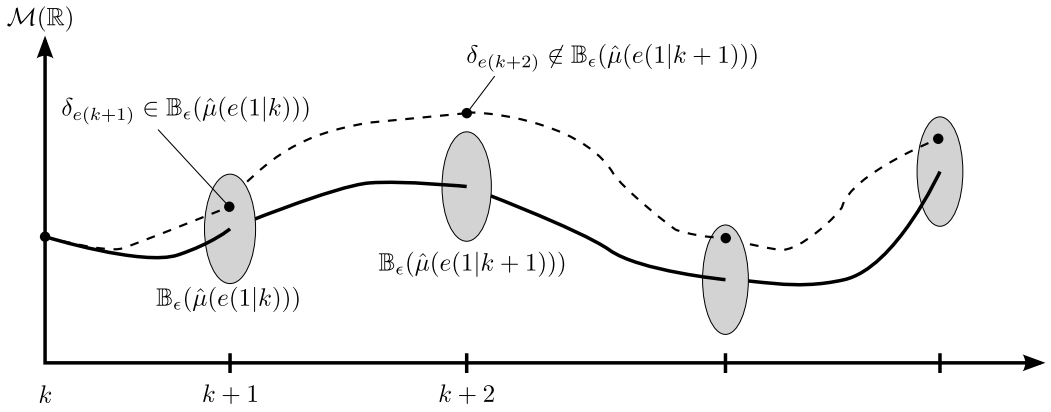

Figure 6.3: Illustration of multi-step Wasserstein balls, i.e., the gray discs centered at the empirical predictive distribution (black line), and the true predictive distribution (dotted line). At time $k+2$, the Dirac delta measure $\delta_{e(k+2)}$ (black dot) lies outside of the Wasserstein ball $\mathbb{B}_\epsilon(\hat{\mu}(e(1|k+1)))$, indicating a loss of feasibility with β probability.

6.2.4 Recursive feasibility

Following the indirect feedback paradigm, we initialize the nominal states and prediction errors with $z(0|k) = z(k)$ and $\hat{e}_j(0|k) = e(k)$ for all $j = 1, \dots, N_\mathrm{s}$. However, since at time k the error $e(k)$ is already a realization of the stochastic process with unbounded support, it is impossible to guarantee robust recursive feasibility, i.e., at time $t = 0$, the distributionally robust CVaR constraint

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(e(0|k)))} \mathrm{CVaR}^{\nu}_{p^i_\mathrm{x}}\big(h_i^\top(z(k) + e(0|k)) - 1\big) \overset{(6.14)}{=} \sup_{\nu \in \mathbb{B}_\epsilon(\delta_{e(k)})} \mathrm{CVaR}^{\nu}_{p^i_\mathrm{x}}\big(h_i^\top x(k) - 1\big) \leq 0$$

is prone to be infeasible, see also Figure 6.3 for an illustration of this issue in view of the space of probability distributions. To render the MPC optimization problem feasible for all possible realizations of $e(k)$, we resort to a soft constraint framework as proposed in [169] and introduce a vector $\Theta = [\theta_0, \dots, \theta_{N-1}]$ of so-called slack variables $\theta_t \geq 0$ for $t = 0, \dots, N-1$. These slack variables are penalized with the function $l_\Theta(\Theta) = c\|\Theta\|_1$, where $c > 0$ is a sufficiently large constant to render l_Θ as an exact penalty function, see [88, Thm. 1] for details. The softening takes into account the right hand side of the CVaR constraints in (6.18) and allows for constraint violations whenever necessary. The soft CVaR constraint set is defined as

$$\mathbb{Z}^{\mathrm{soft}} := \left\{ \begin{bmatrix} z(0|k) \\ \vdots \\ z(N-1|k) \end{bmatrix} \in \mathbb{R}^{nN} \;\middle|\; \begin{array}{l} \exists \tau_{i,t} \in \mathbb{R}, \lambda_{i,t} \in \mathbb{R}_{\geq 0}, s_{i,j,t} \in \mathbb{R}_{\geq 0}, \theta_t \in \mathbb{R}_{\geq 0} \text{ s.t.} \\ -(1 - p^i_\mathrm{x})\tau_{i,t} + \epsilon\lambda_{i,t} + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} s_{i,j,t} \leq \theta_t \\ \big(\gamma_i(z(t|k) + \hat{e}_j(t|k)) + \tau_{i,t}\big)_+ \leq s_{i,j,t} \\ \|h_i^\top\|_p \leq \lambda_{i,t} \\ \forall j \in \{1, \dots, N_\mathrm{s}\}\, \forall i \in \{1, \dots, r\} \\ \forall t \in \{0, \dots, N-1\} \end{array} \right\}.$$

In order to ensure stability of the control scheme, we impose a zero terminal constraint

$$z(N|k) = 0.$$

Remark 6.11. *In the related literature, the issue of recursive feasibility is usually not discussed. The authors of [45] circumvent the feasibility issue entirely by not constraining the initial state at time $t = 0$, which implies that no guarantees can be given. Recently, Micheli, Summers, and Lygeros [124] proposed a Wasserstein DR-MPC scheme similar to our approach, where recursive feasibility is not explicitly addressed and, similar to [45], the initial state is unconstrained. In [173], the uncertainty is assumed to be compactly supported, so that a robust constraint tightening established recursive feasibility. Finally, in [108], the authors propose a soft-constrained DR-MPC in case of bounded uncertainties.*

6.2.5 Tractable MPC optimization problem

At each time step $k = 0, \dots, N_\mathrm{T} - N$, we solve the following DR-MPC optimization problem.

Problem 6.2.1 (Wasserstein scenario-based DR-MPC)**.**

$$\begin{aligned}
\min_{\substack{z,v,\Theta \\ \tau,\lambda,s}} \quad & l_\Theta(\Theta) + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} l_1(\hat{x}_j(\cdot|k)) + l_2(\hat{u}_j(\cdot|k)) && \text{(6.21a)} \\
\text{s.t.} \quad & \hat{x}_j(t+1|k) = z(t+1|k) + \hat{e}_j(t+1|k) && \text{(6.21b)} \\
& \hat{u}_j(t|k) = v(t|k) + \pi(\hat{e}_j(t|k)) && \text{(6.21c)} \\
& \hat{e}_j(t+1|k) = A\hat{e}_j(t|k) + B\pi(\hat{e}_j(t|k)) + \hat{w}_j(k+t) && \text{(6.21d)} \\
& z(t+1|k) = Az(t|k) + Bv(t|k) && \text{(6.21e)} \\
& v(t|k) \in \mathbb{V} && \text{(6.21f)} \\
& [z(0|k), \dots, z(N-1|k)] \in \mathbb{Z}^{\mathrm{soft}} && \text{(6.21g)} \\
& z(N|k) = 0 && \text{(6.21h)} \\
& z(0|k) = z(k),\ \hat{x}_j(0|k) = x(k),\ \hat{e}_j(0|k) = e(k) && \text{(6.21i)}
\end{aligned}$$

for all $t = 0, \dots, N-1$ and for all $j = 1, \dots, N_\mathrm{s}$.

The control input applied to system (6.1) is given by

$$u(k) = v^*(0|k) + \pi(e(k)). \tag{6.22}$$

The optimization problem 6.2.1 implicitly defines a set of feasible control sequences

$$\mathcal{V}_N(z(k), \Theta) \coloneqq \left\{ v(\cdot|k), \Theta \;\middle|\; \begin{array}{l} z(0|k) = z(k) \\ z(\cdot|k) \in \mathbb{Z}^{\mathrm{soft}} \times \{0\} \\ v(\cdot|k) \in \prod_{t=0}^{N-1} \mathbb{V} \end{array} \right\},$$

Algorithm 9 Scenario-based DR-MPC

Require: Sample trajectories $\hat{\mathcal{W}}_{\text{MPC}} = \{\hat{\hat{w}}_j\}_{j=1}^{N_\text{s}} \subseteq \hat{\mathcal{W}}$ (Assumption 6.1)
1: Initialize: $z(0) = x(0)$, $e(0) = 0$ and $k = 0$
2: **repeat**
3: Measure state $x(k)$ of system (6.1)
4: Solve DR-MPC optimization problem 6.2.1 based on $\hat{\mathcal{W}}_{\text{MPC}}$ samples
5: Apply control input (6.22) to system (6.1)
6: Increment time $k \leftarrow k + 1$
7: **until** $k = N_\text{T} - N$
8: Apply remaining open-loop control input for $k = N_\text{T} - N, \dots, N_\text{T} - 1$

a set of feasible initial states $\mathbb{Z}_N^{\text{soft}} = \{z | \exists \Theta : \mathcal{V}_N(z, \Theta) \neq \emptyset\}$ and a set of strictly feasible initial states $\mathbb{Z}_N = \{z \mid \mathcal{V}_N(z, \Theta) \neq \emptyset, \|\Theta\| = 0\}$. In Algorithm 9 we summarize the closed-loop control procedure. Note that in step 8 of Algorithm 9, the open-loop control input is applied for the remaining N time steps. This can be circumvented by defining an extended task horizon $N_{\text{T,ext}} = N_\text{T} + N$, so that feedback is available up to time $k = N_\text{T}$. As a consequence, we need to gather trajectories of length $N_{\text{T,ext}}$ during the offline phase (Assumption 6.1).

6.2.6 Theoretical analysis

In the following, we provide the main results of this section.

Theorem 6.1. *Let Assumption 6.1 hold and consider system* (6.1) *under control law* (6.22) *resulting from optimization problem 6.2.1. If* $x(0) \in \mathbb{Z}_N^{\text{soft}}$*, then the Wasserstein DR-MPC optimization problem 6.2.1 is recursively feasible for all* $0 \leq k \leq N_\text{T} - N$*.*

Proof. The proof can be found in Section 6.5. □

In Theorem 6.1, we establish recursive feasibility of the DR-MPC optimization problem 6.2.1 independent of the choice of the tube controller (6.22). However, depending on the tube controller $\pi(\cdot)$ we can give different guarantees, see Figure 6.4 for a brief overview of the following results.

Corollary 6.1. *Let Assumptions 6.1 and 6.4 hold and consider system* (6.1) *under control law* (6.22) *resulting from the DR-MPC optimization problem 6.2.1. Then the resulting input* $u(k)$ *satisfies the hard input constraints* (6.3)*.*

Proof. Hard input constraint satisfaction for $u(k)$ follows immediately from control input (6.22), recursive feasibility (Theorem 6.1) and constraint tightening (Assumption 6.4), i.e., $u(k) \in \mathbb{U}$, since $\pi(e) \in \mathcal{E}_\text{u} \ \forall e \in \mathbb{R}^n$. □

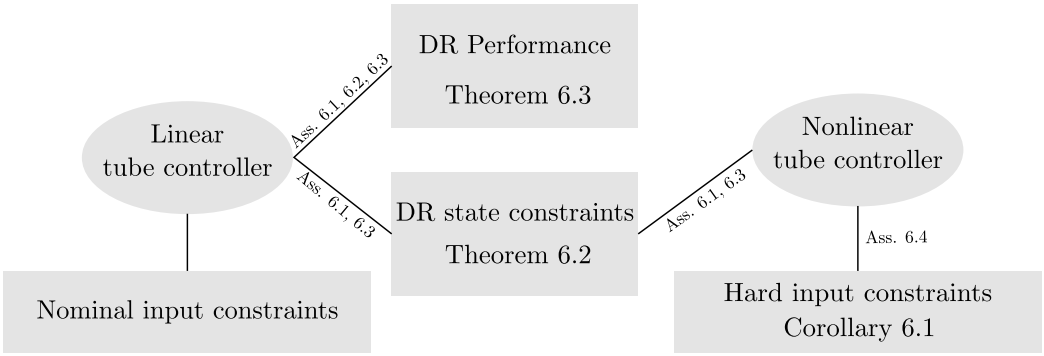

Figure 6.4: Overview of the guarantees and implications of the proposed framework.

Due to the soft constraints, we cannot guarantee satisfaction of the chance constraints (6.2) in closed-loop at all times k. However, under the assumption of strict feasibility at time k, we can state the following result.

Theorem 6.2. *Let Assumptions 6.1 and 6.3 hold and consider a strictly feasible state $x(k) \in \mathbb{Z}_N$, then the state $x(k+1)$ of system* (6.1) *resulting from control law* (6.22) *and the DR-MPC optimization problem 6.2.1 satisfies the state chance constraints* (6.2) *with a probability of at least $1-\beta$ for $\beta \in (0,1)$.*

Proof. The proof can be found in Section 6.5. □

Remark 6.12. *Theorem 6.2 establishes a conditional 1-step ahead constraint satisfaction guarantee, which is typically given in SMPC [57]. Note that the strictly feasible CVaR constraint set $\mathbb{Z}_N$ is an inner approximation of the chance constraint set [132]. Thus, even if $x(k) \in \mathbb{Z}_N^{soft}$, the chance constraints are still likely to be fulfilled.*

Theorem 6.3. *Let Assumptions 6.1, 6.2, 6.3 hold and let $\pi(e)$ be a linear map. Define the true cost $J = \mathbb{E}_{\mu_{W(k)}}(l_1(x(\cdot|k)) + l_2(u(\cdot|k)))$ and let $\hat{J}^*$ be the objective function* (6.21a) *evaluated with the optimal solution of the DR-MPC optimization problem 6.2.1. Then, with a probability of at least $1-\beta$ the following distributionally robust performance bound holds*

$$\mathbb{P}^{N_s}(J \leq \hat{J}^*) \geq 1-\beta.$$

Proof. The proof can be found in Section 6.5. □

6.2.7 Numerical example

As an example, we consider a temperature regulation task of a four room building model taken from [76] with state vector $x = [T_1, T_2, T_3, T_4]$, where T_i [°C] denotes the temperature of each room $i = 1, \ldots, 4$. The input vector u [kW] consists of the four heat flows, i.e.,

the heating/cooling power of each HVAC (heating, ventilation, and air conditioning) unit, while the ambient temperature is given by T_0. We consider the following dynamics

$$x(k+1) = Ax(k) + Bu(k) + w(k),$$

where the model parameters are taken from [76]. We model the ambient temperature as a correlated Gaussian process with $\Sigma_{ij} = 0.1 + 2\exp(-(i-j)^2/60)$ for all $i, j = 0, \ldots, N_\mathrm{T} - 1$ and sinusoidal mean $5\sin((k+6)/4) + 19$. Hence, we cannot verify distributionally robust performance, but chance constraints.

The system is subject to input constraints on the cooling/heating power $\|u\|_\infty \leq 4.5\,\mathrm{kW}$ and individual chance constraints on the room temperature

$$\begin{aligned}\mathbb{P}(x_i(k) \geq 20.4) &\geq 0.9 \quad \forall i \in \{1, \ldots, 4\}\\ \mathbb{P}(x_i(k) \leq 21.6) &\geq 0.9 \quad \forall i \in \{1, \ldots, 4\}.\end{aligned}$$

Starting from the initial condition $x(0) = [20.75 \quad 20.50 \quad 20.65 \quad 20.60]^\top$, we regulate the system to the setpoint $x_s = [21 \quad 21 \quad 21 \quad 21]^\top$ over a task horizon of $N_\mathrm{T} = 48$ hours.

Simulation setup We consider a stage cost composed of a weighted 2-norm for the states and a 1-norm for the control input

$$l_1(x) = \left\| \begin{matrix} Q^{1/2}(x(0|\cdot) - z_s) \\ \vdots \\ Q^{1/2}(x(N-1|\cdot) - z_s) \end{matrix} \right\|_2, l_2(u) = R \left\| \begin{matrix} u(0|\cdot) \\ \vdots \\ u(N-1|\cdot) \end{matrix} \right\|_1,$$

where $Q = 0.01I$ and $R = 1$. The CVaR constraints are parameterized with $p_\mathrm{x}^i = 0.9$ for all $i = 1, \ldots, 4$ and the prediction horizon is $N = 12$ hours. To obtain feedback over

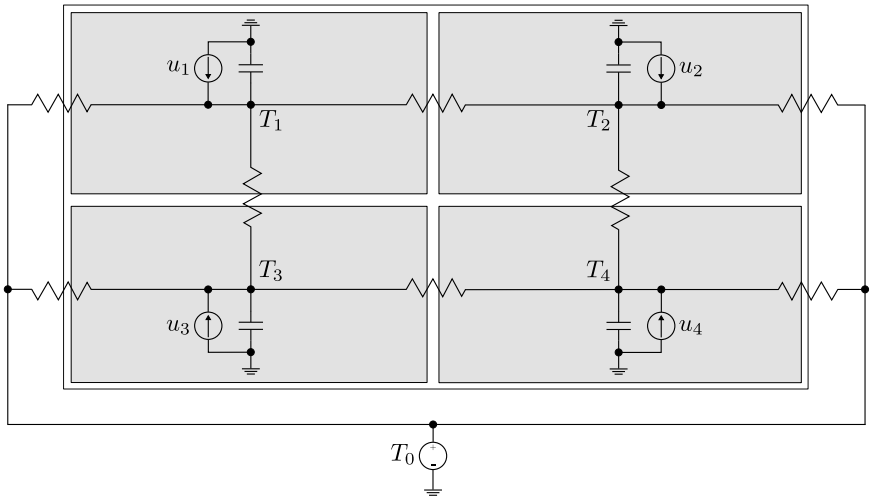

Figure 6.5: Representation of a four room building model as a resistor-capacitor network with states (temperatures) T_i and inputs (heating/cooling power) u_i.

the entire task horizon N_{T}, we implemented the controller for an extended task horizon of $N_{\mathrm{T,ext}} = N_{\mathrm{T}} + N = 60$ hours. The Wasserstein penalty is regularized with the infinity norm, while the tube controller $\pi(\cdot)$ is an LQR with weights $Q_\pi = 10^3 I$ and $R_\pi = I$ saturated at ± 1 kW. The penalty weight for the slack variables is $c = 10^6$.

Results We carried out 300 Monte-Carlo simulations of the system with different noise realizations, see Figure 6.6. It can be seen that for each disturbance realization, the hard input constraints are satisfied. In Table 6.1, for different Wasserstein radii ϵ and sample sizes N_{s}, we compare the resulting empirical worst-case constraint satisfaction (largest in-time constraint violation). It is observed that for $\epsilon = 0$ the chance constraints for $N_{\mathrm{s}} = 10$ and 20 are empirically violated, which underlines the statement that the SAA performs poorly for small sample sizes (Remark 6.2). The chance constraint satisfaction rate can be increased by either the sample size N_{s} (higher sample accuracy) or the Wasserstein radius ϵ (higher robustness to sampling errors). Moreover, the Wasserstein radius ϵ can be reduced

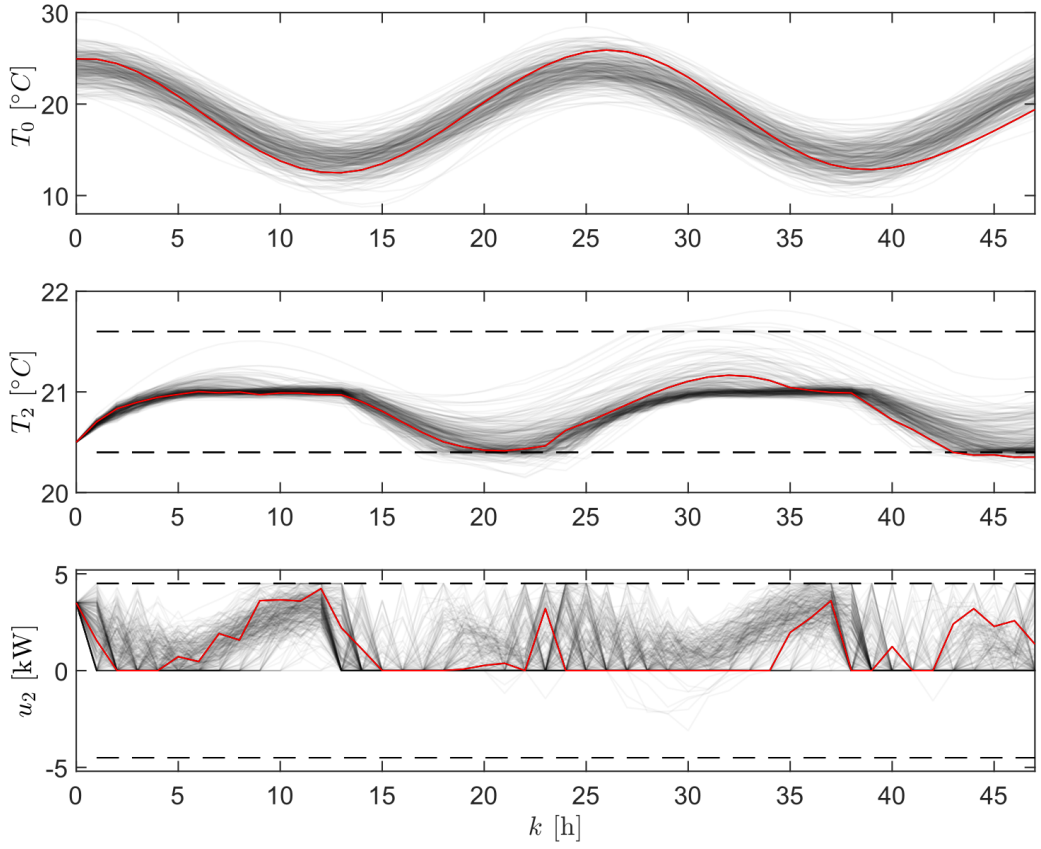

Figure 6.6: 300 realizations of the ambient temperature (Top), Room temperature (Middle) and Heating/Cooling power (Bottom). The red lines depict one particular realization.

Table 6.1: Impact of Wasserstein radius ϵ and sample size N_s on empirical satisfaction of the constraint $x_2 \geq 20.4$.

ϵ	$N_\mathrm{s} = 10$	$N_\mathrm{s} = 20$	$N_\mathrm{s} = 50$
0	87.8%	89.1%	91.2%
10^{-5}	89.5%	90.4%	92.5%
10^{-4}	91.3%	92.2%	94.6%
10^{-3}	93.1%	95.2%	97.1%

Table 6.2: Impact of N_s on average and maximum time to solve the MPC optimization problem 6.2.1.

	$N_\mathrm{s} = 10$	$N_\mathrm{s} = 20$	$N_\mathrm{s} = 30$	$N_\mathrm{s} = 40$	$N_\mathrm{s} = 50$
average	0.79 s	1.17 s	1.65 s	2.11 s	2.53 s
maximum	0.92 s	1.30 s	1.89 s	2.30 s	2.69 s

Table 6.3: Impact of N_s on strict feasibility, maximum and expected value of the slack variables conditioned on $\|\Theta\|_\infty > 0$.

$\epsilon = 1e-3$	$N_\mathrm{s} = 10$	$N_\mathrm{s} = 20$	$N_\mathrm{s} = 50$
$\mathbb{P}(\|\Theta\|_\infty = 0)$	98.69%	98.99%	99.32%
$\mathbb{E}(\|\Theta\|_\infty \mid \|\Theta\|_\infty > 0)$	$1.268 \cdot 10^{-3}$	$1.131 \cdot 10^{-3}$	$1.092 \cdot 10^{-3}$
$\max(\|\Theta\|_\infty)$	$4.871 \cdot 10^{-3}$	$4.806 \cdot 10^{-3}$	$4.637 \cdot 10^{-3}$

by increasing N_s, while preserving the prescribed level of chance constraint satisfaction (Assumption 6.3).

In Table 6.2, we show the effect of the sample size on the average and maximum computation time of the MPC optimization problem. We used CVX [67] and ran the simulation on a desktop PC with an Intel i7-9700 CPU and 16gb ram. In order to satisfy the chance constraints empirically via a SAA, we require a sample size of $N_\mathrm{s} = 50$, which is on average 3.2 times slower compared to our distributionally robust approach with $N_\mathrm{s} = 10$. Finally, we give some statistics about the slack variables θ_0 in Table 6.3. We considered the time steps between $15 \leq k \leq 25$ and $40 \leq k \leq 48$ where constraint violations can occur (cf. Figure 6.6). It can be observed that as the sample size increases, the expected value and maximum value of the slack variables decrease, while the probability of slack variables with value zero increases.

6.3 Analytical indirect feedback DR-MPC using DR-PRS

In this section, we propose an alternative form of the Wasserstein DR-MPC optimization problem 6.2.1 by decoupling the CVaR constraints entirely from the optimization problem using DR-PRS. We consider the same problem setup as in Section 6.1, but replace the constraints (6.2) and (6.3) with more general joint chance constraints of the form

$$\mathbb{P}(Hx(k) \leq h | x(0)) \geq p_{\mathrm{x}} \quad \forall k \in \{0, \dots, N_{\mathrm{T}}\} \tag{6.23a}$$

$$\mathbb{P}(Lu(k) \leq l | x(0)) \geq p_{\mathrm{u}} \quad \forall k \in \{0, \dots, N_{\mathrm{T}}\}, \tag{6.23b}$$

where $h \in \mathbb{R}^r_{>0}$, $l \in \mathbb{R}^q_{>0}$ and $r, q \in \mathbb{N}$ denote the number of individual half-spaces. The level of chance constraint satisfaction is regulated with $p_{\mathrm{x}}, p_{\mathrm{u}} \in (0, 1)$.

To immunize the chance constraints (6.23) against distributional ambiguity, we introduce the following distributionally robust surrogates

$$\inf_{\nu \in \hat{\mathcal{P}}} \mathbb{P}(x(k) \in \mathbb{X} | x(0)) \geq p_{\mathrm{x}} \tag{6.24a}$$

$$\inf_{\nu \in \hat{\mathcal{P}}} \mathbb{P}(u(k) \in \mathbb{U} | x(0)) \geq p_{\mathrm{u}}, \tag{6.24b}$$

where $\hat{\mathcal{P}} \subseteq \mathcal{P}(\Xi)$ is an ambiguity set defined on the space of all Borel probability measures $\mathcal{P}(\Xi)$, while Ξ denotes an arbitrary measurable space. In particular, we consider any data-driven ambiguity set, such that the true probability distribution $\mu \in \mathcal{P}(\Xi)$ belongs to $\hat{\mathcal{P}}$ with $1-\beta$ confidence, i.e., $\mathbb{P}^{N_{\mathrm{s}}}(\mu \in \hat{\mathcal{P}}) \geq 1-\beta$. In the following, we introduce the concept of distributionally robust PRS for arbitrary ambiguity sets $\hat{\mathcal{P}}$, while subsequently we replace $\hat{\mathcal{P}}$ with a Wasserstein ambiguity set to obtain a tractable optimization problem. Throughout the remainder of this section, we consider the predictive dynamics (6.6) with a linear tube controller $\pi(e) = Ke$, where K is a stabilizing gain for the matrix pair (A, B).

6.3.1 Constraint tightening via distributionally robust PRS

In the following, the chance constraints (6.24) are reformulated in terms of DR-PRS.

Definition 6.4. *A set $\mathcal{R}$ is a distributionally robust Probabilistic Reachable Set (DR-PRS) of probability level p w.r.t. to an ambiguity set $\hat{\mathcal{P}}$ for system* (6.5b) *with $e(0) = 0$ if*

$$\mathbb{P}(e(k) \in \mathcal{R}) \geq p \quad \forall \nu \in \hat{\mathcal{P}} \quad \forall k \in \mathbb{N}.$$

For correlated stochastic processes, it is possible to define k-step DR-PRS similar to Definition 2.7 over the task horizon N_{T}.

Definition 6.5. *A set $\mathcal{R}(k)$ for $k \in \mathbb{N}$ is a k-step distributionally robust Probabilistic Reachable Set (k-step DR-PRS) of probability level p w.r.t. to an ambiguity set $\hat{\mathcal{P}}$ for system* (6.5b) *with $e(0) = 0$ if*

$$\mathbb{P}(e(k) \in \mathcal{R}(k)) \geq p \quad \forall \nu \in \hat{\mathcal{P}}.$$

Remark 6.3.1. *If $\hat{\mathcal{P}}$ is a singleton that contains only the true distribution μ, then the original PRS (Def. 2.6) is recovered.*

For the remainder of this section, we assume that the disturbance $w(k) \sim \mu$ is zero-mean and i.i.d. for all $k \in \mathbb{N}$ and the task horizon is $N_{\mathrm{T}} = \infty$. Under this assumption, the distribution of the closed-loop error (6.5b) converges to a stationary distribution μ_e in the limit, which allows us to state the following result.

Lemma 6.3.2 (Union bound approximation)**.** *Consider system* (6.5b) *with $\pi(e(k)) = Ke$ and let $e(k) \sim \mu_e$ be an identically distributed random variable for all $k \in \mathbb{N}$. For some confidence level $\beta \in (0,1)$ and sample size $N_{\mathrm{s}} \in \mathbb{N}$, let $\hat{\mathcal{P}}$ be an ambiguity set such that $\mathbb{P}^{N_{\mathrm{s}}}(\mu_e \in \hat{\mathcal{P}}) \geq 1-\beta$ and let $p \in (0,1)$. If $\eta^* = \mathrm{col}_{i\in\{1,\dots,r\}}(\eta_i^*)$ is the optimal solution to*

$$\eta_i^* = \min_{\tilde{\eta}\in\mathbb{R}} \quad \tilde{\eta} \tag{6.25a}$$

$$\text{s.t.} \quad \inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}([H]_i e \leq \tilde{\eta}) \geq 1 - \tilde{p}_i \tag{6.25b}$$

with $\sum_{i=1}^{r} \tilde{p} \leq 1-p$, then with a probability of no less than $1-\beta$ the set $\mathcal{R} = \{e \,|\, He \leq \eta^*\}$ is a DR-PRS of level p for system (6.5b) *initialized with $e(0) = 0$.*

Proof. The proof can be found in Section 6.5. □

Corollary 6.2. *Let $\mathcal{R}_{\mathrm{x}} = \{e \mid He \leq \eta_{\mathrm{x}}^*\}$ and $\mathcal{R}_{\mathrm{u}} = \{Ke \mid LKe \leq \eta_u^*\}$ be two polytopic DR-PRS for the states and inputs. System* (6.1) *satisfies the distributionally robust chance constraints* (6.24) *for all $k \in \mathbb{N}$ conditioned on $e(0) = x(0) - z(0) = 0$ with a probability of at least $1-\beta$, if the nominal system* (6.5a) *satisfies the constraints $z(k) \in \mathbb{Z}$ and $v(k) \in \mathbb{V}$ with*

$$\mathbb{Z} \coloneqq \mathbb{X} \ominus \mathcal{R}_{\mathrm{x}}, \quad \mathbb{V} \coloneqq \mathbb{U} \ominus \mathcal{R}_{\mathrm{u}}.$$

Proof. The proof can be found in Section 6.5. □

Synthesis The main challenge in the DR-PRS synthesis problem (6.25) is the underlying distributionally robust chance constraint (6.25b). In the following, we replace the general ambiguity set $\hat{\mathcal{P}}$ with a Wasserstein ball centered at the empirical distribution $\hat{\mu}_e$. Note that other reformulations are possible as well, e.g., if $\hat{\mathcal{P}}$ is replaced with a moment-based ambiguity set, then the results from [177] can be used.

The feasible set of (6.25) is defined by the distributionally robust chance constraint (6.25b), which is in general non-convex. Thus, we additionally have to convexify the feasible set by replacing (6.25b) with a distributionally robust CVaR constraint

$$\sup_{\nu\in\mathbb{B}_\epsilon(\hat{\mu}_e)} \mathrm{CVaR}^{\nu}_{1-\tilde{p}_i}([H]_i e - \tilde{\eta}) \leq 0,$$

[5] The notation $[H]_i$ denotes the i-th row of the matrix H.

such that the convexified optimization problem is given by

$$\min_{\tilde{\eta}\in\mathbb{R}} \quad \tilde{\eta} \tag{6.26a}$$

$$\text{s.t.} \quad \sup_{\nu\in\mathbb{B}_\epsilon(\hat{\mu}_e)} \mathrm{CVaR}^{\nu}_{1-\tilde{p}_i}([H]_i e - \tilde{\eta}) \leq 0. \tag{6.26b}$$

Using Proposition 6.1.1, we can state the following result.

Theorem 6.4. *Let $e \sim \mu_e$ be an identically distributed random variable and let $\hat{\mathcal{P}} = \mathbb{B}_\epsilon(\hat{\mu}_e)$ be a Wasserstein ambiguity set (Def. 6.2) centered at $\hat{\mu}_e = N_\mathrm{s}^{-1}\sum_{j=1}^{N_\mathrm{s}} \delta_{\hat{e}_j}$ with $N_\mathrm{s} \in \mathbb{N}$, $\beta \in (0,1)$ satisfying Assumption 6.3. For the q-Wasserstein distance define the dual p-norm, such that $1/p + 1/q = 1$. Let η_i^* be the optimal solution of optimization problem*

$$\eta_i^* = \min_{\substack{\tilde{\eta}\in\mathbb{R} \\ \tau_i,\lambda,s}} \quad \tilde{\eta} \tag{6.27a}$$

$$\text{s.t.} \quad \begin{aligned} &-\tilde{p}_i\tau + \lambda\epsilon + \tfrac{1}{N_\mathrm{s}}\textstyle\sum_{j=1}^{N_\mathrm{s}} s_j \leq 0 \\ &([H]_i\hat{e}_j - \tilde{\eta} + \tau_i)_+ \leq s_j \\ &\|[H]_i\|_p \leq \lambda \quad \forall j = \{1,\dots,N_\mathrm{s}\}, \end{aligned} \tag{6.27b}$$

then η_i^* is a feasible solution to problem (6.26).

Proof. The proof can be found in Section 6.5. □

Remark 6.13. *Lemma 6.3.2 and Theorem 6.4 can now be used together to construct DR-PRS, i.e., by replacing (6.25) with (6.27). In view of this, DR-PRS with different interpretations are possible, e.g., by defining $\mathcal{R}_{\mathrm{CVaR}} = \{e|He \leq \eta^*\}$, we define the DR-PRS in the CVaR sense, whereas with $\mathcal{R} = \{e|He \leq \tau^*\}$ with $\tau^* = \mathrm{col}_{i\in\{1,\dots,r\}}(\tau_i^*)$, we restore the original DR-PRS as in Definition 6.4. Note that the roles of τ and η are switched due to the replacement of the VaR constraint with the CVaR constraint. However, the quantity $\mathcal{R}_{\mathrm{CVaR}}$ in the proposed indirect feedback case seems nonsensical, since the MPC does not actively use the properties of CVaR during decision making (as in Problem 6.2.1), i.e., we cannot penalize the magnitude of constraint violations. Nonetheless, Theorem 6.4 is important because it gives us a tractable optimization problem that yields the distributionally robust VaR τ^* as a byproduct.*

We can impose individual chance constraints of the type (6.2) by removing the condition $\sum_{i=1}^{r}\tilde{p}_i \leq 1-p$ in Lemma 6.3.2. In this case, the individual violation probabilities $\tilde{p}_i$ for all $i = 1,\dots,r$ are readily given by $\tilde{p}_i = 1-p$.

Remark 6.14. *The previous results can easily be extended to correlated disturbances over a finite task horizon via k-step DR-PRS. To do so, we impose Assumption 6.1 and pre-sample the error trajectories along*

$$\hat{e}_j(k+1) = (A+BK)\hat{e}_j(k) + \hat{w}_j(k) \quad \forall k \in \{0,\dots,N_\mathrm{T}-1\},$$

where $\hat{e}_j(0) = 0$ *for all* $j = 1, \dots, N_\mathrm{s}$. *Then we can easily compute* k*-step DR-PRS* $\mathcal{R}(k) = \{e | He(k) \leq \tau^*(k)\}$ *by solving the following optimization problem for all* $k = 0, \dots, N_\mathrm{T} - 1$ *and for all* $i = 1, \dots, r$

$$\eta_i^*(k) = \min_{\substack{\tilde{\eta}_i(k) \in \mathbb{R} \\ \tau_i(k), \lambda, s}} \quad \tilde{\eta}_i(k) \tag{6.28a}$$

$$\text{s.t.} \quad \begin{aligned} &-\tilde{p}_i \tau_i(k) + \lambda \epsilon + \tfrac{1}{N_\mathrm{s}} \textstyle\sum_{j=1}^{N_\mathrm{s}} s_j \leq 0 \\ &([H]_i \hat{e}_j(k) - \tilde{\eta}_i(k) + \tau_i(k))_+ \leq s_j \\ &\|[H]_i\|_p \leq \lambda \quad \forall j = \{1, \dots, N_\mathrm{s}\}, \end{aligned} \tag{6.28b}$$

where $\tau^*(k) = \mathrm{col}_{i \in \{1, \dots, r\}}(\tau_i^*(k))$.

6.3.2 Objective function

We consider a quadratic distributionally robust cost function

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_e)} \mathbb{E}_\nu \left(\sum_{t=0}^{N-1} \|x(t|k)\|_R^2 + \|u(t|k)\|_R^2 \right), \tag{6.29}$$

where Q, R are symmetric positive definite weighting matrices. To analytically evaluate the cost function (6.29), we use the assumption that $w(k)$ is zero-mean i.i.d., which renders the closed-loop error $e(k)$ resulting from (6.5b) with $e(0) = 0$ zero-mean for all $k \in \mathbb{N}$. However, this is generally not true for the predicted error $e(t|k)$ resulting from the initialization $e(0|k) = e(k)$. Thus, the predicted state and input mean are given by $\bar{x}(t|k) = z(t|k) + A_K^t e(k)$ and $\bar{u}(t|k) = v(t|k) + K A_K^t e(k)$, such that (6.29) evaluates to

$$\sum_{t=0}^{N-1} \left(\|\bar{x}(t|k)\|_Q^2 + \|\bar{u}(t|k)\|_R^2 \right) + \sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_e)} \mathbb{E}_\nu \left(\sum_{t=0}^{N-1} \|e(t|k)\|_{Q + K^\top R K}^2 \right). \tag{6.30}$$

It can be seen that the last term in (6.30), i.e., the variance cost, does not depend on the optimization variables z, v since K is a fixed controller gain. Thus, it can be neglected in the receding horizon cost function.

Remark 6.3.3. *Since the last term of* (6.30) *is independent of* z, v*, the out-of-sample performance of* (6.29) *cannot be improved through the MPC optimization problem. This observation is closely connected to Remark 6.5 and Lemma 6.1 in case of the scenario-based DR-MPC, where the Wasserstein regularization is independent of* x *and* u *and thus does not affect the MPC cost.*

Remark 6.3.4. *We do not penalize the terminal state* $z(N|k)$ *in the cost function* (6.29) *since we use a zero-terminal constraint* (6.31g)*. This can, however, easily be relaxed to a terminal set/controller approach, cf. Section 4.3 or [79] for related results.*

6.3.3 MPC optimization problem

Using the results from the previous section, in particular the constraint tightening from Corollary 6.2, we can state the following MPC optimization problem.

Problem 6.3.5 (Tube-based Wasserstein indirect feedback MPC)**.**

$$\min_{\mathcal{V},\mathcal{Z}} \quad \sum_{t=0}^{N-1} \left(\|\bar{x}(t|k)\|_Q^2 + \|\bar{u}(t|k)\|_R^2\right) \tag{6.31a}$$

$$\text{s.t.} \quad z(t+1|k) = Az(t|k) + Bv(t|k) \tag{6.31b}$$

$$\bar{x}(t+1|k) = A\bar{x}(t|k) + B\bar{u}(t|k) \tag{6.31c}$$

$$\bar{e}(t|k) = \bar{x}(t|k) - z(t|k) \tag{6.31d}$$

$$\bar{u}(t|k) = K\bar{e}(t|k) + v(t|k) \tag{6.31e}$$

$$(z(t|k), v(t|k)) \in \mathbb{Z} \times \mathbb{V} \tag{6.31f}$$

$$z(N|k) = 0, \tag{6.31g}$$

$$x(0|k) = x(k), z(0|k) = z(k), \tag{6.31h}$$

for all $t = 0, \ldots, N-1$. [6]

The result of the optimization problem are the optimal nominal state and control input sequences $\mathcal{Z} = \{z^*(t|k)\}_{t=0}^N$ and $\mathcal{V} = \{v^*(t|k)\}_{t=0}^{N-1}$, where only the first element $v^*(0|k)$ is implemented with control law (6.22) to system (6.1) and the rest is discarded. Recursive feasibility and closed-loop constraint satisfaction with probability $1-\beta$ follow from the same arguments as in Proposition 4.3.2 and Remark 4.8 by setting $\mathbb{Z}_{\mathrm{f}}^{\mathrm{s}} = \{0\}$.

6.3.4 Numerical example

This section is dedicated to a numerical example. To compare the results with an analytical stochastic MPC, we consider a simple double integrator with dynamics

$$x(k+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(k) + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} u(k) + w(k),$$

where w is a normally distributed zero-mean i.i.d. disturbance with the true distribution

$$w \sim \mathcal{N}(0, \Sigma_{\mathrm{w}}), \quad \Sigma_{\mathrm{w}} = \begin{bmatrix} 0.25 & 0.5 \\ 0.5 & 1 \end{bmatrix}.$$

We impose an individual chance constraint on the second state $\mathbb{P}(x_2(k) \leq 5) \geq 0.8$. Furthermore, we chose the stabilizing controller $K = [-0.2, -0.6]$. The weighting matrices are defined as $Q = \mathrm{diag}(100, 0.01)$, $R = 0.1$ and the prediction horizon is $N = 10$.

[6] Variables with a bar denote the mean values, i.e., $\bar{e}(t|k) = \mathbb{E}(e(t|k)|x(k))$.

DR-PRS guarantees

First, we quantify the impact of the Wasserstein radius on the out-of-sample performance and confidence level $1-\beta$ of the DR-PRS. To this end, we compute the true VaR of probability level $p_{\mathrm{x}} = 0.8$, which is given by $\tau^* = \mathbb{E}(e) + [\Sigma_{\mathrm{e}}]_{2,2} f(p_{\mathrm{x}})$, where Σ_{e} is the solution of the Lyapunov equation $\Sigma_{\mathrm{e}} = (A+BK)\Sigma_{\mathrm{e}}(A+BK)^\top + \Sigma_{\mathrm{w}}$ and $f(p_{\mathrm{x}})$ is the inverse cumulative distribution function (quantile function) of the standard normal distribution. Since $\mathbb{E}(e) = 0$ by design, the true VaR is given by $\tau^* = 1.2023$. Since the exact stationary error distribution is known, we can sample from it to produce i.i.d. data for the following experiment, i.e., $e_2^{(i)} \sim \mathcal{N}(0, [\Sigma_{\mathrm{e}}]_{2,2})$, where $e_2^{(i)}$ is the i-th sample of the second error state.

The Wasserstein DR-PRS is obtained using 5-fold cross validation for different training set sizes $N_{\mathrm{s}} \in \{10, 100, 400\}$, where for each N_{s} we perform $j = 1, \ldots, N_{\mathrm{mc}} = 300$ Monte-Carlo runs over independent data sets. Each run evaluates for different Wasserstein radii the out-of-sample performance w.r.t. a set of independent validation data of size $N_{\mathrm{v}} = 10^5$, i.e.,

$$p(N_{\mathrm{s}}, \epsilon) = \frac{1}{N_{\mathrm{mc}} N_{\mathrm{v}}} \sum_{j=1}^{N_{\mathrm{mc}}} \sum_{i=1}^{N_{\mathrm{v}}} \mathbb{1}\{e_2^{(i)} \leq \tau^{(j)}(N_{\mathrm{s}}, \epsilon)\}, \tag{6.32}$$

where $\tau^{(j)}(N_{\mathrm{s}}, \epsilon)$ denotes the distributionally robust VaR estimate obtained from 5-fold cross validation of problem (6.27) for the j-th training data set of size N_{s} and radius ϵ. A second indicator is the empirical confidence of the distributionally robust VaR estimate averaged over the N_{mc} Monte-Carlo runs

$$1 - \beta(N_{\mathrm{s}}, \epsilon) = \frac{1}{N_{\mathrm{mc}}} \sum_{j=1}^{N_{\mathrm{mc}}} \mathbb{1}_{\{\tau^{(j)}(N_{\mathrm{s}}, \epsilon) \geq \tau^*\}}. \tag{6.33}$$

Results In Figure 6.7, we plot the out-of-sample performance (6.32) and the corresponding confidence estimate (6.33) for different Wasserstein radii and samples sizes. For $N_{\mathrm{s}} = 10$, the confidence of the distributionally robust VaR estimates are below 90% over the entire range of ϵ, which reflects the poor sample accuracy and the necessity of an even larger Wasserstein radius. By increasing the sample size, we achieve an empirical confidence of 100% for smaller Wasserstein radii, which underlines the consistency result from Assumption 6.3. From this we can conclude that by increasing N_{s} we can reduce ϵ while maintaining a certain level of confidence $1-\beta$, which was similarly observed in Section 6.2.7.

The trade-off between sampling accuracy and robustness can be seen by comparing the out-of-sample performance p with the confidence level, i.e., to achieve a DR-PRS with $1-\beta \approx 1$ confidence and a data size $N_{\mathrm{s}} = 100$, we need a Wasserstein radius $\epsilon = 0.088$, which corresponds to the out-of-sample performance (constraint satisfaction rate) $p = 0.879$. Note that we aim to achieve $p_{\mathrm{x}} = 0.8$. If we compare the same indicators for a sample size $N_{\mathrm{s}} = 400$, it can be seen that we can reduce the Wasserstein radius to $\epsilon = 0.072$ resulting in the out-of-sample performance $p = 0.856$ with $1-\beta \approx 1$ confidence. In the limit, as $N_{\mathrm{s}} \to \infty$, we may select $\epsilon = 0$ and the optimal value of optimization problem (6.27)

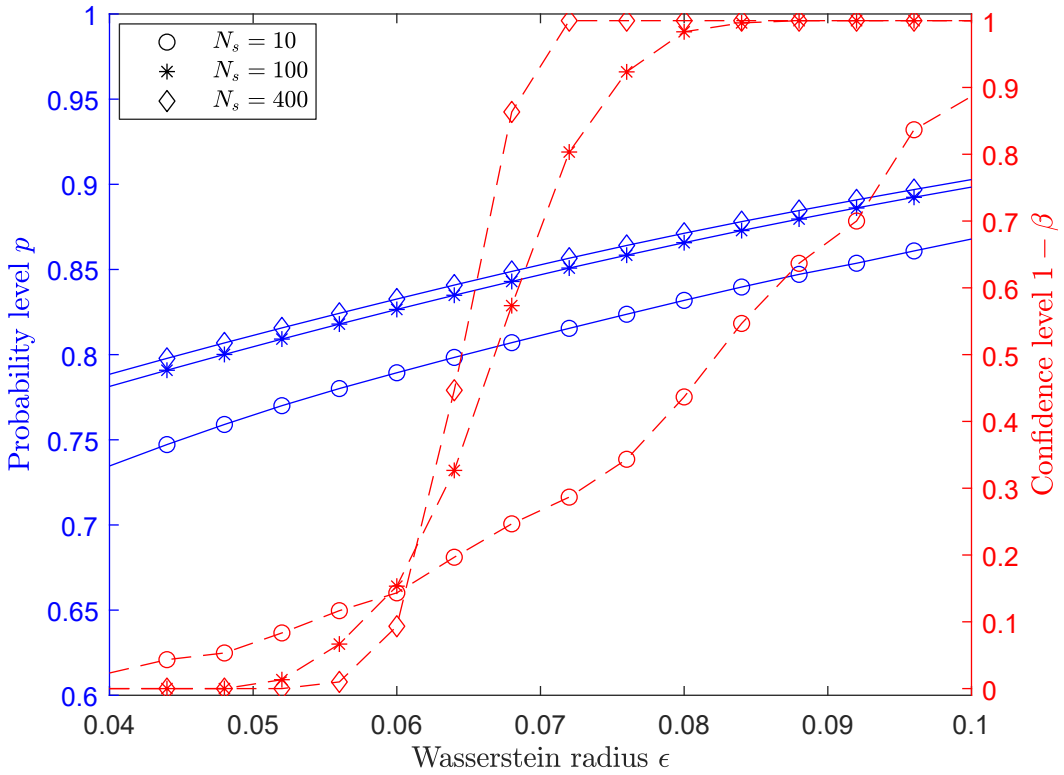

Figure 6.7: Out-of-sample performance (blue) with empirical confidence estimates (red) averaged over 300 independent training data sets.

converges with $1-\beta = 1$ confidence to the true VaR, i.e., $\lim_{N_\mathrm{s}\to\infty} \tau(N_\mathrm{s}, 0) = \tau^*$. This follows from the fact that the SAA is asymptotically consistent.

SMPC implementation

Finally, we perform for different controller configurations 10^4 Monte-Carlo simulations of the closed-loop system starting at $x(0) = [-25 \quad 0]^\top$ for 15 time steps. As performance indicators, we consider the expected cumulative closed-loop cost

$$\mathbb{E}(l(x,u)) = \frac{1}{10^4} \sum_{i=1}^{10^4} \sum_{k=0}^{15} \|x(k)\|_Q^2 + \|u(k)\|_R^2,$$

as well as the lowest in-time empirical chance constraint satisfaction. In Table 6.4, we compare the expected closed-loop cost in percentage to the linear controller $u = Kx$. As expected, the constrained controllers attain a higher transient cost compared to the unconstrained linear controller. In view of this, starting with $N_\mathrm{s} = 10$, the DR-MPC decreases the closed-loop cost by gaining more data, while for $N_\mathrm{s} \to \infty$ we obtain the SMPC with

Table 6.4: Performance comparison for different controllers and sample sizes.

		DR-MPC			
	$u = Kx$	$N_\mathrm{s} = 10$	$N_\mathrm{s} = 100$	$N_\mathrm{s} = 400$	SMPC
$\mathbb{E}(l(x,u))$	100%	132.74%	128.97%	125.58%	123.81%
$\mathbb{P}(\lvert[x_1]_2\rvert \leq 5)$	7.98%	90.38%	88.27%	85.87%	84.56%

exact knowledge of the distribution, and thus the lowest closed-loop cost of all constrained controllers. A similar trend can be seen in the chance constraint satisfaction rates, where the SMPC achieves the tightest constraint satisfaction w.r.t. the desired level of 80%, while the DR-MPC leads to slightly more conservative control actions. Hence, the controller gains confidence to operate closer to the actual chance constraint as more data is gathered. However, it is worth noting that each constrained controller verifies the chance constraint of probability level 0.8 in closed-loop.

6.4 Summary

In this chapter, we studied a class of Wasserstein distributionally robust MPC schemes for linear stochastic systems subject to additive uncertainty.

In Section 6.2, we proposed a scenario-based indirect feedback DR-MPC that uses distributionally robust CVaR constraints to robustify the chance constraints to sampling errors. In our analysis, we included correlated stochastic processes and investigated whether distributionally robust performance bounds and/or distributionally robust chance constraint guarantees can be given. The resulting MPC optimization problem is proven to be recursively feasible by considering a soft constraint formulation, while the closed-loop system verifies the chance constraints with $1 - \beta$ probability. The proposed algorithm is demonstrated on a four-room temperature control task.

In Section 6.3, the concept of distributionally robust PRS and k-step DR-PRS for general ambiguity sets is introduced, while a design for Wasserstein ambiguity sets is subsequently presented. Under a zero-mean i.i.d. assumption on the disturbance, we are able to analytically reformulate the cost function and obtain a constant constraint tightening for arbitrarily large task horizons. A numerical example is carried out to investigate the impact of the sample size and the Wasserstein radius on the VaR estimate. Additionally, we compare different controller configurations with a stochastic MPC implementation based on full distributional knowledge.

6.5 Proofs

Proof of Lemma 6.1

The proof starts by writing the predictive state equation (6.10a) together with (6.11a) in explicit form

$$x(t|k) = z(t|k) + \underbrace{A_{\mathrm{K}}^{t} e(0|k) + \sum_{i=0}^{t-1} A_{\mathrm{K}}^{t-1-i} w(i|k)}_{e(t|k)}. \tag{6.34}$$

Similarly, we can express the input equation (6.10b) together with (6.11b) as

$$u(t|k) = v(t|k) + K A_{\mathrm{K}}^{t} e(0|k) + \sum_{i=0}^{t-1} K A_{\mathrm{K}}^{t-1-i} w(i|k). \tag{6.35}$$

As it can be seen by (6.34)-(6.35), the state and input sequences are both linear in $w(\cdot|k)$, thus, we can combine the functions l_1 and l_2 to $\Phi : \mathbb{X}^N \times \mathbb{U}^N \times \mathbb{X} \times \mathbb{W}^{N-1} \to \mathbb{R}_{\geq 0}$, i.e.,

$$\Phi(W) \coloneqq l_1(w(0|k), \ldots, w(N-2|k)) + l_2(w(0|k), \ldots, w(N-2|k)).$$

To simplify the notation, we have omitted the arguments z, v, and e, since the expectation problem (6.9) considers only the linear mapping $W \mapsto \Phi$, where $W = \{w(k+t)\}_{t=0}^{N-2} \in \mathbb{W}^{N-1}$ denotes the predicted disturbance sequence. Thus, (6.9) is equal to

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_W)} \mathbb{E}_\nu(\Phi(W)). \tag{6.36}$$

By definition, the functions l_1 and l_2 are proper, convex and Lipschitz continuous, so is $\Phi(\cdot)$, since it is a sum of nonnegative convex functions [21]. Furthermore, $L_\phi > 0$ denotes the Lipschitz constant of $\Phi(\cdot)$ w.r.t. the q-norm. Since Assumption 6.1 and 6.2 hold and $\mathbb{W} = \mathbb{R}^n$, we can apply Proposition 6.1.1 with $\Xi = \mathbb{W}^{N-1}$ to (6.36) and find that

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_W)} \mathbb{E}_\nu(\Phi(W)) = \inf_{\lambda \geq 0} \lambda\epsilon + \frac{1}{N_{\mathrm{s}}} \sum_{j=1}^{N_{\mathrm{s}}} \sup_{W \in \mathbb{R}^{n(N-1)}} (\Phi(W) - \lambda \|W - \hat{W}_j\|_q), \tag{6.37}$$

where $\hat{W}_j = \{\hat{w}_j(k+t)\}_{t=0}^{N-2}$. Now we follow the proof of [126, Thm. 6.3], which, given that $\Xi = \mathbb{R}^{n(N-1)}$, establishes the equivalence

$$\begin{aligned}
&\sup_{W \in \mathbb{R}^{n(N-1)}} (\Phi(W) - \lambda \|W - \hat{W}_j\|_q) \\
&= \sup_{\theta \in \Theta} \inf_{\|\zeta\|_{q,*} \leq \lambda} \sup_{W \in \mathbb{R}^{n(N-1)}} \left((\theta + \zeta)^\top W - \Phi^*(\theta) - \zeta^\top \hat{W}_j \right) \\
&\overset{\zeta = -\theta}{=} \sup_{\theta \in \Theta} \inf_{\|\theta\|_{q,*} \leq \lambda} \left(\theta^\top \hat{W}_j - \Phi^*(\theta) \right) \\
&= \begin{cases} \Phi(\hat{W}_j) & \text{if } \sup_{\theta \in \Theta} \|\theta\|_{q,*} \leq \lambda \\ \infty & \text{otherwise,} \end{cases}
\end{aligned} \tag{6.38}$$

where the first equality uses the conjugate function $\Phi^*(\theta)$ with associated effective domain $\Theta = \{\theta \in \mathbb{R}^{n(N-1)} : \Phi^*(\theta) < \infty\}$ and the definition of the dual norm. The second equality carries out the supremum over W and the last equality uses the definition of the biconjugate function $\Phi^{**} = \Phi$ due to convexity and continuity. From [126, Prop. 6.5], we know that $\|\theta\|_{q,*} \leq L_\phi$ and from [99, Remark 3] that $\sup_{\theta\in\Theta} \|\theta\|_{q,*} = L_\phi$. Equations (6.37) - (6.38) then imply that the infimum over $\lambda \geq L_\phi$ is attained at $\lambda = L_\phi$. It remains to substitute the $\hat{W}_j$-dependent data trajectories $\hat{x}_j(\cdot|k)$ and $\hat{u}_j(\cdot|k)$ resulting from (6.34) - (6.35) together with (6.38) and $\lambda = L_\phi$ into (6.37), which, after reversing the steps of the proof yields that (6.9) is equal to

$$L_\phi \epsilon + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} \left(l_1(\hat{x}_j(\cdot|k)) + l_2(\hat{u}_j(\cdot|k))\right) = L_\phi \epsilon + \mathbb{E}_{\hat{\mu}_{W(k)}} \left(l_1(x(\cdot|k)) + l_2(u(\cdot|k))\right).$$

This concludes the proof. □

Proof of Lemma 6.2

At the beginning of the proof we neglect the time indices t, k and half-space constraint index i to simplify the notation. First, we recall the empirical predictive error distribution (6.14), i.e., $\hat{\mu}_{E(k)} = \hat{\mu}(e(0|k)) \times \ldots, \hat{\mu}(e(N|k))$. By definition of the CVaR (Def. 6.3), we can write the distributionally robust constraint (6.17) as

$$\sup_{\nu\in\mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \inf_{\tau\in\mathbb{R}} \Big(-(1-p)\tau + \mathbb{E}_\nu\big\{(\gamma(z+e)+\tau)_+\big\}\Big) \leq 0, \tag{6.39}$$

where we substituted $x = z + e$. From the min-max inequality follows that we can exchange the sup and the inf

$$\begin{aligned} &\sup_{\nu\in\mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \inf_{\tau\in\mathbb{R}} \Big(-(1-p)\tau + \mathbb{E}_\nu\big\{(\gamma(z+e)+\tau)_+\big\}\Big) \\ &\leq \inf_{\tau\in\mathbb{R}} -(1-p)\tau + \sup_{\nu\in\mathbb{B}_\epsilon(\hat{\mu}(e(t|k)))} \mathbb{E}_\nu\big((\gamma(z+e)+\tau)_+\big). \end{aligned} \tag{6.40}$$

Since $(\gamma(z+e)+\tau)_+ = \max(0, \gamma(z+e)+\tau)$ is the nonnegative pointwise maximum of an affine function, it is proper, convex and lower semicontinuous [21]. Thus, we can apply Proposition 6.1.1 to express the supremum over ν as

$$\inf_{\lambda\geq 0} \lambda\epsilon + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} \sup_{e\in\mathbb{R}^n} \left((\gamma(z+e)+\tau)_+ - \lambda\|e - \hat{e}_j\|_q\right).$$

To resolve the max-plus function, we replace $\gamma(\cdot)$ with its definition and distinguish between the following two cases. Suppose $h_i^\top(z+e)-1+\tau>0$, then

$$\begin{aligned}
&\sup_{e\in\mathbb{R}^n}\left(h^\top(z+e)-1+\tau-\lambda\|e-\hat{e}_j\|_q\right)\\
&=\sup_{e\in\mathbb{R}^n}\left(h^\top(z+e)-1+\tau-\sup_{\|\zeta_j\|_p\le\lambda}\zeta_j^\top(e-\hat{e}_j)\right)\\
&=\inf_{\|\zeta_j\|_p\le\lambda}\left(h^\top z-1+\zeta_j^\top\hat{e}_j+\tau+\sup_{e\in\mathbb{R}^n}([h^\top-\zeta_j^\top]e)\right)\\
&\overset{\zeta_j=h}{=}h^\top(z+\hat{e}_j)-1+\tau,
\end{aligned}$$

where the first equality uses the definition of the dual norm, the second equality the min max theorem [13, Prop. 5.5.4] and the third equality carries out the supremum over e, where the infimum is dropped because h is not an optimization variable. However, we still require the constraint $\|h\|_p\le\lambda$. On the other hand, if $h^\top(z+e)-1+\tau\not>0$, it is equal to 0 by definition and thus

$$\sup_{e\in\mathbb{R}^n}(-\lambda\|e-\hat{e}_j\|_q)=\inf_{\|\zeta_j\|_p\le\lambda}\sup_{e\in\mathbb{R}^n}\zeta_j^\top(\hat{e}_j-e)=0.$$

In what follows, we resort to an epigraph formulation and define for each sample j an auxiliary variable s_j, such that $s_j\ge(h^\top(z+\hat{e}_j)-1+\tau)_+$. After combining the above results we arrive at

$$(6.40)=\begin{cases}\inf\limits_{\tau\in\mathbb{R},\lambda\ge0} -(1-p)\tau+\lambda\epsilon+\dfrac{1}{N_\mathrm{s}}\displaystyle\sum_{j=1}^{N_\mathrm{s}}s_j\le0\\ \quad\text{s.t.}\quad (h^\top(z+\hat{e}_j)-1+\tau)_+\le s_j\\ \qquad\quad \|h\|_p\le\lambda\\ \qquad\quad \forall j=\{1,\dots,N_\mathrm{s}\},\end{cases}$$

which, after invoking (6.39) for all constraints $i=1,\dots,r$ for each time step $t=0,\dots,N-1$ yields

$$\left\{\begin{bmatrix}z(0|k)\\ \vdots\\ z(N-1|k)\end{bmatrix}\in\mathbb{R}^{nN}\;\middle|\;\begin{array}{l}\inf\limits_{\tau\in\mathbb{R},\lambda\ge0} -(1-p_x^i)\tau_{i,t}+\epsilon\lambda_{i,t}+\dfrac{1}{N_\mathrm{s}}\displaystyle\sum_{j=1}^{N_\mathrm{s}}s_{i,j,t}\le0\\ \left(h_i^\top\left[z(t|k)+\hat{e}_j(t|k)\right]-1+\tau_{i,t}\right)_+\le s_{i,j,t}\\ \|h_i\|_p\le\lambda_{i,t}\\ \forall j\in\{1,\dots,N_\mathrm{s}\}\,\forall i\in\{1,\dots,r\}\\ \forall t\in\{0,\dots,N-1\}\end{array}\right\}\subseteq\mathbb{X}_{\mathrm{CVaR}}.$$

Similar to [80, Prop. V.1] the infimum can be replaced with the existence of variables τ,λ,s satisfying the constraints if and only if the infimum constraint holds true. The "$\Rightarrow$" part can be split into two cases: (i) If the infimum is achieved, then the optimizer satisfies the constraints. (ii) If the infimum is not achieved, then it is $-\infty$ and the first constraint is trivially satisfied. Thus we can find variables τ,λ,s that satisfy the remaining constraints. The "$\Leftarrow$" part is obvious. This concludes the proof. □

Proof of Theorem 6.1

Assume that at time $k = 0$ a feasible solution to the MPC optimization problem 6.2.1 with $z(0) = x(0) \in \mathbb{Z}_N^{\text{soft}}$ exists, i.e., $v^*(t|k)$ for $t = 0, \ldots, N-1$ and states $z^*(t|k)$ for $t = 0, \ldots, N$.

Applying the control input (6.22) to system (6.1) results in the state $x(0|k+1) = x(k+1)$ and $z(0|k+1) = z^*(1|k)$ (initial constraint (6.21i)), for which we consider the shifted candidate sequence $v(t|k+1) = v^*(t+1|k)$ for $t = 0, \ldots, N-2$ appended with $v(N-1|k+1) = 0$. Applying the shifted input sequence to the nominal dynamics (6.21e) yields the shifted state sequence $z(t|k+1) = z^*(t+1|k)$ for $t = 0, \ldots, N-1$. At time $t = N$, we apply the terminal control input $v(N-1|k+1) = 0$, which results in $z(N|k+1) = Az(N-1|k+1) + Bv(N-1|k+1) = 0$, since $z(N-1|k+1) = z^*(N|k) = 0$ by (6.21h). Thus, $z(N|k+1)$ verifies the terminal constraint (6.21h). Since $v(t|k+1) \in \mathbb{V}$ for $t = 0, \ldots, N-2$ and $v(N-1|k+1) = 0$, we have that $v(\cdot|k+1)$ verifies the input constraints (6.21f).

Note that the state constraints (6.21g) are relaxed as soft constraints. In view of this, we can always find slack variables θ_t for $t = 0, \ldots, N-1$, such that $[z(0|k+1), \ldots, z(N-1|k+1)] \in \mathbb{Z}^{\text{soft}}$, i.e., (6.21g) is verified at time $k+1$. This concludes the proof. □

Proof of Theorem 6.2

Let $\mu_{E(k)}$ be the true conditional error distribution (6.8) of $E(k) = \{e(k), e(1|k), \ldots, e(N-1|k)\}$ and suppose that $\mu_{E(k)} = \mu(e(k)) \times \mu(e(1|k)) \times \ldots \times \mu(e(N-1|k))$ is a product of unknown marginal distributions. Now we recall the empirical distribution (6.14), i.e., $\hat{\mu}_{E(k)} = \delta_{e(k)} \times \hat{\mu}(e(1|k)) \times \ldots, \hat{\mu}(e(N-1|k))$ and define

$$\tau_i(1|k) \coloneqq \sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}(e(1|k)))} \text{CVaR}_{p_x^i}^{\nu}\big(\gamma_i(\underbrace{z^*(1|k) + e(1|k)}_{=\hat{x}^*(1|k)})\big)$$

$$\tau_i(k+1) \coloneqq \text{CVaR}_{p_x^i}^{\mu(e(k+1))}\big(\gamma_i(\underbrace{z^*(1|k) + e(k+1)}_{\overset{(6.21i)}{=} x(k+1)})\big),$$

where $\tau_i(1|k)$ is the 1-step ahead predicted distributionally robust CVaR computed over the Wasserstein ambiguity set centered at $\hat{\mu}(e(1|k))$ and $\tau_i(k+1)$ denotes the CVaR based on the true distribution $\mu(e(k+1))$. By invoking Assumption 6.1 and 6.3, we have

$$\mathbb{P}^{N_s}\left(\mu_{W(k)} \in \mathbb{B}_\epsilon(\hat{\mu}_{W(k)}\right) \geq 1 - \beta.$$

This guarantee directly carries over to the error sequence $E(k)$ by noting that $\mu_{E(k)}$ is fundamentally obtained from (6.7) along with (6.5b), which preserves the light-tailedness of the underlying distribution. This can be verified by explicitly writing the density functions for each term in (6.5b), i.e., for linear tube controllers we have $e(k+1) = (A+BK)e(k)+w(k)$, which essentially adds two light-tailed distributions (convolution integral), since $e(0) = 0$ and $e(1) = w(0)$, whereas for saturated tube controllers, the closed-loop error is governed

by $e(k+1) = Ae(k) + \tilde{e}_u(k) + w(k)$, where the input error $\tilde{e}_u(k) = B\pi(e(k))$ forms an additional bounded random variable, which is light-tailed by definition. Thus, we convolve three light-tailed distributions, which remain light-tailed. In view of this, we have

$$\mathbb{P}^{N_\mathrm{s}}\left(\mu(e(k+1)) \in \mathbb{B}_\epsilon(\hat{\mu}(e(1|k)))\right) \geq 1-\beta,$$

which implies that the true CVaR $\tau_i(k+1)$ can be bounded as

$$\mathbb{P}^{N_\mathrm{s}}\big(\tau_i(k+1) \leq \tau_i(1|k)\big) \geq 1-\beta. \tag{6.41}$$

Now, consider that whenever $x(k) \in \mathbb{Z}_N$, the slack variables satisfy $\|\Theta\| = 0$, which yields that (6.21g) verifies the hard constraints (6.18), i.e.,

$$[z^*(0|k), \ldots, z^*(N-1|k)] \in \mathbb{Z}.$$

This implies that $\tau_i(1|k) \leq 0$ and thus, $\tau_i(k+1)$ in (6.41) is bounded from above by 0 with $1-\beta$ probability for all constraints $i = 1, \ldots, r$ w.r.t. the true measure $\mathbb{P}$. The claim follows by pointing out that the CVaR majorizes the VaR and forms a sufficient condition for chance constraint satisfaction in (6.2), cf. Remark 6.7. □

Remark 6.15. *The inequality* (6.41) *can be explained with Figure 6.3, where whenever the Dirac distribution* $\mu(e(k)) = \delta_{e(k)}$ *is contained in the ambiguity set, the bound* (6.41) *holds true. For example, at time step* $k+1$ *the bound is verified, while at time* $k+2$ *the Dirac distribution is not contained in the ambiguity set, resulting in* $\tau_i(k+1) > \tau_i(1|k)$.

Proof of Theorem 6.3

From Assumption 6.3 we have that $\mathbb{P}^{N_\mathrm{s}}(\mu_{W(k)} \in \mathbb{B}_\epsilon(\hat{\mu}_{W(k)})) \geq 1-\beta$. Hence, J is upper bounded by the distributionally robust cost function (6.9) with confidence $1-\beta$. Using Assumptions 6.1, 6.2 and linearity of $\pi(e)$ allows for application of Lemma 6.1, which establishes an equivalence between (6.9) and the SAA cost function (6.12), i.e.,

$$J \leq L_\phi \epsilon + \mathbb{E}_{\hat{\mu}_{W(k)}}\left(l_1(x(\cdot|k)) + l_2(u(\cdot|k))\right). \tag{6.42}$$

Since the MPC cost function (6.21a) is equivalent to the the right hand side of (6.42) plus an additional non-negative penalty term l_Θ, and the expectation problem is solved w.r.t. the disturbance sequence $W(k)$, the upper bound (6.42) holds for all feasible solutions of the DR-MPC optimization problem 6.2.1 with $1-\beta$ probability. The claim follows by choosing the optimal solution, which yields $\hat{J}^*$. □

Proof of Lemma 6.3.2

We start the proof by recalling the definition of a DR-PRS (Def. 6.4) and an equivalence result from Zymler et al. [177], which states that the condition $\mathbb{P}(e(k) \in \mathcal{R}) \geq p \quad \forall \nu \in \hat{\mathcal{P}}$

can similarly be written as

$$\inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(e(k)\in\mathcal{R}) \geq p \Longleftrightarrow \sup_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(e(k)\notin\mathcal{R}) \leq 1-p. \tag{6.43}$$

Now, consider the polytopic distributionally robust JCC $\inf_{\mu\in\hat{\mathcal{P}}} \mathbb{P}(He\leq\eta)\geq p$, which can equivalently be written as an intersections of ICCs

$$\inf_{\mu\in\hat{\mathcal{P}}} \mathbb{P}\Big(\bigcap_{i=1}^{r}\big([H]_i e(k)\leq\eta_i\big)\Big)\geq p \overset{(6.43)}{\Longleftrightarrow} \sup_{\mu\in\hat{\mathcal{P}}} \mathbb{P}\Big(\bigcup_{i=1}^{r}\big([H]_i e(k)\geq\eta_i\big)\Big)\leq 1-p,$$

where $[H]_i$ denotes the i-th row of H and η_i the i-th element of η. Applying the union bound yields

$$\sup_{\mu\in\hat{\mathcal{P}}} \mathbb{P}\Big(\bigcup_{i=1}^{r}\big([H]_i e(k)\geq\eta_i\big)\Big)\leq \sum_{i=1}^{r}\sup_{\mu\in\hat{\mathcal{P}}} \mathbb{P}([H]_i e(k)\geq\eta_i)\leq \sum_{i=1}^{r}\tilde{p}_i,$$

where the individual violation probabilities $\tilde{p}_i\in(0,1)$ satisfy the condition $\sum_{i=1}^{r}\tilde{p}_i\leq 1-p$. Hence, we can approximate the distributionally robust JCC with a set of r distributionally robust ICCs. By forming the complementary probability, we have

$$\sup_{\mu\in\hat{\mathcal{P}}} \mathbb{P}\big([H]_i e(k)\geq\eta_i\big)\leq\tilde{p}_i \Longleftrightarrow \inf_{\mu\in\hat{\mathcal{P}}} \mathbb{P}\big([H]_i e(k)\leq\eta_i\big)\geq 1-\tilde{p}_i.$$

The values η_i for all $i=1,\ldots,r$ that verify the distributionally robust ICCs are the so-called distributionally robust Value-at-Risk, which can be obtained as the optimal solution η_i^* of problem (6.25) for $i=1,\ldots,r$. By stacking $\eta^*=\mathrm{col}_{i\in\{1,\ldots,r\}}(\eta_i^*)$, we can thus define $\mathcal{R}=\{e\mid He\leq\eta^*\}$. Since the ambiguity set is defined based on N_s samples of $e(k)\sim\mu_e$, the guarantee $\mathbb{P}^{N_\mathrm{s}}(\mu\in\hat{\mathcal{P}})\geq 1-\beta$ implies that $\mathcal{R}$ is a DR-PRS of probability level $1-\beta$ for (6.5b) initialized with $e(0)=0$. □

Proof of Corollary 6.2

Consider the state constraints (6.24) in polytopic form $Hx(k)\leq h$, which can equivalently be written as

$$\inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(Hx(k)\leq h|x(0)) = \inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(Hz(k)\leq h-He(k)|x(0))\geq p_x,$$

where we substituted $x=z+e$. This is equivalent to

$$\exists\tilde{\eta}_x\in\mathbb{R}^r \;:\; Hz(k)\leq h-\tilde{\eta}_x \quad\text{and}\quad \inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(He(k)\leq\tilde{\eta}_x|x(0))\geq p_x.$$

Now, set $\tilde{\eta}_x=\eta_x^*$, where η_x^* is the optimal distributionally robust Value-at-Risk (Lemma 6.3.2). Then we have that $\inf_{\nu\in\hat{\mathcal{P}}} \mathbb{P}(He(k)\leq\eta_x^*|x(0))\geq p_x$ is verified and only the first condition needs to be evaluated, i.e.,

$$Hz(k)\leq h-\eta_x^* \Longleftrightarrow z(k)\in\mathbb{Z}=\mathbb{X}\ominus\mathcal{R},$$

where $\mathcal{R}$ is the DR-PRS of level p_x associated with η_x^*. Since $\mathcal{R}$ is only a DR-PRS with $1-\beta$ confidence, we have that the original chance constraints are verified with a probability no worse than $1-\beta$. This concludes the proof for the state chance constraints (6.24a). The input chance constraints are similarly derived starting from $Lu(k) \leq l$ with separation $Lv(k) + LKe(k) \leq l$ and input DR-PRS $\mathcal{R}_u$. □

Proof of Theorem 6.4

Consider constraint (6.26b), which substituted into the definition of the distributionally robust CVaR leads to

$$\begin{aligned} &\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_e)} \inf_{\tau \in \mathbb{R}} \big(-\tilde{p}_i \tau + \mathbb{E}_\nu \left(([H]_i e - \tilde{\eta} + \tau)_+ \right) \big) \\ \leq &\inf_{\tau \in \mathbb{R}} \big(-\tilde{p}_i \tau + \sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_e)} \mathbb{E}_\nu \left([H]_i e - \tilde{\eta} + \tau)_+ \right) \big), \end{aligned} \tag{6.44}$$

where the inequality follows from the min max inequality. Now we follow the proof of Lemma 6.2 to conclude that

$$\sup_{\nu \in \mathbb{B}_\epsilon(\hat{\mu}_e)} \inf_{\tau \in \mathbb{R}} \big(-\tilde{p}_i \tau + \mathbb{E}_\nu \left(([H]_i \hat{e}_j - \tilde{\eta} + \tau)_+ \right) \big) \leq \begin{cases} \exists \tau \in \mathbb{R}, \lambda \in \mathbb{R}_{\geq 0}, s_j \in \mathbb{R}_{\geq 0} \text{ s.t.} \\ -\tilde{p}_i \tau + \lambda \epsilon + \frac{1}{N_\mathrm{s}} \sum_{j=1}^{N_\mathrm{s}} s_j \leq 0 \\ ([H]_i \hat{e}_j - \tilde{\eta} + \tau)_+ \leq s_j \\ \|[H]_i\|_p \leq \lambda, \quad \forall j = \{1, \ldots, N_\mathrm{s}\}, \end{cases}$$

i.e., the distributionally robust CVaR constraint (6.26b) is conservatively verified. Now, substitute the latter inequality into (6.26b), which results in optimization problem (6.27). Due to the min max inequality (6.44), the optimal solution η_i^* is an inner approximation of the original CVaR constraint (6.26b), which renders η_i^* feasible in (6.26). □

7 Moment-based Distributionally Robust Model Predictive Control

In the previous chapter, we have introduced two distributionally robust MPC schemes with Wasserstein ambiguity sets. In the following, we propose a DR-MPC for constrained linear systems with additive sub-Gaussian i.i.d. noise, where we introduce data-driven moment-based ambiguity in Section 7.1.1. The additional sub-Gaussianity assumption allows us to derive an explicit number of samples to ensure a user-defined confidence level, such that the true distribution belongs the ambiguity set with probability $1 - \beta$. In contrast, for Wasserstein ambiguity sets this is not possible, since the radius ϵ still depends on some unknown parameters of the true distribution. Therefore, we can only estimate the data-driven Wasserstein radius empirically using machine learning tools, while in the case of moment-based ambiguity sets we can provide more rigorous results.

In Section 7.2, we introduce a simplified affine disturbance feedback policy, which allows us to reformulate the distributionally robust chance constraints as second-order cone constraints, while the cost can be analytically reformulated. The main results of this chapter are given in Section 7.3, while numerical examples are carried out in Section 7.4.

For related work on distributionally robust MPC, please refer to Section 1.2.2. This chapter is based on the publication [117] [1].

7.1 Problem description

In this chapter, we consider a discrete linear time-invariant system

$$x(k+1) = Ax(k) + Bu(k) + Ew(k) \tag{7.1}$$

with the state $x \in \mathbb{X} \subseteq \mathbb{R}^{n_\mathrm{x}}$, input $u \in \mathbb{U} \subseteq \mathbb{R}^{n_\mathrm{u}}$, disturbance $w \in \mathbb{W} \subseteq \mathbb{R}^{n_\mathrm{w}}$ and matrices A, B and E of conformal dimensions.

We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ for an infinite sequence $\bar{w} : \Omega \to \mathbb{W}^\infty$ of random variables $w : \Omega \to \mathbb{W}$, i.e., $\bar{w} = \{w(k)\}_{k=0}^\infty$, where $w(k) \sim \mu^*$ is assumed to be zero-mean i.i.d. for all $k \in \mathbb{N}$. We assume that we have access to $i = 1, \ldots, N_\mathrm{s}$ random samples $\hat{w}^i \sim \mu^*$.

[1] C. Mark and S. Liu. "Recursively Feasible Data-Driven Distributionally Robust Model Predictive Control With Additive Disturbances". In: *IEEE Control Systems Letters* 7 (2023), pp. 526–531 ©2022 IEEE.

We further assume that the matrix pair (A, B) is stabilizable, the matrix E has full column rank and that perfect state measurement is available at each time instant k.

The system dynamics is subject to $n_\mathrm{r} \in \mathbb{N}$ and $n_\mathrm{s} \in \mathbb{N}$ individual state and input chance constraints

$$\mathbb{P}(h_r^\top x(k) \leq 1) \geq p_\mathrm{r}^\mathrm{x} \quad \forall k \in \mathbb{N} \quad \forall r \in \{1, \ldots, n_\mathrm{r}\} \tag{7.2a}$$
$$\mathbb{P}(l_s^\top u(k) \leq 1) \geq p_s^\mathrm{u} \quad \forall k \in \mathbb{N} \quad \forall s \in \{1, \ldots, n_\mathrm{s}\}, \tag{7.2b}$$

where $p_\mathrm{r}^\mathrm{x}, p_\mathrm{s}^u \in (0, 1)$ denote the required levels of chance constraint satisfaction.

Predictive dynamics To distinguish between closed-loop and predicted states, we introduce the predictive dynamics over a time horizon of length $N \in \mathbb{N}$

$$\bar{x}_k = \bar{A} x_{0|k} + \bar{B} \bar{u}_k + \bar{E} \bar{w}_k, \tag{7.3}$$

where $\bar{x}_k = [x_{0|k}^\top, x_{1|k}^\top, \ldots, x_{N|k}^\top]^\top \in \mathbb{R}^{(N+1)n_\mathrm{x}}$ denotes the predicted state sequence, $\bar{u}_k = [u_{0|k}^\top, \ldots, u_{N-1|k}^\top]^\top \in \mathbb{R}^{Nn_\mathrm{u}}$ the input sequence and $\bar{w}_k = [w_{0|k}^\top, \ldots, w_{N-1|k}^\top]^\top \in \mathbb{R}^{Nn_\mathrm{w}}$ the disturbance sequence. The matrices $\bar{A} \in \mathbb{R}^{(N+1)n_\mathrm{x} \times n_\mathrm{x}}$, $\bar{B} \in \mathbb{R}^{(N+1)n_\mathrm{x} \times Nn_\mathrm{u}}$ and $\bar{E} \in \mathbb{R}^{(N+1)n_\mathrm{x} \times Nn_\mathrm{w}}$ are defined as

$$\bar{A} \coloneqq \begin{bmatrix} I \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \bar{B} \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ B & 0 & \ldots & 0 \\ AB & B & \ldots & 0 \\ \vdots & \ddots & \ddots & 0 \\ A^{N-1}B & \ldots & AB & B \end{bmatrix}, \bar{E} \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ E & 0 & \ldots & 0 \\ AE & E & \ldots & 0 \\ \vdots & \ddots & \ddots & 0 \\ A^{N-1}E & \ldots & AE & E \end{bmatrix}.$$

Predictive chance constraints To cope with the vector-valued predictions (7.3), we lift the half-space matrices h_r and h_s from (7.2) into the dimension of the vectors $\bar{x}_k$ and $\bar{u}_k$ and define the individual chance constraints on the predicted states and inputs as

$$\mathbb{P}(h_{t,r}^\top \bar{x}_k \leq 1 \mid x(k)) \geq p_\mathrm{r}^\mathrm{x} \quad \forall t \in \{0, \ldots, N-1\} \tag{7.4a}$$
$$\mathbb{P}(l_{t,s}^\top \bar{u}_k \leq 1 \mid x(k)) \geq p_s^\mathrm{u} \quad \forall t \in \{0, \ldots, N-1\}, \tag{7.4b}$$

where $h_{t,r} \in \mathbb{R}^{(N+1)n_\mathrm{x}}$ and $l_{t,s} \in \mathbb{R}^{Nn_\mathrm{u}}$ describe the $r = 1, \ldots, n_\mathrm{r}$ state and $s = 1, \ldots, n_\mathrm{s}$ input chance constraints at time t. Given an initial value $x_{0|k}$, we opt to solve the following finite horizon stochastic optimal control problem

$$\min_{\bar{u}_k} \quad \mathbb{E}_{\mu^*}\left(\|x_{N|k}\|_P^2 + \sum_{t=0}^{N-1} \|x_{t|k}\|_Q^2 + \|u_{t|k}\|_R^2 \middle| x(k) \right) \tag{7.5a}$$
$$\text{s.t.} \quad \bar{x}_k = \bar{A} x_{0|k} + \bar{B} \bar{u}_k + \bar{E} \bar{w}_k$$
$$\mathbb{P}(h_{t,r}^\top \bar{x}_k \leq 1 \mid x(k)) \geq p_\mathrm{r}^\mathrm{x} \quad \forall t \in \{0, \ldots, N-1\} \quad \forall r \in \{1, \ldots, n_\mathrm{r}\} \tag{7.5b}$$
$$\mathbb{P}(l_{t,s}^\top \bar{u}_k \leq 1 \mid x(k)) \geq p_s^\mathrm{u} \quad \forall t \in \{0, \ldots, N-1\} \quad \forall s \in \{1, \ldots, n_\mathrm{s}\}, \tag{7.5c}$$

where $Q \succeq 0$, $R \succ 0$ and $P \succ 0$ are positive definite symmetric weighting matrices and P additionally satisfies the Lyapunov inequality

$$(A + BK)^\top P(A + BK) + Q + K^\top RK \preceq P \tag{7.6}$$

for some linear controller matrix $K \in \mathbb{R}^{n_\mathrm{u} \times n_\mathrm{x}}$.

Problem (7.5) represents an infinite-dimensional optimization problem due to the control input u and the additive disturbance w. This issue will be tackled in Section 7.2 by using a simplified affine disturbance feedback (SADF) parameterization. Furthermore, the cost function (7.5a) and chance constraints (7.5b)-(7.5c) are evaluated w.r.t. the true, but unknown distribution μ^*. Thus, we instead pose a distributionally robust optimization problem that uses a moment-based ambiguity set $\mathcal{P}$, where each distribution $\mu \in \mathcal{P}$ lies within some distance to the sample covariance $\hat{\Sigma} = N_\mathrm{s}^{-1} \sum_{i=1}^{N_\mathrm{s}} \hat{w}^i (\hat{w}^i)^\top$ under the assumption that $\mathbb{E}_\mu(w) = 0$. In particular, the ambiguity set represents the uncertainty of the empirical estimator and is parameterized as proposed by Delage and Ye [49], i.e.,

$$\mathcal{P} \coloneqq \left\{ \mu \in \mathcal{M}(\mathbb{W}) \;\middle|\; \begin{array}{l} \mathbb{E}_\mu(w) = 0 \\ \mathbb{E}_\mu(ww^\top) \preceq \kappa_\beta \hat{\Sigma} \end{array} \right\}, \tag{7.7}$$

where $\mathcal{M}(\mathbb{W})$ denotes the set of all Borel probability measures with finite variance defined on $(\mathbb{W}, \mathcal{B}(\mathbb{W}))$. For some confidence level $\beta \in (0, 1)$, we define a constant $\kappa_\beta \geq 1$, such that $\mathbb{P}(\mu^* \in \mathcal{P}) \geq 1 - \beta$.

Remark 7.1. *Moment-based ambiguity sets have two important advantages over Wasserstein ambiguity sets. First, the resulting optimization problem does not increase in complexity with the sample size and second, the ambiguity radius is estimated with reasonable accuracy based on known information of the distribution via concentration inequalities. Therefore, we can derive a controller that is several orders of magnitude faster than the Wasserstein scenario-based DR-MPC presented in Section 6.2, while the resulting guarantees are theoretically stronger compared to the analytical Wasserstein DR-MPC from Section 6.3.*

7.1.1 Data-driven ambiguity set

In the following, we derive an explicit value for the constant κ_β under the assumption of sub-Gaussianity of the random variables $w(k)$, which was similarly done by Coppens, Schuurmans, and Patrinos [44]. This extends the results from Delage and Ye [49], who provide an explicit value κ_β for bounded random variables.

Definition 7.1 ([167])**.** *A random variable ξ is sub-Gaussian with variance proxy σ^2 if $\mathbb{E}(\xi) = 0$ and its moment generating function satisfies*

$$\mathbb{E}(e^{\lambda \xi}) \leq e^{\frac{\sigma^2 \lambda^2}{2}} \quad \forall \lambda \in \mathbb{R}.$$

We denote this by $\xi \sim subG(\sigma^2)$.

Proposition 7.1.1. *Let $w \in \mathbb{W} \subseteq \mathbb{R}^{n_\mathrm{w}}$ be a zero-mean sub-Gaussian random variable with $\mathbb{E}(ww^\top) = \Sigma$. Let $\{\hat{w}^i\}_{i=1}^{N_\mathrm{s}}$ be N_s i.i.d. samples obtained from the true distribution of w and define $\hat{\Sigma} = N_\mathrm{s}^{-1} \sum_{i=1}^{N_\mathrm{s}} \hat{w}_i (\hat{w}^i)^\top$ as the empirical covariance matrix. Let $\epsilon \in (0, 0.5)$, $\beta \in (0,1)$, $c_1(\sigma, \epsilon) = \sigma^2/(1-2\epsilon)$, $c_2(\beta, \epsilon, n_\mathrm{w}) = n_\mathrm{w} \log(1 + 2/\epsilon) + \log(2/\beta)$, then for all $N_\mathrm{s} \in \mathbb{N}$ satisfying*

$$N_\mathrm{s} \geq \left\lceil 2c_1 c_2 \left(8c_1 + 4\sqrt{4c_1^2 + c_1} + 1 \right) \right\rceil, \tag{7.8}$$

the covariance bound $\Sigma \preceq \frac{1}{1-\gamma(N_\mathrm{s}, \beta/2)} \hat{\Sigma}$ holds with a probability of at least $1 - \beta$, where

$$\gamma(N_\mathrm{s}, \beta) \coloneqq c_1(\sigma, \epsilon) \left(\sqrt{\frac{32 c_2(\beta, \epsilon, n_\mathrm{w})}{N_\mathrm{s}}} + \frac{2 c_2(\beta, \epsilon, n_\mathrm{w})}{N_\mathrm{s}} \right).$$

Proof. The proof can be found in Section 7.6. □

For a fixed $\beta \in (0,1)$, the mapping from $\epsilon \mapsto N_\mathrm{s}$ in condition (7.8) is convex on the interval $\epsilon \in (0, 0.5)$, cf. Figure 7.1. Thus, to obtain the smallest number N_s satisfying (7.8), we solve a nonlinear (convex) optimization problem

$$\epsilon^* = \underset{\epsilon \in (0, 0.5)}{\operatorname{argmin}} \quad 2c_1 c_2 \left(8c_1 + 4\sqrt{4c_1^2 + c_1} + 1 \right). \tag{7.9}$$

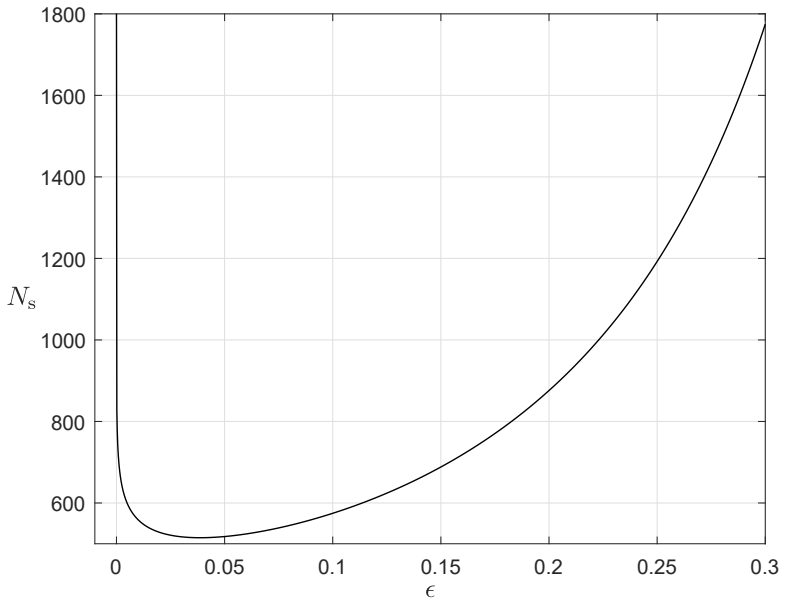

Figure 7.1: Number of samples N_s as a function of ϵ for $n_\mathrm{w} = 2$ and $\beta = 0.05$ on the interval $\epsilon \in [10^{-8}, 0.3]$.

Finally, by setting $\kappa_\beta = 1/(1-\gamma(N_\mathrm{s}, \beta/2))$, we obtain from optimization problem (7.9) and Proposition 7.1.1 an explicit number of samples N_s to ensure that the ambiguity set (7.7) contains the true distribution μ^* with $1-\beta$ confidence.

Corollary 7.1. *Select $\beta \in (0,1)$, $\epsilon \in (0, 0.5)$ and $\tilde{N}_\mathrm{s} \in \mathbb{N}$, such that (7.8) holds. Then, for $N_\mathrm{s} \geq \tilde{N}_\mathrm{s}$ the function κ_β is monotonically decreasing and $\kappa_\beta \to 1$ as $N_\mathrm{s} \to \infty$.*

Proof. By definition of $\gamma(N_\mathrm{s}, \beta)$ in Proposition 7.1.1, we have that $\gamma(N_\mathrm{s}, \beta) < 1$ for $N_\mathrm{s} \geq \tilde{N}_\mathrm{s}$ and that $\gamma(N_\mathrm{s}, \beta) \to 0$ of order $\mathcal{O}(1/\sqrt{N_s})$. Consequently, $1 > \gamma(N_\mathrm{s}, \beta) > \gamma(N_\mathrm{s}+i, \beta) \geq 0$ for all $i \in \mathbb{N}$. For $N_\mathrm{s} = \tilde{N}_\mathrm{s}$ the function $\kappa_\beta = 1/(1-\gamma(\tilde{N}_\mathrm{s}, \beta)) \gg 1$, which in view of $\gamma(N_\mathrm{s}, \beta)$ implies that κ_β monotonically decreases to $\kappa_\beta = 1$ as $N_\mathrm{s} \to \infty$. □

Remark 7.2. *The result of Proposition 7.1.1 is a special case of [44, Thm. 8] with known first moment information. Consequently, we require fewer samples compared to the mean and variance ambiguity set proposed by [44] to achieve the $1-\beta$ confidence in (7.7).*

7.2 Controller design

We resort to a SADF parameterization [171] of the form

$$u_{t|k} = v_{t|k} + \sum_{i=0}^{t-1} M_{t-i|k} w_{i|k}, \tag{7.10}$$

where $v_{i|k} \in \mathbb{R}^{n_\mathrm{u}}$ is the predicted control input and $M_{i-t|k} \in \mathbb{R}^{n_\mathrm{u} \times n_\mathrm{x}}$ are feedback matrices, both of which are decision variables in the resulting MPC optimization problem. Thus, $u_{i|k}$ depends affinely on the past i disturbance $w_{0|k}, \ldots, w_{i-1|k}$. To streamline the presentation, we consider the matrix $\bar{M}_k \in \mathbb{R}^{Nn_\mathrm{u} \times Nn_\mathrm{w}}$ and the vector $\bar{v}_k \in \mathbb{R}^{Nn_\mathrm{u}}$

$$\bar{M}_k \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ M_{1|k} & 0 & \ldots & 0 \\ \vdots & \ddots & \ddots & 0 \\ M_{N-1|k} & \ldots & M_{1|k} & 0 \end{bmatrix}, \; \bar{v}_k \coloneqq \begin{bmatrix} v_{0|k} \\ v_{1|k} \\ \vdots \\ v_{N-1|k} \end{bmatrix},$$

such that $\bar{u}_k = \bar{v}_k + \bar{M}_k \bar{w}_k$. Before we proceed, we reformulate the predicted state sequence (7.3) by means of the SADF policy (7.10), i.e.,

$$\begin{aligned} \bar{x}_k &= \bar{A} x_{0|k} + \bar{B}\bar{u}_k + \bar{E}\bar{w}_k \\ &\overset{(7.10)}{=} \bar{A} x_{0|k} + \bar{B}\bar{v}_k + (\bar{B}\bar{M}_k + \bar{E})\bar{w}_k \\ &= \bar{z}_k + (\bar{B}\bar{M}_k + \bar{E})\bar{w}_k. \end{aligned} \tag{7.11}$$

The vector $\bar{v}_k$ can be interpreted as the predicted control input that corresponds to the predicted state trajectory $\bar{z}_k = \bar{A} x_{0|k} + \bar{B}\bar{v}_k$.

Remark 7.3. *The simplified affine disturbance feedback parameterization contains* $(N-1)n_{\mathrm{u}}n_{\mathrm{w}}$ *decision variables that grow linearly in the prediction horizon, whereas the original affine disturbance feedback grows quadratically in* N *with* $N(N-1)n_{\mathrm{u}}n_{\mathrm{w}}/2 + n_{\mathrm{u}}$ *decision variables [171].*

7.2.1 Distributionally robust chance constraints

In the following, we replace the individual chance constraints (7.5b)-(7.5c) with distributionally robust chance constraints of the form

$$\inf_{\mu\in\mathcal{P}} \mathbb{P}(h_{t,r}^\top \bar{x}_k \leq 1 \mid x_{0|k}) \overset{(7.11)}{=} \inf_{\mu\in\mathcal{P}} \mathbb{P}(h_{t,r}^\top(\bar{z}_k + [\bar{B}\bar{M}_k + \bar{E}]\bar{w}_k) \leq 1 \mid x_{0|k}) \geq p_{\mathrm{r}}^{\mathrm{x}}, \tag{7.12}$$

while the input constraints are given by

$$\inf_{\mu\in\mathcal{P}} \mathbb{P}(l_{t,s}^\top \bar{u}_k \leq 1 \mid x_{0|k}) \overset{(7.10)}{=} \inf_{\mu\in\mathcal{P}} \mathbb{P}(l_{t,s}^\top(\bar{v}_k + \bar{M}_k\bar{w}_k) \leq 1 \mid x_{0|k}) \geq p_s^{\mathrm{u}}. \tag{7.13}$$

By definition of the ambiguity set (7.7), we have that $\mathbb{E}_\mu(w) = 0$, $\sup_{\mu\in\mathcal{P}} \mathbb{E}_\mu(ww^\top) = \kappa_\beta\hat{\Sigma}$ and w is i.i.d. for all times k. Therefore, since the random vector $\bar{w}_k \in \mathbb{R}^{Nn_{\mathrm{w}}}$ contains N-times the random variable w, the vector is zero-mean and the worst-case covariance matrix is given by

$$\hat{\Sigma}_N \coloneqq \sup_{\mu\in\mathcal{P}} \mathbb{E}_\mu(\bar{w}_k\bar{w}_k^\top) = I_{\mathrm{N}} \otimes {}^2\kappa_\beta\hat{\Sigma}. \tag{7.14}$$

To obtain tractable expressions for the distributionally robust chance constraints (7.12)-(7.13), we apply [26, Thm 3.1] and obtain deterministic second-order cone (SOC) constraints of the form

$$h_{t,r}^\top \bar{z}_k \leq 1 - \sqrt{\kappa_\beta}\sqrt{\frac{p_{\mathrm{r}}^{\mathrm{x}}}{1-p_{\mathrm{r}}^{\mathrm{x}}}}\|h_{t,r}^\top(\bar{B}\bar{M}_k + \bar{E})\hat{\Sigma}_N^{1/2}\|_2 \tag{7.15}$$

$$l_{t,s}^\top \bar{v}_k \leq 1 - \sqrt{\kappa_\beta}\sqrt{\frac{p_{\mathrm{s}}^{u}}{1-p_{\mathrm{s}}^{u}}}\|l_{t,s}^\top\bar{M}_k\hat{\Sigma}_N^{1/2}\|_2. \tag{7.16}$$

Remark 7.4. *Note that in our setting, the SOC constraints* (7.15) - (7.16) *are equal to distributionally robust CVaR constraints, which follows from [177, Thm 2.2] and the fact that* (7.12) *and* (7.13) *are linear (and thus concave) in the uncertainty* w*, cf. [103, Thm. 4]. Therefore, it is expected that the chance constraints are conservatively satisfied.*

[2]The operator $\otimes$ denotes the Kronecker product. We use this to describe block diagonal matrices, e.g., $I_2 \otimes Q = \mathrm{blkdiag}(Q, Q)$, where I_2 is a 2×2 identity matrix and Q an arbitrary matrix.

7.2.2 Distributionally robust cost function

Similar to the previous section, the cost function (7.5a) is robustified to distributional uncertainty. First, we cast the quadratic cost function (7.5a) into a vector-matrix representation by defining block diagonal weighting matrices $\bar{Q} = \text{blkdiag}(I_N \otimes Q, P)$ and $\bar{R} = I_N \otimes R$, such that

$$\begin{aligned} J_k(k) &= \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\|x_{N|k}\|_P^2 + \sum_{t=0}^{N-1} \|x_{t|k}\|_Q^2 + \|u_{t|k}\|_R^2 \,\middle|\, x_{0|k} \right) \\ &= \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\bar{x}_k^\top \bar{Q} \bar{x}_k + \bar{u}_k^\top \bar{R} \bar{u}_k \mid x_{0|k} \right). \end{aligned} \tag{7.17}$$

Then we use the assumption that $\mathbb{E}(\bar{w}_k) = 0$ and substitute the SADF policy (7.10) together with the state prediction (7.11) into (7.17), resulting in

$$J_k(k) = \underbrace{\bar{z}_k^\top \bar{Q} \bar{z}_k + \bar{v}_k^\top \bar{R} \bar{v}_k}_{\text{mean part}} + \underbrace{\sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\bar{w}_k^\top \left[(\bar{B}\bar{M}_k + \bar{E})^\top \bar{Q} (\bar{B}\bar{M}_k + \bar{E}) + \bar{M}_k^\top \bar{R} \bar{M}_k \right] \bar{w}_k \right)}_{\text{variance part}}.$$

Next, we use the trace trick, which involves applying the trace operator to the above equation. By using linearity of the trace, we can separate the mean from the variance part. Furthermore, due to linearity of the expectation and the cyclic invariance property of the trace, we can factor out $\bar{w}_k$ from the variance part, so that

$$\begin{aligned} J_k(k) &= \text{tr}\left(\bar{z}_k^\top \bar{Q} \bar{z}_k + \bar{v}_k^\top \bar{R} \bar{v}_k \right) + \text{tr}\left(\sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu (\bar{w}_k \bar{w}_k^\top) \left[(\bar{B}\bar{M}_k + \bar{E})^\top \bar{Q} (\bar{B}\bar{M}_k + \bar{E}) + \bar{M}_k^\top \bar{R} \bar{M}_k \right] \right) \\ &\overset{(7.14)}{=} \bar{z}_k^\top \bar{Q} \bar{z}_k + \bar{v}_k^\top \bar{R} \bar{v}_k + \text{tr}\left(\hat{\Sigma}_N \left[(\bar{B}\bar{M}_k + \bar{E})^\top \bar{Q} (\bar{B}\bar{M}_k + \bar{E}) + \bar{M}_k^\top \bar{R} \bar{M}_k \right] \right), \end{aligned} \tag{7.18}$$

where the second equality used the fact that the trace of a scalar is equal to the scalar itself. For the cost $J_k(\cdot)$, we use the convention that the subscript denotes the time on which the expected value is conditioned on, while the argument denotes the closed-loop time instant at which the underlying MPC optimization problem is solved.

Remark 7.5. *The cost function* (7.18) *is formulated for the mean and variance of the states. Note that due to the SADF parameterization, the gain matrix $\bar{M}_k$ is a decision variable, which allows us to minimize the state and input variance. This aspect was neglected in other proposed MPC schemes, e.g., Section 6.3, where we used a fixed feedback gain. We refer to paragraph Control parameterization in Section 2.3.1 for more details on this topic.*

7.2.3 Terminal constraints

We enforce stability of the controller by imposing constraints at the end of the prediction horizon, where we make the following assumption.

Assumption 7.1. *There exists a terminal controller $\pi_{\mathrm{f}}(z) = Kz$ and a terminal set $\mathbb{Z}_{\mathrm{f}}$, such that for all $z \in \mathbb{Z}_{\mathrm{f}}$*

$$
\begin{aligned}
&(A + BK)z \in \mathbb{Z}_{\mathrm{f}} \\
&h_{\mathrm{r}}^{\top} z \leq 1 - \sqrt{\frac{p_{\mathrm{r}}^{\mathrm{x}}}{1 - p_{\mathrm{r}}^{\mathrm{x}}}} \left\| h_r^{\top} \hat{\Sigma}_{\infty}^{1/2} \right\|_2 \ \forall r \in \{1, \ldots, n_{\mathrm{r}}\} \\
&l_{\mathrm{s}}^{\top} K z \leq 1 - \sqrt{\frac{p_{\mathrm{r}}^{\mathrm{x}}}{1 - p_{\mathrm{r}}^{\mathrm{x}}}} \left\| l_s^{\top} K \hat{\Sigma}_{\infty}^{1/2} \right\|_2 \ \forall s \in \{1, \ldots, n_{\mathrm{s}}\},
\end{aligned}
$$

where $\hat{\Sigma}_{\infty} = (A + BK)\hat{\Sigma}_{\infty}(A + BK)^{\top} + \kappa_{\beta} E \hat{\Sigma} E^{\top}$ and $h_r \in \mathbb{R}^{n_{\mathrm{x}}}$, $l_s \in \mathbb{R}^{n_{\mathrm{u}}}$ denote the half-space matrices from (7.2)*.*

The first condition of Assumption 7.1 ensures that the terminal set is invariant for the nominal system under the terminal controller, whereas the second and third conditions enforce the distributionally robust chance constraints for all $z \in \mathbb{Z}_{\mathrm{f}}$ under the worst-case stationary covariance matrix $\hat{\Sigma}_{\infty}$.

Remark 7.6. *Assumption 7.1 can be ensured with methods proposed in [42, Sec. 2.4.2], i.e., by using an ellipsoidal terminal set $\mathbb{Z}_{\mathrm{f}} = \{z \mid z^{\top} P z \leq \alpha\}$ as an α-scaled sublevel set of the terminal cost function $V_{\mathrm{f}}(z) = \|z\|_P^2$. It remains to find a scalar α such that the terminal state and input chance constraints (inequality constraints in Assumption 7.1) are satisfied, which can be easily determined with a linear program, e.g., similar to Problem 5.3.5.*

7.2.4 Interpolated initial constraint

The final and most crucial point in ensuring recursive feasibility is the selection of a suitable initial condition for the MPC optimization problem. In the following, we adopt a recently proposed initialization scheme from Köhler and Zeilinger [93], where we constrain $x_{0|k}$ on a line between $x(k)$ and the guaranteed feasible solution $z_{1|k-1}$, i.e.,

$$
x_{0|k} = (1 - \lambda_k) x(k) + \lambda_k z_{1|k-1}, \tag{7.19}
$$

where $\lambda_k \in [0, 1]$. The advantage is that only one optimization problem needs to be solved, where $\lambda_k = 1$ reflects the guaranteed feasible solution (Mode 2) and $\lambda_k = 0$ the feedback strategy (Mode 1). Moreover, this allows a more natural definition of the cost function to prove that the closed-loop performance is not worse than for a linear controller $u = Kx$, cf. [93], where K is the controller gain associated with (7.6).

7.2.5 Optimization problem

At each time instant $k \in \mathbb{N}$, we solve the following MPC optimization problem.

Problem 7.2.1. *(Moment-based DR-MPC)*

$$\min_{\bar{v}_k, \bar{M}_k, \lambda_k} \quad (7.18) \tag{7.20a}$$

$$\text{s.t.} \quad \bar{z}_k = \bar{A} x_{0|k} + \bar{B} \bar{v}_k \tag{7.20b}$$

$$(7.15), (7.16), (7.19), \lambda_k \in [0, 1] \tag{7.20c}$$

$$z_{N|k} \in \mathbb{Z}_{\mathrm{f}}. \tag{7.20d}$$

The solution of the MPC optimization problem 7.2.1 is the optimal SADF pair $(\bar{v}_k^*, \bar{M}_k^*)$ and the states $\bar{z}_k^*$. To obtain the control input at time k, we recall [171, Thm. 1], which establishes an equivalence between the SADF parameterization (7.10) and a state feedback parameterization. By linear superposition, we can thus establish also an equivalence to the error feedback parameterization $u_{t|k}^{\mathrm{ef}} = g_{t|k} + \sum_{i=0}^{t} K_{t-i|k}(x_{i|k} - z_{i|k})$. In other words, the state and input trajectories $(\bar{x}_k, \bar{u}_k)$ resulting from the SADF parameterization with $(\bar{v}_k^*, \bar{M}_k^*)$ are equivalent to the ones obtained from the error feedback parameterization with $(\bar{g}_k^*, \bar{K}_k^*)$, where

$$\bar{K}_k \coloneqq \begin{bmatrix} K_{0|k} & 0 & \dots & 0 & 0 \\ K_{1|k} & K_{0|k} & \dots & 0 & 0 \\ \vdots & \ddots & \ddots & \vdots & 0 \\ K_{N-1|k} & \dots & K_{1|k} & K_{0|k} & 0 \end{bmatrix}, \bar{g}_k \coloneqq \begin{bmatrix} g_{0|k} \\ g_{1|k} \\ \vdots \\ g_{N-1|k} \end{bmatrix}.$$

Similar to [171], the optimal error feedback pair $(\bar{g}_k^*, \bar{K}_k^*)$ is obtained by

$$\bar{K}_k^* = (I + \bar{M}_k^* \bar{E}^\dagger \bar{B})^{-1} \bar{M}_k^* \bar{E}^\dagger \tag{7.21a}$$

$$\bar{g}_k^* = (I + \bar{M}_k^* \bar{E}^\dagger \bar{B})^{-1} (\bar{v}_k^* - \bar{M}_k^* \bar{E}^\dagger A z_{0|k}^*), \tag{7.21b}$$

while the input to system (7.1) is defined with the error feedback parameterization

$$u(k) = u_{0|k}^{\mathrm{ef}} = g_{0|k}^* + K_{0|k}^*(x(k) - z_{0|k}^*). \tag{7.22}$$

Remark 7.7. *Note that the chance constraints* (7.12) - (7.13) *depend on the information available at time k. In view of the initial condition* (7.19)*, this implies that whenever the MPC optimization problem 7.2.1 is feasible with $\lambda_k = 0$, the probability operator in* (7.12) - (7.13) *is conditioned on time k, resulting in closed-loop constraint satisfaction, while for $\lambda_k \in (0, 1]$ the constraints are verified in prediction, i.e., conditioned on the last time instant $k - \tau$ when problem 7.2.1 was feasible with $\lambda_{k-\tau} = 0$.*

By leaving λ_k un-penalized in the objective function (7.20a)*, we mimic a so-called hybrid scheme [57] with the intention to minimize the open-loop cost despite feasibility of $x(k)$. This can lead to an increase in constraint violations in presence of unmodeled disturbances, as we will demonstrate in Section 7.4. This approach is related to the indirect feedback scheme presented in Chapter 6.3, where the MPC cost function* (7.20a) *similarly penalizes*

the mean predictions rather than the nominal ones, i.e., due to the interpolating initial constraint (7.19).

However, adding an additional penalty term $c\lambda_k^2$ *with* $c > 0$ *to the objective function* (7.20a) *causes the MPC controller to favor feedback initialization* $x_{0|k} = x(k)$ *with the intention of introducing as much feedback as possible into the constraints, i.e., conditioning the probability operator in* (7.12) - (7.13) *on time* k *as often as possible. In this case, the behavior of the controller is closely related to a direct feedback MPC as outlined in Section 2.3.1, where a drawback is the deterioration of closed-loop performance, since the initial state cannot be freely chosen and the optimization problem therefore has fewer degrees of freedom, see Section 7.4 for a numerical comparison.*

7.3 Theoretical properties

In the following, we state the main results of this chapter.

Proposition 7.3.1. *Let Assumption 7.1 hold. If at time* $k = 0$ *the MPC optimization problem 7.2.1 admits a feasible solution with* $\lambda_0 = 0$*, then it is recursively feasible for all* $k \in \mathbb{N}$*.*

Proof. The proof can be found in Section 7.6. □

The following theorem establishes a quadratic stability result of the closed-loop system (7.1) under control law (7.22).

Theorem 7.1. *Let Assumption 7.1 hold and choose* $\beta \in (0, 1)$*,* $\epsilon \in (0, 0.5)$ *and* N_s*, such that* (7.8) *holds true and let* $\hat{\Sigma}$ *be the corresponding empirical covariance matrix. Suppose that at time* $k = 0$ *a feasible solution to problem 7.2.1 exists. Then, for all* $k \in \mathbb{N}$ *the optimal cost* $J_k^*(k+1)$ *satisfies*

$$J_k^*(k+1) - J_k^*(k) \leq -\mathbb{E}(\|x(k)\|_Q^2 + \|u(k)\|_R^2 | x(k)) + \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top).$$

Furthermore, with a probability of at least $1-\beta$ *the closed-loop system achieves the following asymptotic average bound*

$$\mathbb{P}\left(\lim_{T\to\infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E}_{\mu^*}(\|x(k)\|_Q^2 + \|u(k)\|_R^2 | x(0)) \leq \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top) \right) \geq 1 - \beta. \tag{7.23}$$

Proof. The proof can be found in Section 7.6. □

Remark 7.8. *By adding an additional penalty term* $c\lambda_k^2$ *to the cost function (Remark 7.7), the performance bound* (7.23) *contains an additional term* c*, i.e.,* $\kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top) + c$*. This additional constant is related to the Lipschitz-based arguments we have used in direct feedback SMPC schemes, e.g., as in Theorem 4.1, which renders the performance bound worse than that from the linear terminal controller. See also [93] for a more in-depth discussion.*

7.4 Numerical example

In this section, we carry out a numerical example and consider a simple double integrator system

$$x(k+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(k) + \begin{bmatrix} 1 \\ 0.5 \end{bmatrix} u + \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} w(k),$$

where $w(k) \sim \mathcal{N}(0, \Sigma)$ with $\Sigma = 0.01^2 I$. For the ambiguity set, we select $\beta = 0.05$, $\epsilon = 0.0428$. Since $w(k)$ is a zero-mean Gaussian it follows that $\xi = \Sigma^{1/2} w$ is sub-Gaussian with variance $\sigma^2 = 1$. According to Proposition 7.1.1, we require $N_\mathrm{s} \geq 516$ samples to give the guarantee that $\mathbb{P}(\mu^* \in \mathcal{P}) \geq 1 - \beta$. For the MPC cost function, we choose the weighting matrices $Q = \left[\begin{smallmatrix} 10 & 0 \\ 0 & 10 \end{smallmatrix}\right]$, $R = 1$ and $P = \left[\begin{smallmatrix} 20.5988 & 5.9161 \\ 5.9161 & 14.2284 \end{smallmatrix}\right]$. We impose a single chance constraint $\mathbb{P}(x_2(k) \leq 1) \geq p^\mathrm{x}$ and use an ellipsoidal terminal set $\mathbb{Z}_\mathrm{f} = \{z \,|\, z^\top P z \leq \alpha\}$, where $\alpha = 0.5293$ is obtained from $N_\mathrm{s} = 517$ samples. We keep α constant for each experiment and select a prediction horizon of $N = 10$. Note that the choice of α is quite conservative, i.e., for $N_\mathrm{s} = 10^3$ the resulting terminal set is already 20.4 times larger, while under exact moment information we can enlarge the terminal set about 21.9 times.

Performance and constraint satisfaction Starting at an initial condition $x(0) = [6, 0]^\top$, we performed 10^3 Monte-Carlo simulations of the closed-loop system for different sample sizes $550 \leq N_\mathrm{s} \leq 10^6$. As it can be seen in Figure 7.2 (**left**), the expected cost converges asymptotically to the optimal cost derived with exact moment information as the sample size N_s increases. This result is in line with Corollary 7.1, which indicates a convergence rate of $\mathcal{O}(1/\sqrt{N_\mathrm{s}})$ (i.e., fast convergence for small N_s and increasingly slower for larger N_s). As for the chance constraints, it can be seen in Figure 7.2 (**right**) that as the number of samples increases, the controller becomes more confident to operate closer to the constraint. In Table 7.1, we compare for different prescribed probability levels p^x and sample sizes N_s the achieved empirical constraint satisfaction rate averaged over 10^4 Monte-Carlo runs. The discrepancy between the prescribed and empirical satisfaction rate follows from the conservatism of the distributionally robust chance constraints, cf. Remark 7.4. Finally, the MPC optimization problem 7.2.1 is reliably solved in 6 milliseconds on average on a desktop PC with an Intel Core i7-9700k processor, Yalmip [105] and MOSEK [4].

Table 7.1: Effect of sample size N_s on the worst-case empirical probability of satisfying the constraint $\mathbb{P}(x_2 \leq 1) \geq p^\mathrm{x}$. ©2022 IEEE.

p^x	$N_\mathrm{s} = 520$	$N_\mathrm{s} = 800$	$N_\mathrm{s} = 10^5$	$N_\mathrm{s} = 10^6$
0.7	100%	99.25%	86.95%	85.99%
0.8	100%	99.95%	93.81%	93.29%
0.9	100%	100%	99.17%	98.83%

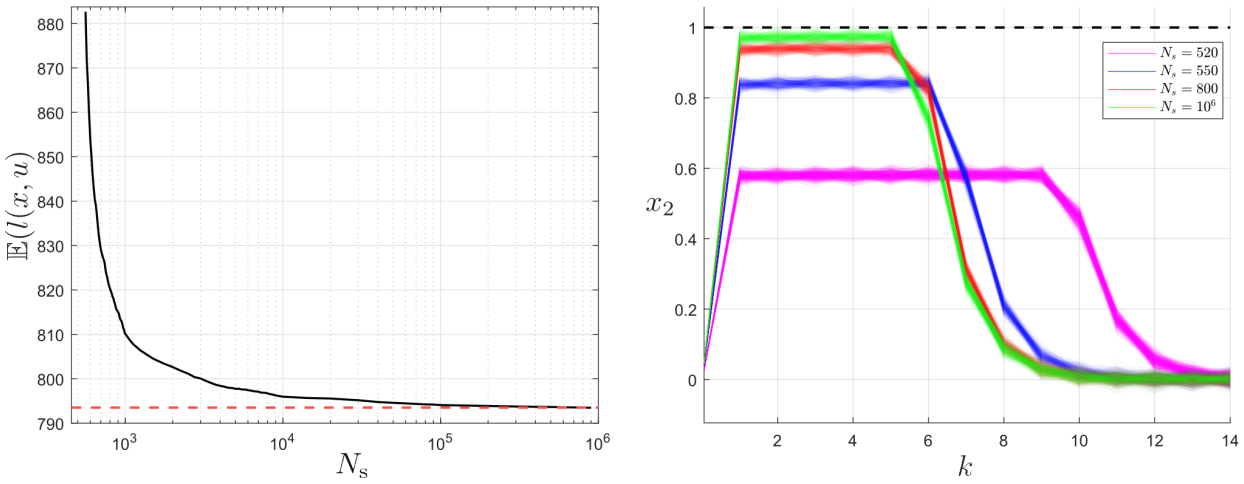

Figure 7.2: (**Left**) Expected closed-loop cost $l(x,u) = \sum_{k=1}^{15} \|x(k)\|_Q^2 + \|u(k)\|_R^2$ for different sample sizes N_s computed over 10^3 Monte-Carlo simulations (black) and optimal cost under exact moment information (red). (**Right**) Closed-loop trajectories for different N_s with $p_x = 0.9$. The black dotted line denotes the constraint $x_2 \leq 1$. ©2022 IEEE.

Unmodeled disturbances In the following, we investigate the benefits of adding a penalty term for λ_k to the objective function (7.20a). To this end, we keep the same simulation setup as before and introduce an unmodeled larger disturbance at time step $k = 5$ with $w(5) \sim \mathcal{N}(0, 6\Sigma)$. We add $c\lambda_k^2$ to the objective function (7.20a) to force the MPC optimization problem to prefer the feedback initialization over open-loop cost reduction (Remark 7.7) and opt to satisfy the chance constraint with 70% probability.

Table 7.2 reveals that penalization of λ_k increases the constraint satisfaction rate by sacrificing transient closed-loop performance compared to the unpenalized case $c = 0$ (Remark 7.8). Additionally, for $c > 0$ the chance constraint is empirically verified, whereas $c = 0$ violates the prescribed level of 70%.

Table 7.2: Comparison of different controller configurations for unmodeled disturbances. ©2022 IEEE.

	$c = 0$	$c = 10$	$c = 10^3$	$c = 10^6$
$\mathbb{E}(l(x,u))$	783.20	784.47	784.50	784.74
$\mathbb{P}(x_2(5) \leq 1)$	68.76%	73.37%	73.91%	75.08%

7.5 Summary

In this chapter, we have presented a DR-MPC framework for linear systems with sub-Gaussian additive disturbances under data-driven moment-based ambiguity sets, providing guarantees on closed-loop performance and recursive feasibility. We used a simplified affine disturbance feedback parameterization to obtain a tractable MPC optimization problem, where the distributionally robust chance constraints are reformulated as SOC constraints. The quadratic cost function is minimized subject to the worst-case distribution contained in the ambiguity set. We provided a design procedure for data-driven ambiguity sets, where we derive an explicit number of samples, such that a user-specified confidence bound holds true. We carried out a simple numerical example of a double integrator system to demonstrate the impact of the sample size on the conservatism of the controller.

7.6 Proofs

Proof of Proposition 7.1.1

The proof follows from [44, Thm. 8]. Define a random variable $\xi = \Sigma^{-1/2} w \sim \mathrm{subG}(\sigma^2)$ such that

$$\begin{aligned}\mathbb{E}(\xi) &= \Sigma^{-1/2}\mathbb{E}(w) = 0\\ \mathbb{E}(\xi\xi^\top) &= \Sigma^{-1/2}\mathbb{E}(ww^\top)\Sigma^{-1/2} = I\end{aligned}$$

and let $\tilde{I} = N_\mathrm{s}^{-1}\sum_{i=1}^{N_\mathrm{s}} \hat{\xi}^i(\hat{\xi}^i)^\top$ be the empirical covariance matrix of ξ. Consider now the empirical covariance matrix $\hat{\Sigma} = N_\mathrm{s}^{-1}\sum_{i=1}^{N_\mathrm{s}} \hat{w}^i(\hat{w}^i)^\top$ of the actual random variable w, which, after substitution of $\hat{w}^i = \Sigma^{1/2}\hat{\xi}^i$ equals

$$\hat{\Sigma} = \Sigma^{1/2}\left[N_\mathrm{s}^{-1}\sum_{i=1}^{N_\mathrm{s}} \hat{\xi}^i(\hat{\xi}^i)^\top\right]\Sigma^{1/2} = \Sigma^{1/2}\tilde{I}\Sigma^{1/2}. \tag{7.24}$$

From [81, Lem. A.1] we have with probability of at least $1-\beta$ that $\|\tilde{I} - I\|_2 \leq \gamma(N_\mathrm{s}, \beta/2)$, which is equivalent to

$$(1-\gamma(N_\mathrm{s},\beta/2))I \preceq \tilde{I} \preceq (1+\gamma(N_\mathrm{s},\beta/2))I. \tag{7.25}$$

Since we are only interested in an upper bound for the covariance matrix Σ, e.g., as required by (7.7), we find from the left inequality in (7.25) that

$$I \preceq \frac{1}{1-\gamma(N_\mathrm{s},\frac{\beta}{2})}\tilde{I} \overset{(7.24)}{\Longrightarrow} \Sigma \preceq \frac{1}{1-\gamma(N_\mathrm{s},\frac{\beta}{2})}\hat{\Sigma},$$

where we used the fact that condition (7.8) implies $\gamma(N_\mathrm{s},\beta/2) < 1$. Finally, condition (7.8) follows from assuming that $1-\gamma(N_\mathrm{s},\beta/2) > 0$, which is a quadratic inequality in the sample size $\sqrt{N_\mathrm{s}}$.

Proof of Proposition 7.3.1

Suppose that at time k problem 7.2.1 is feasible with $\lambda_k^* = 0$, $(\bar{v}_k^*, \bar{M}_k^*)$, $\bar{z}_k^*$ and equivalently with the error feedback parameterized input $\bar{u}_k = \bar{g}_k^* + \bar{K}_k^*(\bar{x}_k - \bar{z}_k^*)$ with $(\bar{g}_k^*, \bar{K}_k^*)$ due to [171, Thm. 1]. Now we construct the usual shifted candidate sequence $u_{t|k+1} = u_{t+1|k}$ for $t = 0, \dots, N-2$ and append the terminal controller $u_{N-1|k+1} = Kz_{N|k}^*$. The shifted mean states and controller gains satisfy $(z_{t|k+1}, K_{t|k+1}) = (z_{t+1|k}^*, K_{t+1|k}^*)$ for $t = 0, \dots, N-1$ appended with $(z_{N|k}, K_{N|k}) = ((A+BK)z_{N|k}^*, K)$. Recursive feasibility is then a consequence of Assumption 7.1. By stacking the shifted candidate sequences into the corresponding matrix and vector form, we obtain the triplet $(\bar{g}_{k+1}, \bar{K}_{k+1}, \bar{z}_{k+1})$. A feasible input pair $(\bar{v}_{k+1}, \bar{M}_{k+1})$ for problem 7.2.1 is then simply found by [171, eq. (24)], i.e.,

$$\begin{aligned}\bar{M}_{k+1} &= \bar{K}_{k+1}(I - \bar{B}\bar{K}_{k+1})^{-1}\bar{E} \\ \bar{v}_{k+1} &= \bar{K}_{k+1}(I - \bar{B}\bar{K}_{k+1})^{-1}(\bar{A}z_{0|k}^* + \bar{B}\bar{g}_{k+1}) + \bar{g}_{k+1}\end{aligned}$$

with $\lambda_{k+1} = 1$. This concludes the proof. □

Proof of Theorem 7.1

Consider that at time k a feasible solution exists. Now, at time $k+1$ we establish an expected cost decrease condition in case of $\lambda_{k+1} = 1$, where we consider the cost function (7.17) evaluated under the worst-case distribution $\hat{\mu}$, i.e.,

$$J_k(k) = \mathbb{E}_{\hat{\mu}}\left(\|x_{N|k}\|_P^2 + \sum_{t=0}^{N-1} \|x_{t|k}\|_Q^2 + \|u_{t|k}\|_R^2 \middle| x_{0|k} \right). \tag{7.26}$$

The predicted states $x_{t|k}$ are initialized with $x_{0|k} = x(k)$ and satisfy

$$x_{t+1|k} = Ax_{t|k} + Bu_{t|k} + Ew(t+k),$$

while the control input is given by $u_{t|k} = g_{t|k} + \sum_{i=0}^{t} K_{t-i|k}(x_{i|k} - z_{i|k})$. Due to the quadratic form of (7.26), we can equivalently write the cost function as $\tilde{J}_k(k) = J^{\mathrm{m}}(\tilde{x}_k, \tilde{u}_k) + J^{\mathrm{v}}(\kappa_\beta \hat{\Sigma}, \bar{K}_k)$, where the mean and variance part satisfy

$$\begin{aligned}J^{\mathrm{m}}(\tilde{x}_k, \tilde{u}_k) &= \|\tilde{x}_{N|k}\|_P^2 + \sum_{t=0}^{N-1} \|\tilde{x}_{t|k}\|_Q^2 + \|\tilde{u}_{t|k}\|_R^2 \\ J^{\mathrm{v}}(\kappa_\beta \hat{\Sigma}, \bar{K}_k) &= \mathrm{tr}(P\hat{\Sigma}_{N|k}^{\tilde{x}}) + \sum_{t=0}^{N-1} \mathrm{tr}(Q\hat{\Sigma}_{t|k}^{\tilde{x}} + R\hat{\Sigma}_{t|k}^{\tilde{u}}),\end{aligned}$$

where $\tilde{x}_{t|k} = \mathbb{E}_{\hat{\mu}}(x_{t|k} \mid x(k))$ and $\tilde{u}_{t|k} = \mathbb{E}_{\hat{\mu}}(u_{t|k} \mid x(k))$ denote the mean predictions, whereas $\hat{\Sigma}_{t+1|k}^{\tilde{x}} = (A + BK_{t|k})\hat{\Sigma}_{t|k}^{\tilde{x}}(A + BK_{t|k})^\top + \kappa_\beta E\hat{\Sigma}E^\top$ and $\hat{\Sigma}_{t|k}^{u} = \sum_{i=0}^{t} K_{t-i|k}\hat{\Sigma}_{i|k}^{\tilde{x}}K_{t-i|k}^\top$ denote the predicted state and input variance conditioned on $x(k)$, cf. [93]. Note that due to the

interpolating initial constraint (7.19), the predicted sequences are different to the nominal trajectories $\bar{z}_k$ and $\bar{v}_k$ if $\lambda_k \neq 1$. Therefore, the mean cost $J^{\mathrm{m}}(\tilde{x}_k, \tilde{u}_k)$ is defined based on the state and input mean rather than the nominal states and inputs. In view of this, the state variance of the initial state always satisfies $\hat{\Sigma}^{\tilde{x}}_{0|k} = 0$. Using the feasible candidate solution from Proposition 7.3.1, we can argue by optimality that

$$
\begin{aligned}
J^*_k(k+1) &\overset{\lambda_{k+1}=1}{\leq} J^{\mathrm{m}}(\tilde{x}_{k+1}, \tilde{u}_{k+1}) + J^{\mathrm{v}}(\kappa_\beta \hat{\Sigma}, \bar{K}_{k+1}) \\
&= J^{\mathrm{m}}(\tilde{x}_k, \tilde{u}_k) - \|\tilde{x}_{0|k}\|^2_Q - \|\tilde{u}_{0|k}\|^2_R + \|\tilde{x}_{N|k}\|^2_Q + \|K\tilde{x}_{N|k}\|^2_R - \|\tilde{x}_{N|k}\|^2_P + \|A_{\mathrm{K}}\tilde{x}_{N|k}\|^2_P \\
&+ J^{\mathrm{v}}(\kappa_\beta \hat{\Sigma}, \bar{K}_k) - \mathrm{tr}([Q + K^\top_{0|k} R K_{0|k}]\hat{\Sigma}^{\tilde{x}}_{0|k}) + \mathrm{tr}([Q + K^\top R K]\hat{\Sigma}^{\tilde{x}}_{N|k} + P A_{\mathrm{K}} \hat{\Sigma}^{\tilde{x}}_{N|k} A_{\mathrm{K}}^\top \\
&\quad + PE\kappa_\beta \hat{\Sigma} E^\top - P\hat{\Sigma}^{\tilde{x}}_{N|k}) \\
&\overset{(7.6)}{\leq} J^{\mathrm{m}}(\tilde{x}_k, \tilde{u}_k) + J^{\mathrm{v}}(\kappa_\beta \hat{\Sigma}, \bar{K}_k) - \|\tilde{x}_{0|k}\|^2_Q - \|\tilde{u}_{0|k}\|^2_R \\
&- \mathrm{tr}([Q + K^\top_{0|k} R K_{0|k}]\hat{\Sigma}^{\tilde{x}}_{0|k}) + \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top) \\
&= J^*_k(k) - \mathbb{E}_{\hat{\mu}}(\|x(k)\|^2_Q + \|u(k)\|^2_R \mid x(k)) + \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top),
\end{aligned}
$$

where $A_{\mathrm{K}} = A + BK$. To achieve the asymptotic average cost bound, we use standard arguments in stochastic MPC and obtain

$$
\begin{aligned}
0 &\leq \lim_{T\to\infty} \frac{1}{T}\big(J^*_0(T) - J^*_0(0)\big) \\
&\leq \lim_{T\to\infty} \frac{1}{T} \sum_{k=0}^{T-1} -\mathbb{E}_{\mu^*}(\|x(k)\|^2_Q + \|u(k)\|^2_R | x(0)) + \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top) \\
&\leq \kappa_\beta \mathrm{tr}(PE\hat{\Sigma}E^\top),
\end{aligned}
$$

while the probability bound (7.23) follows by definition of the ambiguity set (7.7), i.e., $\mathbb{P}(\Sigma \leq \kappa_\beta \hat{\Sigma}) \geq 1 - \beta$. □

8 Distributionally Robust MPC in application of wind farms

A large part of green energy production is currently covered by wind farms (WF), where several wind turbines (WT) are placed in close proximity to each other to reduce the cost of cabling and maintenance. One problem that occurs in such an environment is that each wind turbine generates a wake that moves downstream and is characterized by a flow velocity deficit and increased turbulence intensity [8]. The flow velocity deficit directly impacts the power production of downstream turbines [9], while the increased turbulence intensity increases the fatigue loads [18]. In this chapter, a DR-MPC is developed as a supervisory controller for a wind farm with the primary objective of dynamically distributing a required wind farm power reference $P_{\text{ref}}^{\text{wf}}$ to the individual $i = 1, \dots, N_{\text{wt}}$ wind turbines in the field, see Figure 8.1. The WT power references $P_{i,\text{ref}}^{\text{wt}}$ are then tracked by underlying local WT controllers, which operate on a much faster timescale (millisecond range) compared to the WF controller (second range). A secondary objective of the wind farm controller is to reduce fatigue loads of the turbines to increase their overall lifetime.

In Section 8.1, we derive a control-oriented model of the NREL 5MW wind turbine [87], which was developed with the intention of becoming a benchmark system for the development of large-scale wind farm controllers. Throughout this chapter, we use the Matlab/Simulink toolbox SimWindFarm [71] that serves as our simulation environment. In Section 8.2, we extend the DR-MPC formulation from Chapter 7 to include cost functions for output variables, while additionally an optimal wind turbulence predictor is introduced. Section 8.3 is devoted to a numerical example of a wind farm consisting of five wind turbines. This chapter is based on [116] [1].

Related work In [146], the authors consider the same setup as we do and use a stochastic MPC to design a supervisory control system for wind farms, adopting the probabilistic SMPC framework from [60]. However, their approach is based on the assumption that the true wind speed is Gaussian and the moments are known exactly. In [17], a scenario-based SMPC for power reference tracking is developed, where Gaussianity of the wind speed distribution is assumed. Fatigue load reduction is not considered explicitly in this work. The authors of [160] investigated a deterministic MPC approach for wind farm control. Similar to our approach, the goal was to track power and reduce mechanical stress, however,

[1] C. Mark and S. Liu. "Distributionally robust model predictive control for wind farms". In: *arXiv preprint arXiv:2303.03276* (2023). Accepted for presentation at the 22nd IFAC World Congress ©the authors.

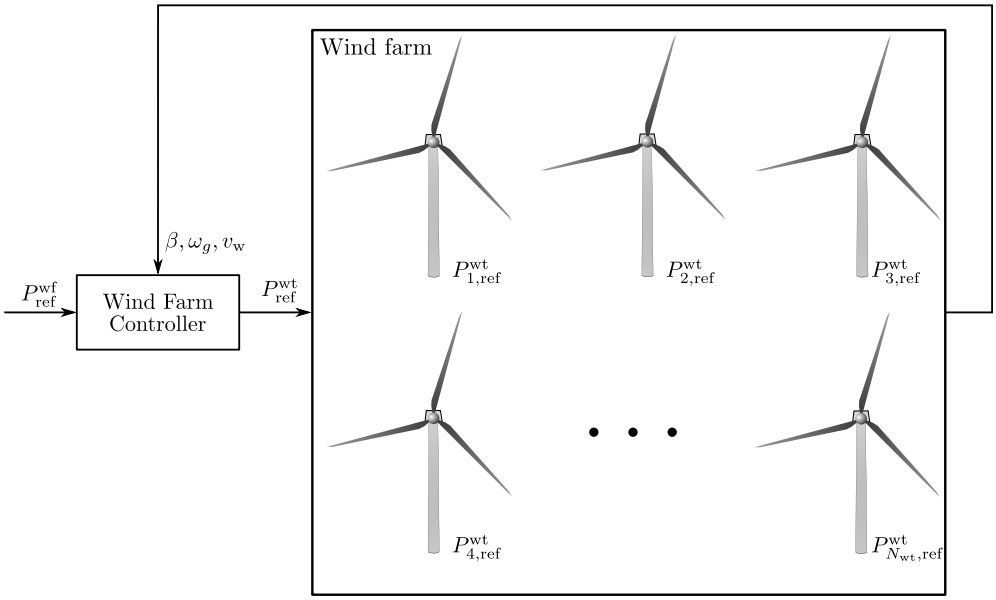

Figure 8.1: Supervisory control of wind farms.

the stochasticity of the wind is neglected and assumed to be constant over the prediction horizon. This approach was extended to a distributed MPC in [161]. In terms of wind turbine control, several papers have been published that address fatigue reduction, such as [55], where a robust MPC was developed for oscillation damping, or [69], where an economic nonlinear MPC was applied to reduce structural and actuator fatigue.

8.1 Control-oriented modeling of wind turbines

In the following, we first present the nonlinear wind turbine model from [158] and then a linearized version similar to [146]. Finally, the control-oriented model is validated with the full-scale nonlinear model provided by SimWindFarm (SWF) [71].

8.1.1 Nonlinear model

We first introduce the mechanical parts of the wind turbine, e.g., the aerodynamics, the generator, and the transmission system, while we then introduce the local NREL wind turbine controller that we include in the supervisory control design.

Aerodynamics The wind momentum is transferred to the rotor in form of an aerodynamic torque given by

$$T_\mathrm{r} = \frac{1}{2}\rho\pi R^3 v_\mathrm{w}^2 C_\mathrm{q}(\lambda,\beta), \quad \lambda = (\omega_\mathrm{r} R)/v_\mathrm{w}, \tag{8.1}$$

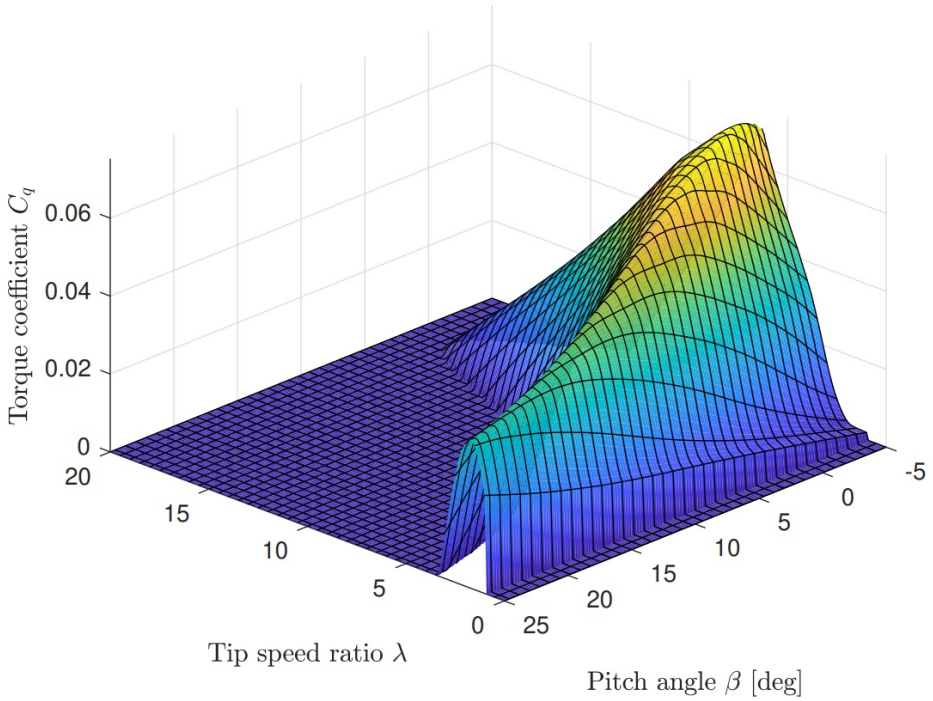

Figure 8.2: Torque coefficient of the NREL 5MW wind turbine.

where $\omega_\mathrm{r}\,[\mathrm{rad/s}]$ is the rotor speed, $\rho\,[\mathrm{kg/m^3}]$ the air density, $R\,[\mathrm{m}]$ the blade radius, $v_\mathrm{w}\,[\mathrm{m/s}]$ the wind speed and $C_\mathrm{q}(\lambda,\beta)$ the torque coefficient as a function of the tip speed ratio λ and the blade pitch angle $\beta\,[^\circ]$. The function $C_\mathrm{q}(\lambda,\beta)$ is typically known from measurements and is available as a look-up table individually for each wind turbine, e.g., Figure 8.2 depicts an example of such a C_q. During the conversion process, part of the wind energy is dissipated by secondary effects acting on the WT rotor. This results in a force acting orthogonally to the rotor plane and leading to a bending moment

$$M_\mathrm{t} = hF_\mathrm{t},$$

where $h\,[\mathrm{m}]$ denotes the tower height, while the exerted force $F_\mathrm{t}\,[\mathrm{N}]$ is called the thrust force, given by the static relation

$$F_\mathrm{t} = \frac{1}{2}\rho\pi R^2 v_\mathrm{w}^2 C_\mathrm{t}(\lambda,\beta), \tag{8.2}$$

where $C_\mathrm{t}(\lambda,\beta)$ is the corresponding thrust coefficient obtained from measurements.

Generator In the NREL wind turbine, the electrical power output is computed by the static equation

$$P_\mathrm{out} = \mu\omega_\mathrm{g} T_\mathrm{g}, \tag{8.3}$$

where μ denotes the generator efficiency, ω_g [rad/s] the generator angular velocity and T_g [Nm] the generator torque. Next, we introduce the transmission system that connects the rotor to the generator via the main shaft.

Transmission The transmission system describes the effects of the aerodynamic rotor torque T_r and the generator torque T_g on the rotor and generator angular velocities ω_r and ω_g, where we follow the lines of [146] and model a low-frequency low speed transmission shaft neglecting the fast torsional dynamics. The rotor dynamics with lumped inertia $\tilde{J} = J_\mathrm{r} + N_\mathrm{g}^2 J_\mathrm{g}$ can be described as

$$\begin{aligned}\dot{\omega}_\mathrm{r} &= \frac{1}{\tilde{J}}(T_\mathrm{r} - N_\mathrm{g} T_\mathrm{g}) \\ \omega_\mathrm{g} &= N_\mathrm{g}\omega_\mathrm{r},\end{aligned}$$

where $J_\mathrm{r}\,[\mathrm{kg/m^2}]$ is the rotor inertia, $J_\mathrm{g}\,[\mathrm{kg/m^2}]$ the generator inertia and N_g the gear ratio, while the main shaft torque $T_\mathrm{s}\,[\mathrm{Nm}]$ is given by

$$T_\mathrm{s} = \frac{N_\mathrm{g}^2 J_\mathrm{g}}{\tilde{J}} T_\mathrm{r} + \frac{N_\mathrm{g} J_\mathrm{r}}{\tilde{J}} T_\mathrm{g}. \tag{8.4}$$

Local NREL turbine controller In the following, we use the results from [158] to describe the local controller dynamics of the NREL wind turbine. The NREL controller consists of two control loops based on the measured generator angular velocity ω_g and the measured pitch angle β. The first computes a power reference P_ref, which is tracked by an underlying torque controller that ensures fast tracking of the torque reference

$$T_\mathrm{g,ref} = \frac{P_\mathrm{ref}}{\mu\omega_\mathrm{g}} = \frac{P_\mathrm{ref}}{\mu N_g \omega_\mathrm{r}}.$$

For the slow-scale supervisory control design, we can assume that $T_\mathrm{g,ref} \approx T_\mathrm{g}$, which implies that $P_\mathrm{out} \approx P_\mathrm{ref}$, where P_out is the electrical power output of the generator (8.3). The second control loop computes a pitch angle reference β_ref that is tracked by an underlying hydraulic pitch actuator, where we again assume that $\beta_\mathrm{ref} \approx \beta$ for the purposes of the supervisory control design. In the wind turbine tracking configuration assumed below, the NREL controller computes β_ref using a gain-scheduled PI controller to regulate the generator angular velocity ω_g to the rated speed ω_g0. To obtain an accurate model, we include the controller for the pitch angle in the supervisory control design, where use the following model of the gain-scheduled PI controller

$$\begin{aligned}\dot{\beta} &= \frac{1}{K_\mathrm{gs}}\left[\left(\frac{K_\mathrm{p}}{\tau_\omega} - K_\mathrm{i}\right)\omega_\mathrm{g}^\mathrm{f} - \frac{K_\mathrm{p}}{\tau_\omega}\omega_\mathrm{g}\right] \\ \dot{\omega}_\mathrm{g}^\mathrm{f} &= \frac{1}{\tau_\omega}\left(\omega_\mathrm{g} - \omega_\mathrm{g}^\mathrm{f}\right),\end{aligned}$$

where τ_ω is a time constant of a low-pass filter that lumps the effect of the sensors, $\omega_\mathrm{g}^\mathrm{f}$ is the corresponding filtered angular velocity of the generator, while $K_\mathrm{p}, K_\mathrm{i}$ and K_gs denote the proportional, integral and gain-scheduled controller gains. The adaptive correction gain K_gs depends on the power reference P_ref and the pitch angle β_ref. Note that asymptotically $\omega_\mathrm{g}^\mathrm{f} = \omega_\mathrm{g}$.

Thus, the overall nonlinear state space model can be written as

$$\dot{\beta} = \frac{1}{K_{\mathrm{gs}}} \left[\left(\frac{K_{\mathrm{p}}}{\tau_\omega} - K_{\mathrm{i}} \right) \omega_{\mathrm{g}}^{\mathrm{f}} - \frac{K_{\mathrm{p}}}{\tau_\omega} \omega_{\mathrm{g}} \right] \tag{8.5a}$$

$$\dot{\omega}_{\mathrm{r}} = \frac{1}{\tilde{J}} \left(\underbrace{\frac{1}{2} \rho \pi R^3 v_{\mathrm{w}}^2 C_{\mathrm{q}}(\lambda, \beta)}_{T_{\mathrm{r}}} - \frac{P_{\mathrm{ref}}}{\mu \omega_{\mathrm{r}}} \right) \tag{8.5b}$$

$$\dot{\omega}_{\mathrm{g}}^{\mathrm{f}} = \frac{1}{\tau_\omega} \left(\omega_{\mathrm{r}} - \frac{1}{N_{\mathrm{g}}} \omega_{\mathrm{g}}^{\mathrm{f}} \right). \tag{8.5c}$$

8.1.2 Linearized model

To use the DR-MPC framework from Chapter 7, we require a linear representation of the dynamics (8.5). To this end, let $x = [\beta, \omega_{\mathrm{r}}, \omega_{\mathrm{g}}^{\mathrm{f}}]$ be the state vector with operating point $x_0 = [\beta_0, \omega_{\mathrm{r}0}, \omega_{\mathrm{g}0}]$, $u = P_{\mathrm{ref}}$ the control input with operating point $u_0 = P_{\mathrm{ref}0}$ and $w = v_{\mathrm{w}}$ the wind speed with mean speed $w_0 = v_{\mathrm{w}0}$.

State equation We perform a first order Taylor approximation of (8.5b) to obtain the linear differential equation

$$\Delta\dot{\omega}_{\mathrm{r}} = \frac{1}{\tilde{J}} \left(T_{\mathrm{r}}^{\beta} \underbrace{(\beta - \beta_0)}_{\Delta\beta} + \left[T_{\mathrm{r}}^{\omega_{\mathrm{r}}} + \frac{P_{\mathrm{ref}0}}{\mu \omega_{\mathrm{r}0}^2} \right] \underbrace{(\omega_r - \omega_{\mathrm{r}0})}_{\Delta\omega_{\mathrm{r}}} - \frac{1}{\mu \omega_{\mathrm{r}0}} \underbrace{(P_{\mathrm{ref}} - P_{\mathrm{ref}0})}_{\Delta P_{\mathrm{ref}}} + T_{\mathrm{r}}^{v_{\mathrm{w}}} \underbrace{(v_{\mathrm{w}} - v_{\mathrm{w}0})}_{\Delta w} \right), \tag{8.6}$$

where we introduce deviation variables $\Delta\beta$, $\Delta\omega_{\mathrm{g}}$ and $\Delta\omega_{\mathrm{r}}$ to stabilize the operating point x_0 instead of the origin, while ΔP_{ref} and Δw denote deviations from the nominal input and mean wind speed. The terms $T_{\mathrm{r}}^{\omega_{\mathrm{r}}}$, T_{r}^{β} and $T_{\mathrm{r}}^{v_{\mathrm{w}}}$ denote the gradients of the rotor torque T_{r} with respect to ω_{r}, β and v_{w}, which are given by

$$\begin{aligned}
T_{\mathrm{r}}^{\omega_{\mathrm{r}}} &= \frac{\partial T_{\mathrm{r}}}{\partial \omega_{\mathrm{r}}} = \frac{1}{2} \rho \pi R^3 v_{\mathrm{w}0}^2 \left(\frac{\partial C_{\mathrm{q}}(\lambda, \beta)}{\partial \lambda} \frac{\partial \lambda}{\partial \omega_{\mathrm{r}}} \right) \Big|_{(\lambda_0, \beta_0)} \overset{\lambda = \frac{\omega_{\mathrm{r}} R}{v_{\mathrm{w}}}}{=} \frac{1}{2} \rho \pi R^3 v_{\mathrm{w}0} R \frac{\partial C_{\mathrm{q}}(\lambda, \beta)}{\partial \lambda} \Big|_{(\lambda_0, \beta_0)} \\
T_{\mathrm{r}}^{\beta} &= \frac{\partial T_{\mathrm{r}}}{\partial \beta} = \frac{1}{2} \rho \pi R^3 v_{\mathrm{w}0}^2 \frac{\partial C_{\mathrm{q}}(\lambda, \beta)}{\partial \beta} \Big|_{(\lambda_0, \beta_0)} \\
T_{\mathrm{r}}^{v_{\mathrm{w}}} &= \frac{\partial T_{\mathrm{r}}}{\partial v_{\mathrm{w}}} = \frac{1}{2} \rho \pi R^3 \left(2 v_{\mathrm{w}0} C_{\mathrm{q}}(\lambda_0, \beta_0) + \left(\frac{\partial C q(\lambda, \beta)}{\partial \lambda} \frac{\partial \lambda}{\partial v_{\mathrm{w}}} \right) \Big|_{(\lambda_0, \beta_0)} \right) \\
&\overset{\lambda = \frac{\omega_{\mathrm{r}} R}{v_{\mathrm{w}}}}{=} \frac{1}{2} \rho \pi R^3 \left(2 v_{\mathrm{w}0} C_{\mathrm{q}}(\lambda_0, \beta_0) - \omega_{\mathrm{r}0} R \frac{\partial C q(\lambda, \beta)}{\partial \lambda} \Big|_{(\lambda_0, \beta_0)} \right).
\end{aligned}$$

It remains to cast (8.6) together with the already linear differential equations (8.5a) and (8.5c) into the state space form

$$\Delta\dot{x} = A\Delta x + B\Delta u + E\Delta w, \tag{8.7}$$

where $\Delta x = x - x_0$, $\Delta u = u - u_0$ with

$$A = \begin{bmatrix} 0 & \frac{-K_\mathrm{p}N_\mathrm{g}}{K_\mathrm{gs}\tau_\mathrm{w}} & \frac{K_\mathrm{p}-K_\mathrm{i}\tau_\mathrm{w}}{K_\mathrm{gs}\tau_\mathrm{w}} \\ \frac{T_\mathrm{r}^\beta}{\tilde{J}} & \frac{1}{\tilde{J}}(T_\mathrm{r}^{\omega_\mathrm{r}} + \frac{P_\mathrm{ref0}}{\mu\omega_\mathrm{r0}^2}) & 0 \\ 0 & \frac{1}{\tau_\mathrm{w}} & -\frac{1}{N_\mathrm{g}\tau_\mathrm{w}} \end{bmatrix}, B = \begin{bmatrix} 0 \\ -\frac{1}{\tilde{J}\mu\omega_\mathrm{r0}} \\ 0 \end{bmatrix}, E = \begin{bmatrix} 0 \\ \frac{T_\mathrm{r}^{v_\mathrm{w}}}{\tilde{J}} \\ 0 \end{bmatrix}.$$

Output equation We define the system output as the tower bending force F_t and the shaft torque T_s, i.e., $y = (F_\mathrm{t}, T_\mathrm{s})$. Linearizing (8.2) and (8.4) around the operating point $y_0 = (F_\mathrm{t0}, T_\mathrm{s0})$ yields

$$\begin{aligned}
\Delta F_\mathrm{t} &= F_\mathrm{t}^\beta \Delta\beta + F_\mathrm{t}^{\omega_\mathrm{r}}\Delta\omega_\mathrm{r} + F_\mathrm{t}^{v_\mathrm{w}}\Delta w \\
\Delta T_\mathrm{s} &= \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^\beta}{J_\mathrm{t}}\Delta\beta + \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^{\omega_\mathrm{r}}}{J_\mathrm{t}} - \frac{J_\mathrm{r}P_\mathrm{ref0}}{J_\mathrm{t}\mu\omega_\mathrm{r0}^2}\Delta\omega_\mathrm{r} + \frac{J_\mathrm{r}}{J_\mathrm{t}\mu\omega_\mathrm{r0}}\Delta P_\mathrm{ref} + \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^{v_\mathrm{w}}}{J_\mathrm{t}}\Delta w,
\end{aligned}$$

where the partial derivatives of the thrust force (8.2) are given by

$$\begin{aligned}
F_\mathrm{t}^{\omega_\mathrm{r}} &= \frac{\partial F_\mathrm{t}}{\partial\omega_\mathrm{r}} = \frac{1}{2}\rho\pi R^2 v_\mathrm{w0}^2 \left(\frac{\partial C_\mathrm{t}(\lambda,\beta)}{\partial\lambda}\frac{\partial\lambda}{\partial\omega_\mathrm{r}}\right)\Big|_{(\lambda_0,\beta_0)} \overset{\lambda=\frac{\omega_\mathrm{r}R}{v_\mathrm{w}}}{=} \frac{1}{2}\rho\pi R^3 v_\mathrm{w0}\frac{\partial C_\mathrm{t}(\lambda,\beta)}{\partial\lambda}\Big|_{(\lambda_0,\beta_0)} \\
F_\mathrm{t}^{\beta} &= \frac{\partial F_\mathrm{t}}{\partial\beta} = \frac{1}{2}\rho\pi R^2 v_\mathrm{w0}^2 \frac{\partial C_\mathrm{t}(\lambda,\beta)}{\partial\beta}\Big|_{(\lambda_0,\beta_0)} \\
F_\mathrm{t}^{v_\mathrm{w}} &= \frac{\partial F_\mathrm{t}}{\partial v_\mathrm{w}} = \frac{1}{2}\rho\pi R^2 \left(2v_\mathrm{w0}C_\mathrm{t}(\lambda_0,\beta_0) + v_\mathrm{w0}^2\left(\frac{\partial C_\mathrm{t}(\lambda,\beta)}{\partial\lambda}\frac{\partial\lambda}{\partial v_\mathrm{w}}\right)\Big|_{(\lambda_0,\beta_0)}\right) \\
&\overset{\lambda=\frac{\omega_\mathrm{r}R}{v_\mathrm{w}}}{=} \frac{1}{2}\rho\pi R^2 \left(2v_\mathrm{w0}C_\mathrm{t}(\lambda_0,\beta_0) - R\omega_\mathrm{r0}\frac{\partial C_\mathrm{t}(\lambda,\beta)}{\partial\lambda}\Big|_{(\lambda_0,\beta_0)}\right).
\end{aligned}$$

The individual output equations are state-space representable with

$$\Delta y = \underbrace{\begin{bmatrix} F_\mathrm{t}^\beta & F_\mathrm{t}^{\omega_\mathrm{r}} & 0 \\ \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^\beta}{J_\mathrm{t}} & \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^{\omega_\mathrm{r}}}{J_\mathrm{t}} - \frac{J_\mathrm{r}P_\mathrm{ref0}}{J_\mathrm{t}\mu\omega_\mathrm{r0}^2} & 0 \end{bmatrix}}_{C}\Delta x + \underbrace{\begin{bmatrix} 0 \\ \frac{J_\mathrm{r}}{J_\mathrm{t}\mu\omega_\mathrm{r0}} \end{bmatrix}}_{D}\Delta u + \underbrace{\begin{bmatrix} F_\mathrm{t}^{v_\mathrm{w}} \\ \frac{J_\mathrm{g}N_\mathrm{g}^2T_\mathrm{r}^{v_\mathrm{w}}}{J_\mathrm{t}} \end{bmatrix}}_{F}\Delta w.$$

Finally, we discretize the continuous time dynamics (8.7) with a sample time of 1 second using exact discretization, which, together with the previous output equation yields the discrete-time LTI system

$$\Delta x(k+1) = A|_{(\lambda_0,\beta_0,K_\mathrm{gs})}\Delta x(k) + B|_{(\lambda_0,\beta_0)}\Delta u(k) + E|_{(\lambda_0,\beta_0)}\Delta w(k) \tag{8.8a}$$

$$\Delta y(k) = C|_{(\lambda_0,\beta_0)}\Delta x(k) + D|_{(\lambda_0,\beta_0)}\Delta u(k) + F|_{(\lambda_0,\beta_0)}\Delta w(k), \tag{8.8b}$$

where the matrices A, B, E, C, D, F depend on the linearization point (λ_0, β_0). In the following, we show how one can obtain the required operating point (λ_0, β_0).

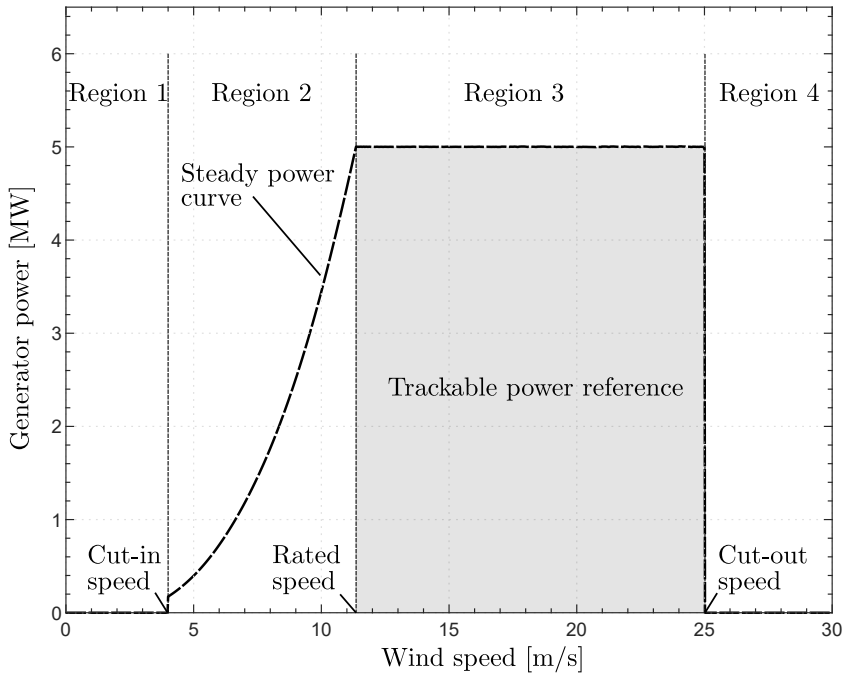

Figure 8.3: Steady power curve of the NREL 5MW wind turbine (dashed line) and area of trackable power references.

8.1.3 Determining the operating point

We characterize the operating point $x_0 = [\beta_0, \omega_{\mathrm{r}0}, \omega_{\mathrm{g}0}]$ of the turbine through its operating region and the power balance between demand and generation. First, note that a wind turbine has four operation regions, as illustrated in Figure 8.3. In Regions 1 and 4, the wind turbine is not operating because there is too little or too much wind. In Region 2, the main objective is to maximize the generated power by keeping the blade pitch angle β at zero, while in Region 3 the objective is to track a power setpoint by controlling the pitch angle β.

In this work, we consider only operating Region 3, i.e., the above rated area, where the rated generator speed $\omega_{\mathrm{g}0}$, as well as the rated rotor speed $\omega_{\mathrm{r}0}$ are known and held constant by the underlying turbine pitch and torque controller. In this way, we can assign power references to each turbine, which lie in the gray area in Figure 8.3. To derive the operating point β_0, we consider the aerodynamic torque (8.1) and find that

$$\frac{P_{\mathrm{ref}0}}{\mu\omega_{\mathrm{r}0}} = \frac{1}{2}\rho\pi R^3 v_{\mathrm{w}0}^2 C_{\mathrm{q}}(\lambda_0, \beta_0), \tag{8.9}$$

where $v_{\mathrm{w}0}$ is some above rated average wind speed and $\lambda_0 = (\omega_{\mathrm{r}0}R)/v_{\mathrm{w}0}$ the current tip-speed ratio. Therefore, the pitch angle β_0 is the only variable in equation (8.9), which can

readily be computed with the following optimization problem

$$\beta_0 = \underset{\beta}{\operatorname{argmin}} \quad \beta \tag{8.10a}$$

$$\text{s.t.} \quad \beta_{\min} \leq \beta \leq \beta_{\max} \tag{8.10b}$$

$$C_{\mathrm{q}}(\lambda_0, \beta) = \frac{2P_{\mathrm{ref}}}{\rho\pi R^3 v_{\mathrm{w0}}^2 \mu\omega_{\mathrm{r0}}}. \tag{8.10c}$$

In practice, the static map $C_{\mathrm{q}}(\lambda, \beta)$ is generated from measurements of the turbine, which, in case of the NREL 5MW wind turbine is encoded in a look-up table. Therefore, since λ_0 is fixed, we can heuristically search for β, such that the equality constraint (8.10c) approximately holds true.

8.1.4 Wind description

The wind w acting on the wind turbine is driven by a stochastic process that can generally be decomposed into a mean part w_0, that changes on a scale of ten minutes to several hours, superimposed with an additional turbulent part Δw, that changes on shorter time scales down to seconds [25, Sec. 2.1]. The turbulence can be characterized with the so-called turbulence intensity

$$T_{\mathrm{I}} = \frac{\sigma_{\Delta w}}{w_0},$$

where $\sigma_{\Delta w}$ denotes the standard deviation of the wind speed variation $\Delta w = w - w_0$ about the mean wind speed w_0, usually defined over ten minutes or one hour [25, Sec. 2.6]. Even though the turbulent wind can be roughly approximated by a Gaussian distribution, the actual wind gusts are non-Gaussian [125]. This highlights the idea of a data-driven moment-based distributionally robust approach, since we only need to estimate the covariance matrix of the zero-mean turbulence, rather than making any unrealistic distributional assumptions (see also [164] for a related distributionally robust approach for a wind turbine blade-pitch control design).

8.1.5 Model validation

To validate our control-oriented model, we make use of the SWF toolbox and set up a single NREL 5MW wind turbine with a constant power reference of 4MW as our nonlinear reference model. We consider a mean wind speed of 15 m/s and different turbulence intensities, e.g., for $T_{\mathrm{I}} = 0.1$, the turbulence Δw has a variance of 2.25. In Table 8.1, we list the resulting root mean square errors (RMSE) for the output and state vectors averaged over four independent wind datasets. As we can see in Table 8.1, a lower turbulence intensity is associated with a lower RMSE. This follows from the fact that in less turbulent scenarios, the effective wind speed is closer to the mean of 15 m/s, which coincides with the linearization point of the state space model (8.8). For higher turbulence intensities, the effective

Table 8.1: RMSE for different turbulence intensities T_I and mean wind speed of 15m/s.

T_I	$\Delta\beta$ [°]	$\Delta\omega_\mathrm{r}$ [$\frac{\mathrm{rad}}{\mathrm{s}}$]	$\Delta\omega_\mathrm{g}$ [$\frac{\mathrm{rad}}{\mathrm{s}}$]	ΔT_s [kNm]	ΔF_t [kN]
0.01	0.056	0.0028	0.276	20.1	3.7
0.05	0.273	0.0121	1.184	41.9	18.0
0.1	0.751	0.0248	2.418	78.6	42.9

wind speed can be driven far away from the linearization point, which therefore renders the control-oriented model inexact. This can be seen in Figure 8.4, where large deviations occur at around $t = 450$ and $t = 760$ seconds. Nonetheless, the control-oriented model adequately captures the slow-scale dynamics of the wind turbine and is used hereafter to model a wind farm.

8.1.6 Wind farm model

The wind farm model is obtained by indexing the state-space model (8.8) for each wind turbine $i = 1, \ldots, N_\mathrm{wt}$, such that

$$\begin{aligned} x^\mathrm{wf} &= \begin{bmatrix} \Delta x_1^\top & \ldots & \Delta x_{N_\mathrm{wt}}^\top \end{bmatrix}^\top \in \mathbb{R}^{n_x} \\ u^\mathrm{wf} &= \begin{bmatrix} \Delta u_1^\top & \ldots & \Delta u_{N_\mathrm{wt}}^\top \end{bmatrix}^\top \in \mathbb{R}^{n_u} \\ w^\mathrm{wf} &= \begin{bmatrix} \Delta w_1^\top & \ldots & \Delta w_{N_\mathrm{wt}}^\top \end{bmatrix}^\top \in \mathbb{R}^{n_w} \\ y^\mathrm{wf} &= \begin{bmatrix} \Delta y_1^\top & \ldots & \Delta y_{N_\mathrm{wt}}^\top \end{bmatrix}^\top \in \mathbb{R}^{n_y}, \end{aligned}$$

where the dynamic matrices of the wind farm model, i.e., $A^\mathrm{wf}, B^\mathrm{wf}, E^\mathrm{wf}, C^\mathrm{wf}, D^\mathrm{wf}, F^\mathrm{wf}$, are obtained by block diagonal stacking of the indexed matrices of (8.8). Thus, the resulting linear state space model has only independent subsystems without dynamic coupling and is given by

$$\Delta x^\mathrm{wf}(k+1) = A^\mathrm{wf}\Delta x^\mathrm{wf}(k) + B^\mathrm{wf}\Delta u^\mathrm{wf}(k) + E^\mathrm{wf}\Delta w^\mathrm{wf}(k) \tag{8.11a}$$

$$\Delta y^\mathrm{wf}(k) = C^\mathrm{wf}\Delta x^\mathrm{wf}(k) + D^\mathrm{wf}\Delta u^\mathrm{wf}(k) + F^\mathrm{wf}\Delta w^\mathrm{wf}(k). \tag{8.11b}$$

Furthermore, state and input constraints are enforced only locally, while an additional coupled input constraint is required to fulfill the reference tracking goal

$$\sum_{i=1}^{N_\mathrm{wt}} \Delta u_i(k) = 0,$$

such that the individual power references add up to the wind farm power reference, i.e., $P_\mathrm{ref}^\mathrm{wf} = \sum_{i=1}^{N_\mathrm{wt}} P_{i,\mathrm{ref0}} + \Delta u_i(k)$, where $P_{i,\mathrm{ref0}}$ is the nominal power reference (operating point) for the i-th wind turbine, cf. Section 8.1.2.

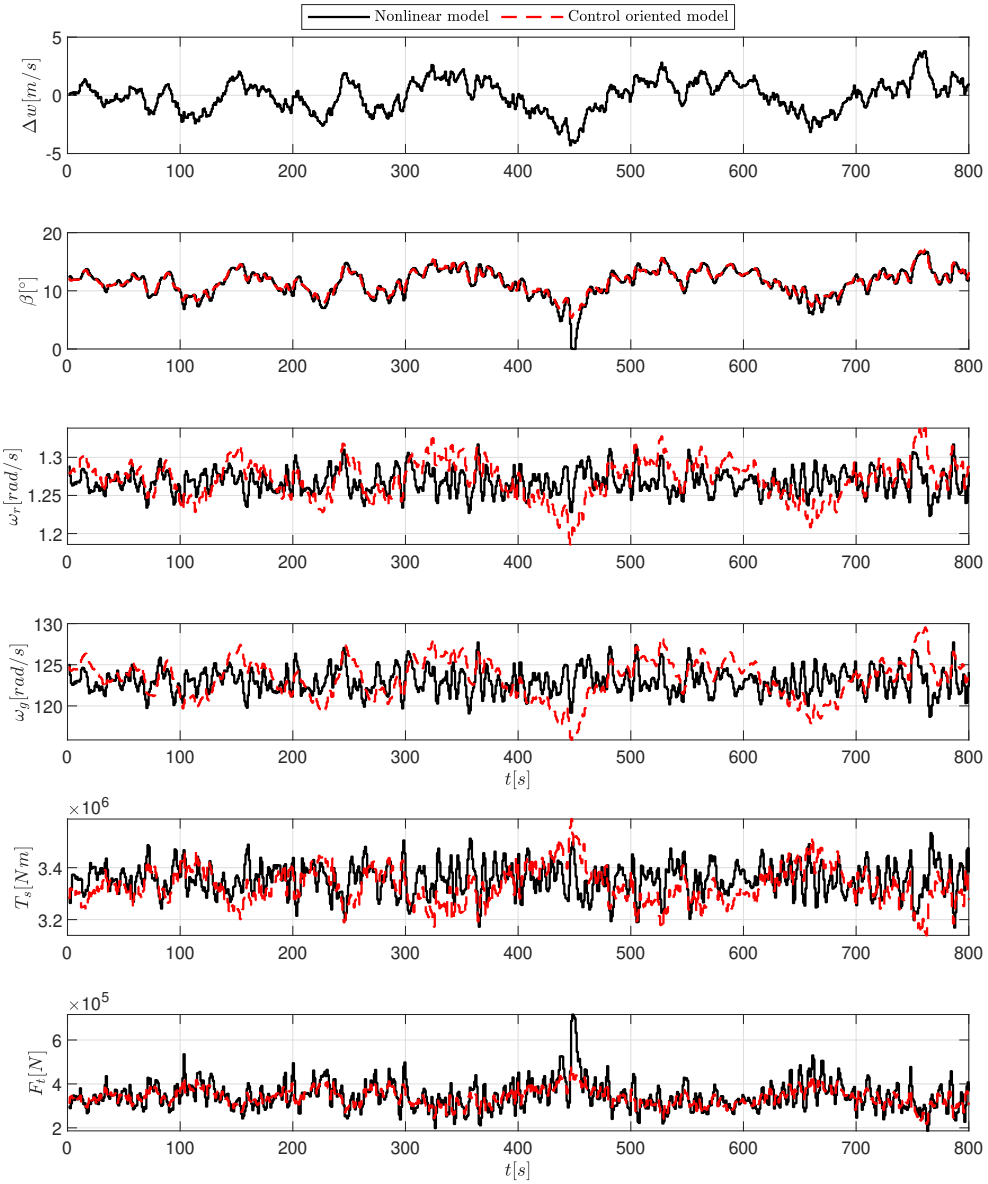

Figure 8.4: Comparison of nonlinear model with the control-oriented model for an average wind speed of $w_0 = 15$ m/s and turbulence intensity of $T_\mathrm{I} = 0.1$. The top plot shows the deviation of the effective wind speed from the average wind speed.

8.2 Distributionally robust Wind Farm MPC

The main objective of a wind farm controller in the above rated region is to distribute the wind power reference provided by the system operator to each wind turbine in the field while minimizing fatigue load [3, 91]. Fatigue loads result from repetitive stress reversals on a specific part of the structure, where typical fatigue prone components are the turbine tower and the generator shaft [159]. Therefore, we formulate an output cost function for (8.11b)

such that we minimize the mechanical load, i.e., tower bending force F_t and shaft torque T_s, of the entire wind farm. Further note that the DR-MPC scheme requires zero-mean i.i.d. additive disturbances, which the correlated wind speed does not satisfy. To render the additive disturbance (wind turbulence) $\Delta w = w - w_0$ zero-mean and i.i.d., we identify an autoregressive moving average (ARMA) model and formulate the DR-MPC problem based on the new stochastic residual part.

8.2.1 ARMA model

An ARMA model represents a stochastic process in terms of two polynomials, where the first one represents the auto-regressive (AR) part and the second one the moving average (MA) part [19]. In particular, an ARMA(p, q) model with p AR terms and q MA terms is given by

$$\Delta w(k) = \sum_{l=1}^{p} a_l \Delta w(k-l) + \sum_{l=1}^{q} b_l \epsilon(k-l) + \epsilon(k),$$

where ϵ is a zero-mean i.i.d white noise. In related work, e.g. [136, 146], the authors make a more stringent assumption for ϵ with the additional assumption that the noise is normally distributed, which in case of wind turbulence data is prone to be wrong, cf. [164]. Therefore, we treat ϵ as a zero-mean white noise with unknown (but finite) variance $\Sigma_\epsilon \succ 0$. In practice, one needs to identify the ARMA model with limited data. Therefore, the empirical variance is typically falsified due to sample inaccuracies, for which we introduce a moment-based ambiguity set that captures the true variance with high probability

$$\mathcal{P}(w_0, T_\mathrm{I}) \coloneqq \left\{ \mu \in \mathcal{M}(\mathbb{R}^n) \;\middle|\; \begin{array}{l} \mathbb{E}_\mu(\epsilon) = 0 \\ \mathbb{E}_\mu(\epsilon\epsilon^\top) \preceq \kappa_\beta^{(w_0,T_\mathrm{I})} \hat{\Sigma}_\epsilon^{(w_0,T_\mathrm{I})} \end{array} \right\}. \tag{8.12}$$

Note that we parameterize the ambiguity set with the mean wind speed and turbulence intensity pair (w_0, T_I). The ambiguity radius $\kappa_\beta^{(w_0,T_\mathrm{I})}$ can readily be found thanks to Proposition 7.1.1.

We identify for each wind turbine $i = 1, \ldots, N_\mathrm{wt}$ an ARMA$(p, p-1)$ model, which can be converted to a canonical form similar to [136], i.e,

$$\begin{aligned} \psi_i(k+1) &= A_{\psi,i}\psi_i(k) + B_{\psi,i}\epsilon_i(k) \\ \Delta w_i(k) &= C_{\psi,i}\psi_i(k), \end{aligned}$$

where the matrices are defined as follows

$$A_{\psi,i} \coloneqq \begin{bmatrix} a_{i,1} & 1 & 0 & \ldots & 0 \\ a_{i,2} & 0 & 1 & & 0 \\ \vdots & \vdots & & \ddots & \\ a_{i,p-1} & 0 & 0 & & 1 \\ a_{i,p} & 0 & 0 & \ldots & 0 \end{bmatrix}, B_{\psi,i} \coloneqq \begin{bmatrix} 1 \\ b_{i,1} \\ \vdots \\ b_{i,p-2} \\ b_{i,p-1} \end{bmatrix}, C_{\psi,i} \coloneqq \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}$$

and the auxiliary state vector ψ_i is given by $\psi_i(k) = \left[\Delta w_i^\top(k), \psi_{i,2}^\top(k), \ldots, \psi_{i,p}^\top(k)\right]$ with

$$\psi_{i,j}(k) = \sum_{l=j}^{p} a_{i,l}\Delta w_i(k+j-l-1) + \sum_{l=j-1}^{p-1} b_{i,l}\epsilon_i(k+j-l-1) \quad \forall i \in \{1,\ldots,N_{\mathrm{wt}}\}.$$

To obtain farm-wide wind predictions, we stack the local matrices and vectors together, such that $A_\psi = \mathrm{diag}(A_{\psi,1},\ldots,A_{\psi,N_{\mathrm{wt}}})$, $B_\psi = \mathrm{diag}(B_{\psi,1},\ldots,B_{\psi,N_{\mathrm{wt}}})$, $C_\psi = \mathrm{diag}(C_{\psi,1},\ldots,C_{\psi,N_{\mathrm{wt}}})$ and $\psi = \mathrm{col}_{i=1}^{N_{\mathrm{wt}}}(\psi_i)$. A N-step prediction of the turbulent wind is readily given by

$$\Delta\bar{w}_k \coloneqq \bar{C}_\psi\bar{A}_\psi\psi(k) + \bar{C}_\psi\bar{B}_\psi\bar{\epsilon}_k, \tag{8.13}$$

where $\bar{C}_\psi \coloneqq \mathrm{diag}(C_\psi,\ldots,C_\psi)$,

$$\bar{A}_\psi \coloneqq \begin{bmatrix} I \\ A_\psi \\ A_\psi^2 \\ \vdots \\ A_\psi^{N-1} \end{bmatrix}, \bar{B}_\psi \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ B_\psi & 0 & \ldots & 0 \\ A_\psi B_\psi & B_\psi & \ldots & 0 \\ \vdots & \ddots & \ddots & 0 \\ A_\psi^{N-2}B_\psi & \ldots & A_\psi B_\psi & B_\psi \end{bmatrix},$$

while the random vector $\bar{\epsilon}_k$ is zero-mean and each element is i.i.d. with variance $\Sigma_\epsilon^{(w_0,T_{\mathrm{I}})}$.

Remark 8.1. *As a byproduct of the ARMA model, a covariance reduction of the new random variable ϵ compared to the original random variable w is usually achieved. This aspect is important for the DR-MPC implementation because the chance constraints and the cost function depend directly on the covariance matrix, i.e., the lower the covariance, the lower the conservatism.*

8.2.2 Output prediction

We recall the input vector $\Delta\bar{u}_k = [\Delta u_{0|k}^\top,\ldots,\Delta u_{N-1|k}^\top]^\top \in \mathbb{R}^{Nn_u}$ and the predicted state vector $\Delta\bar{x}_k = [\Delta x_{0|k}^\top, \Delta x_{1|k}^\top,\ldots,\Delta x_{N|k}^\top]^\top \in \mathbb{R}^{(N+1)n_x}$ defined as

$$\Delta\bar{x}_k = \bar{A}\Delta x_{0|k} + \bar{B}\Delta\bar{u}_k + \bar{E}\Delta\bar{w}_k, \tag{8.14}$$

where the matrices $\bar{A}, \bar{B}, \bar{E}$ are given in Section 7.1. The output equation (8.11b) in a matrix form can be defined as

$$\Delta\bar{y}_k = \bar{C}\Delta\bar{x}_k + \bar{D}\Delta\bar{u}_k + \bar{F}\Delta\bar{w}_k,$$

where $\Delta\bar{y}_k = [\Delta y_{0|k}^\top,\ldots,\Delta y_{N|k}^\top] \in \mathbb{R}^{(N+1)n_y}$ and the matrices are given by

$$\bar{C} \coloneqq \begin{bmatrix} C & 0 & \ldots & 0 & 0 \\ 0 & C & \ldots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & C & 0 \\ 0 & 0 & \ldots & 0 & C \end{bmatrix}, \bar{D} \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ D & 0 & \ldots & 0 \\ 0 & D & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & D \end{bmatrix}, \bar{F} \coloneqq \begin{bmatrix} 0 & 0 & \ldots & 0 \\ F & 0 & \ldots & 0 \\ 0 & F & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & F \end{bmatrix}$$

with dimensions $\bar{C} \in \mathbb{R}^{(N+1)n_y \times (N+1)n_x}$, $\bar{D} \in \mathbb{R}^{(N+1)n_y \times Nn_u}$ and $\bar{F} \in \mathbb{R}^{(N+1)n_y \times Nn_w}$. Next, we include the ARMA prediction (8.13) into the wind farm model and consider the simplified affine disturbance feedback policy in the new random vector $\bar{\epsilon}_k$, i.e.,

$$\Delta \bar{u}_k = \Delta \bar{v}_k + \bar{M}_k \bar{\epsilon}_k. \tag{8.15}$$

Based on this, we can reformulate the predicted output sequence as follows

$$\begin{aligned}
\Delta \bar{y}_k &= \bar{C}\Delta \bar{x}_k + \bar{D}\Delta \bar{u}_k + \bar{F}\Delta \bar{w}_k \\
&\overset{(8.14)}{=} \bar{C}[\bar{A}\Delta x_{0|k} + \bar{B}\Delta \bar{u}_k + \bar{E}\Delta \bar{w}_k] + \bar{D}\Delta \bar{u}_k + \bar{F}\Delta \bar{w}_k \\
&\overset{(8.13)}{=} \bar{C}\bar{A}\Delta x_{0|k} + (\bar{C}\bar{B} + \bar{D})\Delta \bar{u}_k + (\bar{C}\bar{E} + \bar{F})\bar{C}_\psi \bar{A}_\psi \psi(k) + (\bar{C}\bar{E} + \bar{F})\bar{C}_\psi \bar{B}_\psi \bar{\epsilon}_k \\
&\overset{(8.15)}{=} \underbrace{\bar{C}\bar{A}\Delta x_{0|k} + (\bar{C}\bar{B} + \bar{D})\Delta \bar{v}_k + (\bar{C}\bar{E} + \bar{F})\bar{C}_\psi \bar{A}_\psi \psi(k)}_{\Delta \tilde{\bar{y}}_k} \\
&\quad + \underbrace{[\bar{C}\bar{B}\bar{M}_k + \bar{D}\bar{M}_k + (\bar{C}\bar{E} + \bar{F})\bar{C}_\psi \bar{B}_\psi]}_{\Psi_k} \bar{\epsilon}_k.
\end{aligned} \tag{8.16}$$

8.2.3 Cost function

To achieve the primary goal of power tracking and the secondary goal of fatigue load reduction, we formulate a quadratic cost function in the output and input deviations

$$\begin{aligned}
J_k &= \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\Delta \bar{y}_k^\top \bar{Q}_{\mathrm{y}} \Delta \bar{y}_k + \Delta \bar{u}_k^\top \bar{R} \Delta \bar{u}_k \middle| x(k) \right) \\
&\overset{(8.16)}{=} \sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\begin{bmatrix} \bar{\epsilon}_k \\ 1 \end{bmatrix}^\top \left(\begin{bmatrix} \Psi_k \\ \Delta \tilde{\bar{y}}_k \end{bmatrix}^\top \bar{Q}_{\mathrm{y}} \begin{bmatrix} \Psi_k \\ \Delta \tilde{\bar{y}}_k \end{bmatrix} + \begin{bmatrix} \bar{M}_k \\ \Delta \bar{v}_k \end{bmatrix}^\top \bar{R} \begin{bmatrix} \bar{M}_k \\ \Delta \bar{v}_k \end{bmatrix} \right) \begin{bmatrix} \bar{\epsilon}_k \\ 1 \end{bmatrix} \middle| x(k) \right) \\
&= \operatorname{tr} \left(\sup_{\mu \in \mathcal{P}} \mathbb{E}_\mu \left(\begin{bmatrix} \bar{\epsilon}_k \\ 1 \end{bmatrix} \begin{bmatrix} \bar{\epsilon}_k \\ 1 \end{bmatrix}^\top \middle| x(k) \right) \left[\bar{H}_{\mathrm{y},k}^\top \bar{Q}_{\mathrm{y}} \bar{H}_{\mathrm{y},k} + \bar{H}_{\mathrm{u},k}^\top \bar{R} \bar{H}_{\mathrm{u},k} \right] \right),
\end{aligned} \tag{8.17}$$

where $\bar{H}_{\mathrm{y},k} = \begin{bmatrix} \Psi_k^\top & \Delta \tilde{\bar{y}}_k^\top \end{bmatrix}^\top$ and $\bar{H}_{\mathrm{u},k} = \begin{bmatrix} \bar{M}_k^\top & \Delta \bar{v}_k^\top \end{bmatrix}^\top$. The inner distributionally robust expectation problem can be solved w.r.t. the moment-based ambiguity set (8.12), i.e.,

$$\begin{aligned}
\hat{\Sigma}_N^{(w_0, T_{\mathrm{I}})} &:= \sup_{\mu \in \mathcal{P}} \left\{ \begin{bmatrix} \bar{\tilde{\epsilon}}_k \\ 1 \end{bmatrix} \begin{bmatrix} \bar{\tilde{\epsilon}}_k \\ 1 \end{bmatrix}^\top \middle| x(k) \right\} \\
&\overset{\text{i.i.d.}}{=} \operatorname{blkdiag} \left(I_N \otimes \sup_{\mu \in \mathcal{P}(w_0, T_{\mathrm{I}})} \left\{ \begin{bmatrix} \epsilon \\ 1 \end{bmatrix} \begin{bmatrix} \epsilon \\ 1 \end{bmatrix}^\top \middle| x(k) \right\}, 1 \right) \\
&\overset{(8.12)}{=} \operatorname{blkdiag}(I_N \otimes \kappa_\beta^{(w_0, T_{\mathrm{I}})} \hat{\Sigma}_\epsilon^{(w_0, T_{\mathrm{I}})}, 1),
\end{aligned}$$

where the first equality follows from the i.i.d. sequence $\bar{\epsilon}_k$. The MPC cost function is obtained by substituting the latter into (8.17).

8.2.4 MPC optimization problem

At each time step $k \in \mathbb{N}$, we solve the following MPC optimization problem.

Problem 8.2.1 (Distributionally Robust Wind Farm MPC)**.**

$$\min_{\Delta\bar{v}_k, \bar{M}_k, \lambda_k} \quad \operatorname{tr}\big(\hat{\Sigma}_N^{(w_0,T_\mathrm{I})}[\bar{H}_{\mathrm{y},k}^\top \bar{Q}_\mathrm{y} \bar{H}_{\mathrm{y},k} + \bar{H}_{\mathrm{u},k}^\top \bar{R} \bar{H}_{\mathrm{u},k}]\big) \tag{8.18}$$

$$\text{s.t.} \quad \Delta\bar{z}_k = \bar{A}\Delta x_{0|k} + \bar{B}\Delta\bar{v}_k + \bar{E}\bar{C}_\psi \bar{A}_\psi \psi(k) \tag{8.19}$$

$$\tilde{\bar{y}}_k = \bar{C}\bar{A}\Delta x_{0|k} + (\bar{C}\bar{B} + \bar{D})\Delta\bar{v}_k + (\bar{F} + \bar{C}\bar{E})\bar{C}_\psi \bar{A}_\psi \psi(k) \tag{8.20}$$

$$\Delta x_{0|k} = (1-\lambda_k)\Delta x(k) + \lambda \Delta z^*_{1|k-1} \quad \lambda_k \in [0,1] \tag{8.21}$$

$$\bar{h}_t^\top \Delta\bar{z}_k \leq 1 - \sqrt{\frac{p_x}{1-p_x}} \|\bar{h}_t^\top (\bar{B}\bar{M}_k + \bar{E})(\hat{\Sigma}_N^{(w_0,T_\mathrm{I})})^{1/2}\|_2 \quad \forall t \in \{0,\dots,N-1\} \tag{8.22}$$

$$\bar{l}_t^\top \Delta\bar{v}_k \leq 1 - \sqrt{\frac{p_u}{1-p_u}} \|\bar{l}_t^\top \bar{M}_k (\hat{\Sigma}_N^{(w_0,T_\mathrm{I})})^{1/2}\|_2 \quad \forall t \in \{0,\dots,N-1\} \tag{8.23}$$

$$\mathbb{1}^\top \Delta v(k) = 0. \tag{8.24}$$

The optimal solution to Problem 8.2.1 is the SADF pair $(\Delta\bar{v}_k^*, \bar{M}_k^*)$ and the mean state prediction $\Delta\bar{z}_k^*$. Similar to Section 7.2.5, we obtain an equivalent admissible error feedback control policy via

$$\begin{aligned} \bar{K}_k^* &= (I + \bar{M}_k^* \bar{E}^\dagger \bar{B})^{-1} \bar{M}_k^* \bar{E}^\dagger \\ \Delta\bar{g}_k^* &= (I + \bar{M}_k^* \bar{E}^\dagger \bar{B})^{-1} (\Delta\bar{v}_k^* - \bar{M}_k^* \bar{E}^\dagger A \Delta z^*_{0|k}), \end{aligned}$$

while the input to the wind turbines is defined as

$$P_\mathrm{ref}^\mathrm{wt}(k) = u(k) = P_\mathrm{ref0}^\mathrm{wt} + \Delta g^*_{0|k} + K^*_{0|k}(\Delta x(k) - \Delta z^*_{0|k}),$$

where $P_\mathrm{ref0}^\mathrm{wt}$ is a vector of nominal power references for all wind turbines. Moreover, the input chance constraints (8.23) are used to bound the power deviation from the nominal set points, while optionally state chance constraints (8.22) can be imposed, e.g., for the pitch angle or generator speed.

8.3 Simulation results

In the following, we apply our proposed DR-MPC to a wind farm consisting of $N_\mathrm{wt} = 5$ NREL 5 MW wind turbines in a row, see Figure 8.5, where each WT is equidistantly arranged with $d = 400$ m. The wind turbine dynamics use the parameters given in Table 8.2. The gain scheduled parameter K_gs is obtained from the underlying look-up table.

Table 8.2: Parameters of the NREL 5MW wind turbine.

Parameter	Variable	Value	Unit
Blade radius	R	63	m
Rotor inertia	J_r	$35.44 \cdot 10^6$	kg/m^2
Generator inertia	J_g	534.116	kg/m^2
Rated generator speed	ω_{g0}	122.9096	rad/s
Rated rotor speed	ω_{r0}	1.2671	rad/s
Rated power	P_0	$5 \cdot 10^6$	W
Gearbox ratio	N_g	97	-
Generator efficiency	μ	0.944	-
Filter time constant	τ_ω	0.0125	s
Blade control proportional gain	K_p	-0.2143	-
Blade control integral gain	K_i	-0.0918	-
Air density	ρ	1.2231	kg/m^3

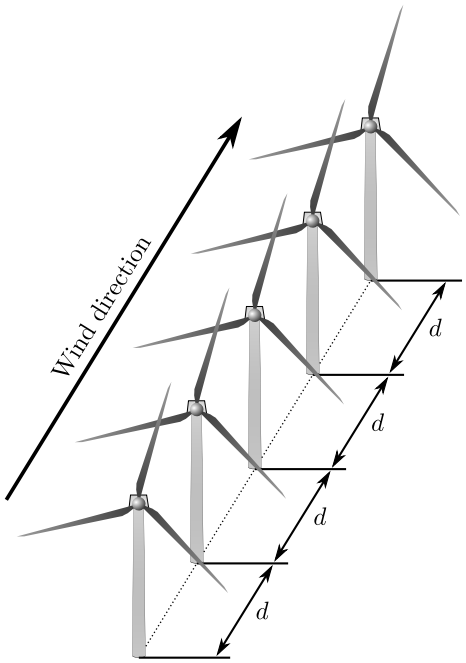

Figure 8.5: Simulation setup with five equidistantly arranged wind turbines.

Controllers We compare the DR-MPC to an open-loop *scheduler* that assigns a constant power references $P^{\mathrm{wt}}_{i,\mathrm{ref}} = 3\mathrm{MW}$ for each wind turbine $i = 1, \ldots, N_{\mathrm{wt}}$ for the entire simulation horizon of $T = 900$ s, i.e., the wind farm should nominally produce $P^{\mathrm{wf}}_{\mathrm{ref0}} = 15$ MW. In addition, we consider the *SWF controller* [71], which dynamically dispatches the power

references based on the available power estimates of each wind turbine

$$P_{i,\text{avail}} = \min\left\{P_0, \frac{1}{2}\pi\rho R^2 w_i^3 C_\text{p}^\text{max}\right\} \quad \forall i \in \{1,\dots,N_\text{wt}\},$$

where w_i is the measured wind speed of turbine i, P_0 the rated power and $C_\text{p}^\text{max} = 0.45$ the maximum power coefficient. Therefore, the SWF controller distributes the power as follows

$$P_{i,\text{ref}}^\text{wt} = \max\left\{0, \min\left\{P_0, \frac{P_\text{ref0}^\text{wf} P_{i,\text{avail}}}{\sum_{i=1}^{N_\text{wt}} P_{i,\text{avail}}}\right\}\right\} \quad \forall i \in \{1,\dots,N_\text{wt}\}.$$

For the DR-MPC, we identify for each wind turbine an ARMA$(3,2)$ model based on an independent wind scenario of 1000 time steps, i.e., $N_s = 1000$ samples. In view of Proposition 7.1.1, we derive an ambiguity radius of $\kappa_\beta^{(w_0,T_\text{I})} = 2.36$ with a confidence of $1-\beta = 0.95$.

Performance We use the following criteria to evaluate the performance of each controller:

- Tracking error $J_\text{p} = \sqrt{\dfrac{1}{T}\displaystyle\sum_{k=0}^{T-1}\sum_{i=1}^{N_\text{wt}} \frac{(P_{i,\text{out}}(k) - P_{i,\text{ref}}^\text{wt}(k))^2}{N_\text{wt} P_0}}$,
- Transmission shaft fatigue $J_\text{s} = \text{std}\left(\dfrac{\sum_{k=0}^{T-1}\sum_{i=1}^{N_\text{wt}} T_{i,\text{s}}(k)}{N_\text{wt} T_\text{s0}}\right)$,
- Tower fatigue $J_\text{t} = \text{std}\left(\dfrac{\sum_{k=0}^{T-1}\sum_{i=1}^{N_\text{wt}} F_{i,\text{t}(k)}}{N_\text{wt} T_\text{t0}}\right)$,

where $P_{i,\text{out}}$, $T_{i,\text{s}}$ and $F_{i,\text{t}}$ denote the power output, main shaft torque and tower bending force of turbine i, while $T_\text{s0} = 2.5 \cdot 10^6$ and $T_\text{t0} = 0.27 \cdot 10^6$ are standardization constants obtained from (8.2) and (8.4) for the nominal operating point of 3 MW. To reduce the tuning effort of the MPC cost function (8.18), we fix the output weight $\bar{Q}_\text{y}$ to

$$\bar{Q}_\text{y} = I_{N+1} \otimes \text{blkdiag}\left(\begin{bmatrix} q_{F_\text{t}} & 0 \\ 0 & q_{T_\text{s}} \end{bmatrix}, \dots, \begin{bmatrix} q_{F_\text{t}} & 0 \\ 0 & q_{T_\text{s}} \end{bmatrix}\right)$$

with standardized weights

$$q_{F_\text{t}} = \frac{1}{F_\text{t0}^2 N} \quad \text{and} \quad q_{T_\text{s}} = \frac{100}{T_\text{s0}^2 N}.$$

Analogously, we define the standardized input weighting matrix as

$$\bar{R} = I_N \otimes \begin{bmatrix} \frac{r}{P_0^2 N} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \frac{r}{P_0^2 N} \end{bmatrix},$$

where $r \in \mathbb{R}_{>0}$. Thus, it remains to tune the parameter r, which introduces a trade-off between tracking performance and fatigue load reduction. We consider a prediction horizon of $N = 5$ seconds for each simulation.

Table 8.3: Performance comparison of scenario 1. The first row denotes the scheduler performance, which we consider as the baseline, i.e., 100%. The other rows denote the performance w.r.t. the scheduler, where numbers smaller than 100% denote a performance increase and numbers greater than 100% a performance decrease.

Method	J_p	J_t	J_s
Scheduler	0.0811	0.2944	0.0740
SWF controller	100%	100%	100%
DR-MPC $r = 1$	99.983%	100.004%	100.003%
DR-MPC $r = 10$	99.983%	100.004%	100.003%
DR-MPC $r = 10^2$	99.972%	100.005%	100.001%
DR-MPC $r = 10^3$	99.995%	100.000%	99.998%

Scenario 1 – Strictly above rated wind speed

In the first scenario, we consider a wind field described by a mean velocity of $w_0 = 20\,\mathrm{m/s}$ and a turbulence intensity of $T_\mathrm{I} = 0.05$, which implies that the turbulent wind has $\sigma^2_{\Delta w} = (T_\mathrm{I} w_0)^2 = 1$ variance. The residual $\bar{\epsilon}_k$ of the identified farm-wide ARMA model (8.13) has the following empirical covariance matrix

$$\hat{\Sigma}_\epsilon^{(20,0.05)} = \mathrm{diag}(0.255,\ 0.270,\ 0.288,\ 0.262,\ 0.274),$$

which reduces the noise variance around 75% compared to $\sigma^2_{\Delta w}$.

The main assumption in the first scenario is that the wind speed never drops below the rated level of 11.4 m/s for all wind turbines, cf. Figure 8.3. Thus, each WT operates only in Region 3 throughout the simulation period. In Figure 8.6, the power output of each turbine is shown. In this operating region, the available power is always at its upper limit, i.e., 5 MW. Therefore, the SWF controller and the Scheduler essentially command constant reference points, while the DR-MPC adjusts the set points dynamically in a prescribed band of $\pm 0.2\,\mathrm{MW}$ around the nominal value of 3 MW with a probability of 90%. This is enforced with the input chance constraint (8.23). Additionally, we add a penalty term $5\lambda^2$ to the cost function (8.18), which enforces that the interpolated initial constraint (8.21) favors the feedback initialization (see Remark 7.7 for a discussion on this matter).

As we can see in Table 8.3, the performance of all considered controllers is almost identical, which is due to the strictly above rated nature of the wind with relatively low turbulence. However, this operating condition can be considered as an extreme event since average wind speeds are typically lower for wind farms, as reported in [127].

Scenario 2 – Above and below rated wind speed

The second scenario assumes a more realistic environment in which some wind turbines temporarily operate in the below rated region (Region 2, cf. Figure 8.3) due to deficiencies

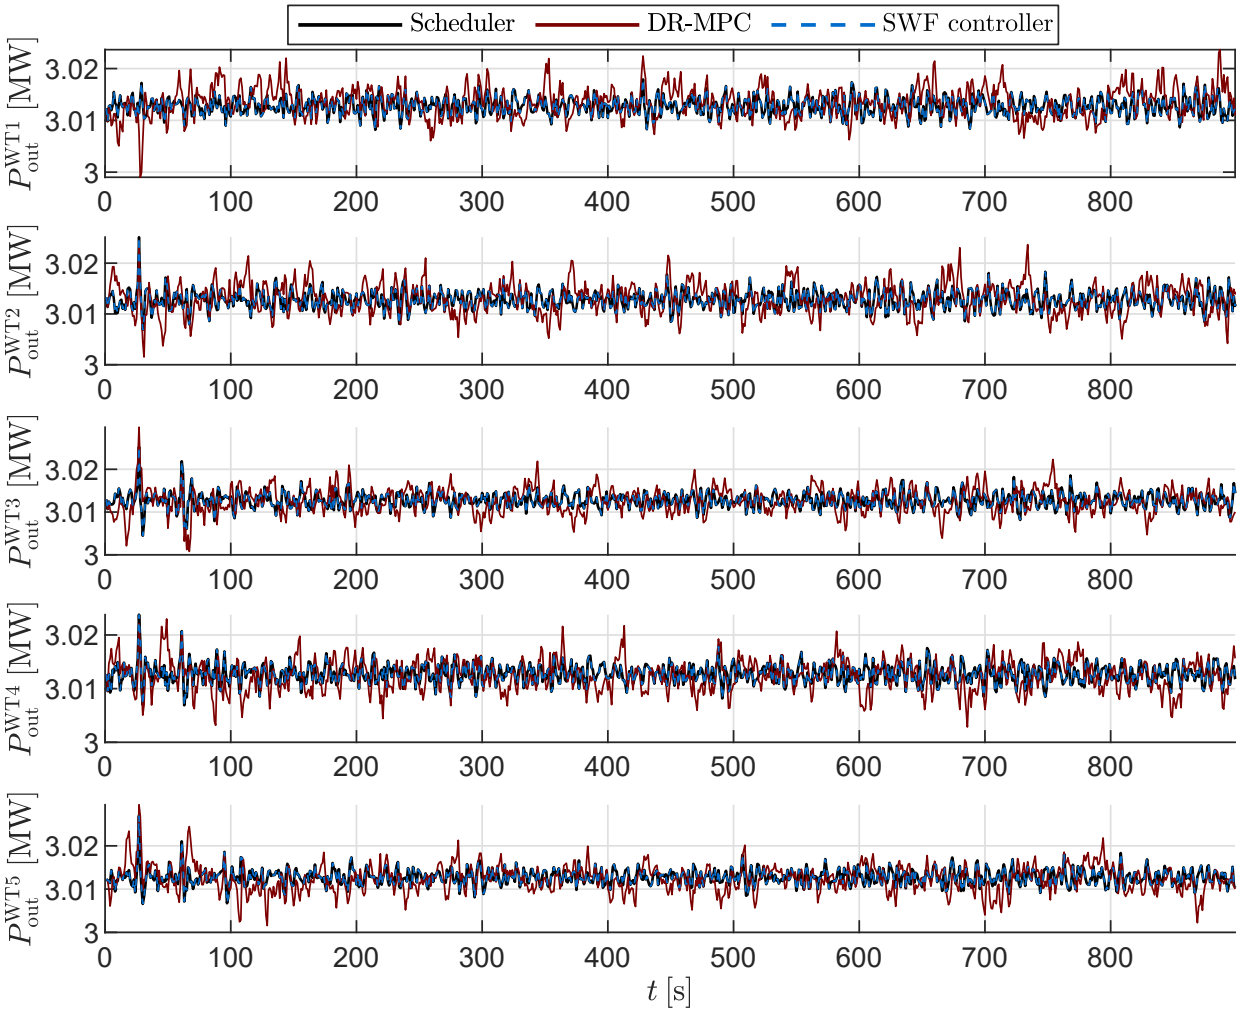

Figure 8.6: Wind turbine power output for scenario 1 with weight $r = 100$.

in wind speed. The wind field has a mean velocity of $w_0 = 12\text{m/s}$ and a turbulence intensity of $T_\text{I} = 0.1$, yielding a turbulence variance of $\sigma^2_{\Delta\text{w}} = 1.44$. The empirical covariance matrix of the ARMA residuals is given by

$$\hat{\Sigma}_\epsilon^{(12,0.1)} = \text{diag}(0.255,\ 0.270,\ 0.288,\ 0.262,\ 0.274).$$

In this scenario, we constrain the input deviations to $\pm 1\,\text{MW}$ around the nominal operating point of $3\,\text{MW}$. In this way, we can dynamically dispatch the power references depending on the available wind speed, while ensuring a power tracking goal and minimizing fatigue load, see Table 8.4 and Figure 8.7. In particular, for $r = 1$, we increase the tracking performance compared to the scheduler by approximately 52.6% and compared to the SWF controller by 4.4%. The tracking performance increase comes at the price of increasing the tower fatigue by 27.8%, while reducing the main shaft fatigue by 10%. A reasonable choice is $r = 500$, which only marginally increases the mechanical stress on the tower, while still increasing the tracking performance by nearly 34%.

Table 8.4: Performance comparison of scenario 2. The first row denotes the scheduler performance, which we consider as the baseline, i.e., 100%. The other rows denote the performance w.r.t. the scheduler, where numbers smaller than 100% denote a performance increase and numbers greater than 100% a performance decrease.

Method	J_p	J_t	J_s
Scheduler	0.0999	0.3217	0.0734
SWF controller	51.79%	131.99%	90.15%
DR-MPC $r = 1$	47.38%	127.85%	89.99%
DR-MPC $r = 500$	65.79%	105.30%	93.41%
DR-MPC $r = 10^3$	84.07%	101.88%	93.67%
DR-MPC $r = 10^4$	98.55%	100.12%	99.21%

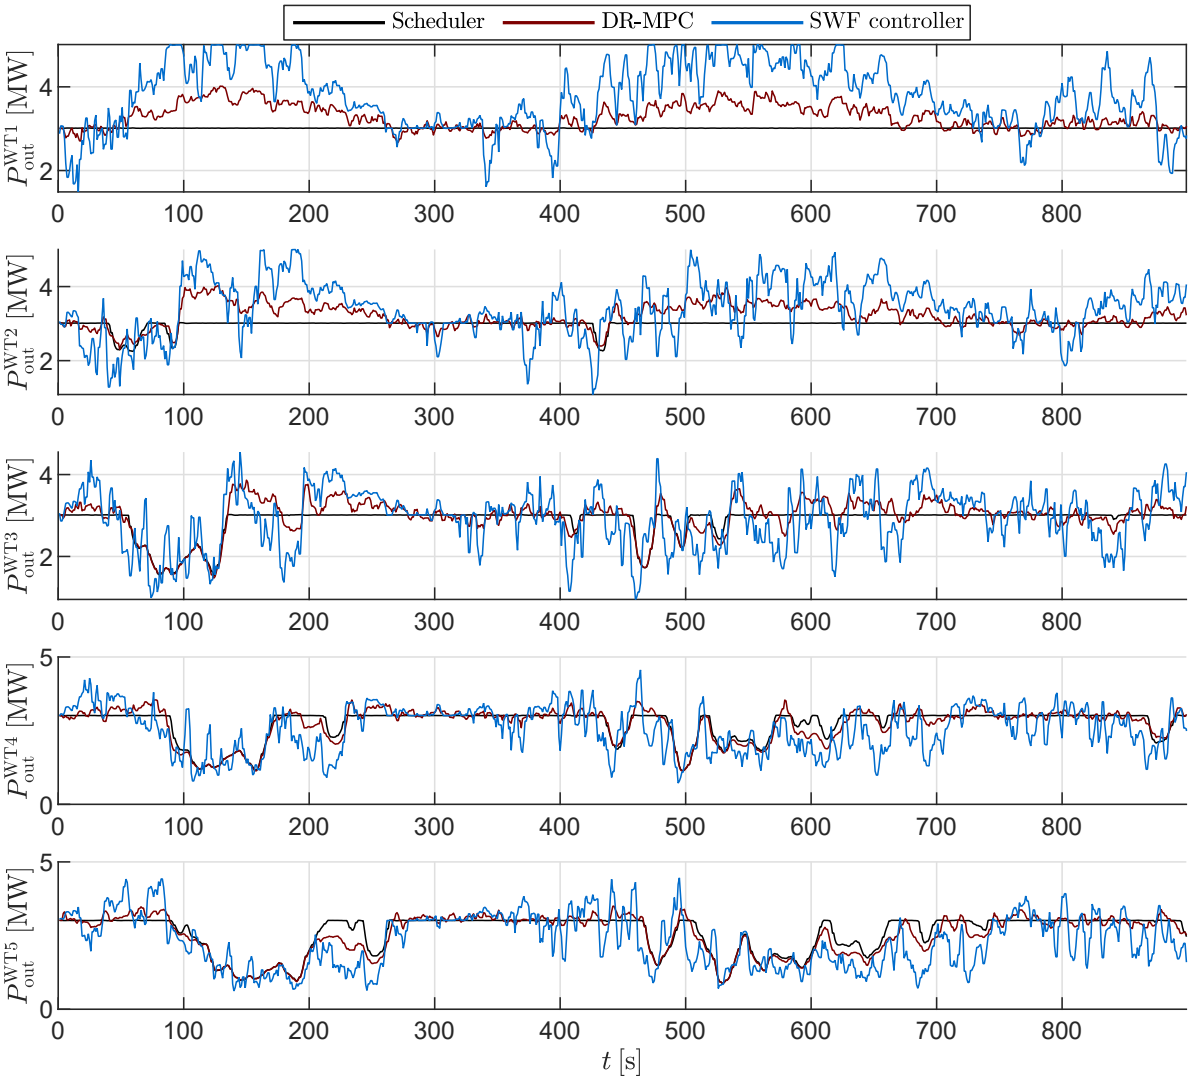

Figure 8.7: Wind turbine power output for scenario 2 with weight $r = 500$ for the DR-MPC.

8.4 Summary

In this chapter, we presented a distributionally robust MPC approach to tackle the problem of coordinating individual wind turbines inside of a wind farm. The main objective hereby was to ensure power tracking, while secondary goals were to reduce the mechanical stress acting on the tower and main shaft. We numerically verified the increase in tracking performance, as well as a reduction in mechanical stress. We considered a simple ARMA model to predict the turbulent wind speed locally for each wind turbine individually, neglecting the broader picture of spatial correlations of the wind field. This can be improved by considering a spatio-temporal wind speed forecast that includes wind measurements from neighboring turbines, e.g., as proposed by [175]. This could further increase the tracking performance, i.e., when a wind deficit is measured at an upstream turbine, it is inevitably passed on to the downstream turbines, allowing us to anticipate the temporary drop in output power. Therefore, the output of the unaffected wind turbines can be increased to compensate for the loss of output power of the others.

9 Conclusion

In the following, we summarize the main contributions of this thesis and give an outlook for future research.

9.1 Summary

Part I

In the first part of this thesis, we considered a class of distributed systems that are representable on a graph, i.e., the overall system consists of several non-overlapping subsystems with neighbor-to-neighbor coupled dynamics and constraints.

Chapter 3 We developed a distributed stochastic MPC for tracking of piece-wise constant output reference signals subject to coupled state and local input chance constraints. The contribution of this chapter is twofold. First, two practical design procedures for distributed PRS are introduced, such that the constraints can be distributedly tightened. Second, we developed an analytical approach to distributed stochastic MPC with distributed PRS constraint tightening, while the online DSMPC algorithm is based on the alternating direction method of multipliers, incorporating a simple stopping condition for practical reasons. This, however, introduces inexactness into the global optimal solution, which can compromise recursive feasibility and closed-loop chance constraint satisfaction. This issue is addressed by explicitly including the inexactness into the design phase of the MPC algorithm via robust constraint tightening. The MPC optimization problem is proven to be recursively feasible, convergent to an asymptotic average performance bound and the closed-loop system verifies the chance constraints. The properties of the resulting controller are demonstrated on two numerical examples.

Chapter 4 In this chapter, we have extended the state feedback design to the output feedback case, where we introduced two distributed stochastic MPC controllers for steady-state regulation for distributed systems, one with direct feedback (Section 4.2) and another based on indirect feedback (Section 4.3).

The first contribution is the extension of the distributed PRS design from Chapter 3 to the output feedback case, i.e., by combining the process and measurement noise into a

new random variable. In Section 4.2, a distributed direct output feedback stochastic MPC algorithm based on analytical approximations and distributed PRS constraint tightening is developed. The MPC algorithm is proven to be recursively feasible and the cost converges to an asymptotic average performance bound, while the closed-loop system verifies the chance constraints. The analytical distributed PRS design is unfortunately associated with a large conservatism, which we address in Section 4.3 with a scenario-based distributed PRS. Hereby, the chance constraints are verified with scenario optimization guarantees, i.e., the closed-loop system verifies the chance constraints with high probability. Another contribution is a distributed sample-based design procedure for the distributed PRS, where no strict assumptions on the distribution are required. The resulting indirect output feedback SMPC is proven to be recursively feasible, while under an additional zero-mean i.i.d assumption the cost converges to an asymptotic average performance bound no worse than from a linear controller. In a numerical example, we contrasted the three proposed distributed PRS designs and compared the direct and indirect DSMPC schemes.

Chapter 5 In this chapter, we studied a class of distributed systems on a graph subject to (local) individual chance constraints and multiplicative noise. The main contribution is an analytical approximation based DSMPC that uses Cantelli's inequality to approximate the local chance constraints, while the expected value cost function is analytically approximated via mean-variance dynamics. The MPC optimization problem uses LMIs to propagate the distributed covariance matrices resulting in a semidefinite program, which we solve with distributed consensus ADMM. In addition, a distributed design procedure for the MPC ingredients is provided, such that the controller can synthesized fully distributed and no central coordination node is required. The MPC optimization problem is proven to be recursively feasible, the closed-loop state to be point-wise convergent, while the chance constraints are verified in prediction for all times.

Part II

In the second part of this thesis, we studied distributionally robust MPC schemes for constrained linear stochastic systems subject to chance constraints and unknown probability distributions.

Chapter 6 In this chapter, we introduced two DR-MPC schemes with Wasserstein ambiguity sets, where we use methods from Wasserstein distributionally robust optimization to robustify the stochastic optimal control problems to distributional uncertainty.

In Section 6.2, we introduced a scenario-based DR-MPC scheme with indirect feedback, where only a potentially small set of disturbance trajectories is assumed to exist over a task horizon. As a first contribution, we investigate for nonlinear and linear tube controllers whether distributionally robust performance and/or chance constraint guarantees can be given. The resulting DR-MPC scheme is proven to be recursively feasible, regardless of

the choice of tube controller, by relying on a soft-constrained framework. In addition, the proposed framework allows for the use of hard input constraints by restricting the control authority of the tube controller. In a numerical example of a four-room temperature control task, the impact of the Wasserstein radius on closed-loop chance constraint satisfaction is shown.

In Section 6.3, an analytical indirect feedback DR-MPC scheme is developed based on distributionally robust PRS. The first contribution of this section is the extension of PRS to the distributionally robust case, where we provide for Wasserstein ambiguity sets a sample-based design procedure. Under a zero-mean i.i.d. assumption, quadratic cost and fixed tube controller, it can be shown that the distributionally robust variance cost is independent of the MPC decision variables, which allows us to omit this part in the receding horizon implementation. The resulting MPC optimization problem is trivially recursively feasible due to the indirect feedback initialization, while the closed-loop system verified the chance constraints with high probability due to the distributionally robust PRS.

Chapter 7 In the previous chapter, we considered Wasserstein ambiguity sets that use tools from Machine Learning to calibrate the ambiguity radius. In this chapter, we introduce moment-based ambiguity sets that have two main advantages over Wasserstein ambiguity sets for a class of i.i.d. sub-Gaussian random variables, i.e., (i) the resulting optimization problem does not increase in complexity with the sample size and (ii) the ambiguity radius is estimated with reasonable accuracy based on known information of the distribution via concentration inequalities. Hence, the resulting guarantees are stronger compared to Wasserstein ambiguity sets. The main contribution of this chapter is a novel DR-MPC scheme for a class of constrained stochastic systems subject to individual chance constraints on the states and inputs, which uses a simplified affine disturbance feedback parameterization to reformulate the cost and chance constraints. Recursive feasibility is established by constraining the initial state on a line between the state feedback and a feasible backup solution. The resulting closed-loop system is proven to converge to an asymptotic average performance bound no worse that from the equivalent linear quadratic regulator. A numerical example demonstrates the performance improvements for an increasing number of disturbance samples.

Chapter 8 In this chapter, a moment-based distributionally robust MPC for coordinated control of wind farms is proposed, which extends the DR-MPC formulation from Chapter 7 to support cost functions for output variables. We incorporate an ARMA prediction model that serves as an optimal turbulent wind predictor and additionally renders the stochastic residual an i.i.d. white noise. We introduce a parameterized moment-based ambiguity set for mean wind speed and turbulence intensity pairs, which uses the theoretical results from Chapter 7 to find a data-driven ambiguity radius. In a numerical study, we investigate the advantages of DR-MPC over a classical scheduler approach, comparing the effective power gain/loss for different wind scenarios.

9.2 Outlook

This thesis addressed stochastic MPC for distributed linear systems as well as distributionally robust MPC, both of which enable numerous future research directions, some of which are presented below.

Distributionally robust distributed MPC An important extension is the combination of both methods, i.e., distributionally robust distributed MPC, since in distributed stochastic systems the main problem of unavailability of exact distributional information still exists. The concept of cooperative distributionally robust distributed optimization is already known to the literature [36] and was applied successfully to dispatch problems [170]. However, the main properties of the receding horizon implementation, such as recursive feasibility and closed-loop chance constraint satisfaction, are so far not investigated. In view of Wasserstein DRO, the separation of the global ambiguity set is non-trivial and needs to be properly addressed in future research, i.e., how one can find a worst-case distribution in a distributed way to solve the underlying expectation problem.

Non-iterative DSMPC In the first part of this thesis, we used an ADMM-based DSMPC algorithm in each chapter, which has a high communication demand. This communication overhead can be reduced by resorting to an event-triggered non-iterative DMPC architecture [12], where the challenge is to incorporate a distributed k-step PRS to ensure closed-loop constraint satisfaction. Another possibility is to use DSMPC with sequential updating, similar to [47], which conceptually fits an indirect feedback scheme that always uses the shifted optimal solution.

Ambiguity set In this thesis, we considered only a small subset of all possible variations of ambiguity sets, cf. [143]. Future research directions could include different distance measures, such as the Sinkhorn distance [84], which basically enhances the Wasserstein distance by an entropic regularization [168]. In addition, the worst-case distribution of a Sinkhorn ambiguity set is continuous, even if the ambiguity set is centered at the empirical distribution [168, Remark 5]. Therefore, the question is whether we can achieve a more meaningful robustification in context of DR-MPC.

Nonlinear systems Throughout this thesis, we considered linear time-invariant systems. In view of this, the extension to nonlinear systems is desirable, but, due to several cumbersome technicalities a non-trivial task to perform, cf. Section 2.2.1. A tractable way of addressing distributionally robust MPC for nonlinear systems is given by [174], where the authors rely on linearizing the dynamics along the trajectories. In addition, recent developments in the field of stochastic MPC using incremental stability [152] enable extensions to distributionally robust formulations that potentially provide stronger closed-loop guarantees than the current state of the art.

Distributionally robust economic MPC Over the last couple of years a lot of research was done in the field of economic MPC [54], where the main idea is to use cost functions that represent an economic interest. In view of this, one can introduce a distributionally robust economic MPC framework to give robust performance certificates for the optimal solution, e.g., as in Chapter 6. As a consequence, the resulting economic steady-state (if it exists), is robust against distributional uncertainty that breaks the so-called optimizer's curse [99]. A possible application of this concept is stochastic optimal power flow, where the cost is typically economically oriented [72].

Wind farm control A potential extension of MPC-based wind farm control is the inclusion of spatio-temporal turbulence prediction [175], which could further improve the closed-loop performance of the wind farm. In particular, predicting the behavior of the wind field allows to anticipate power drops of downstream wind turbines due to wind deficits (and also wake effects).

Bibliography

[1] L. Adam and M. Branda. "Nonlinear Chance Constrained Problems: Optimality Conditions, Regularization and Solvers". In: *Journal of Optimization Theory and Applications* 170.2 (2016), pp. 419–436.

[2] T. W. Anderson. "The Integral of a Symmetric Unimodal Function over a Symmetric Convex Set and Some Probability Inequalities". In: *Proc. of the American Mathematical Society* 6.2 (1955), pp. 170–176.

[3] L. E. Andersson et al. "Wind farm control-Part I: A review on control system concepts and structures". In: *IET Renewable Power Generation* 15.10 (2021), pp. 2085–2108.

[4] M. ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 10.0.* 2022.

[5] P. Artzner et al. "Coherent Measures of Risk". In: *Mathematical Finance* 9.3 (1999), pp. 203–228.

[6] G. Banjac et al. "Infeasibility Detection in the Alternating Direction Method of Multipliers for Convex Optimization". In: *Journal of Optimization Theory and Applications* 183.2 (2019), pp. 490–519.

[7] L. Baringo and A. J. Conejo. "Correlated wind-power production and electric load scenarios for investment decisions". In: *Applied Energy* 101 (2013), pp. 475–482.

[8] R. J. Barthelmie et al. "Modelling and measurements of power losses and turbulence intensity in wind turbine wakes at Middelgrunden offshore wind farm". In: *Wind Energy* 10.6 (2007), pp. 517–528.

[9] R. J. Barthelmie et al. "Quantifying the Impact of Wind Turbine Wakes on Power Output at Offshore Wind Farms". In: *Journal of Atmospheric and Oceanic Technology* 27.8 (2010), pp. 1302–1317.

[10] I. Batina, A. Stoorvogel, and S. Weiland. "Stochastic disturbance rejection in model predictive control by randomized algorithms". In: *Proc. 2001 American Control Conf. (ACC). (Cat. No.01CH37148).* 2001, pp. 732–737.

[11] A. Bemporad et al. "The explicit linear quadratic regulator for constrained systems". In: *Automatica* 38.1 (2002), pp. 3–20.

[12] F. Berkel and S. Liu. "Non-Iterative Distributed Model Predictive Control with Event-Triggered Communication". In: *Proc. American Control Conf. (ACC)*. 2018, pp. 2344–2349.

[13] D. Bertsekas. *Convex Optimization Theory*. Athena Scientific optimization and computation series. Athena Scientific, 2009.

[14] D. Bertsekas. *Dynamic Programming and Optimal Control: Volume I*. Athena scientific optimization and computation series. Athena Scientific, 2012.

[15] G. Betti, M. Farina, and R. Scattolini. "Realization issues, tuning, and testing of a distributed predictive control algorithm". In: *Journal of Process Control* 24.4 (2014), pp. 424–434.

[16] P. Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics. Wiley, 2013.

[17] S. Boersma et al. "Stochastic Model Predictive Control: uncertainty impact on wind farm power tracking". In: *Proc. American Control Conf. (ACC)*. 2019, pp. 4167–4172.

[18] J. Bossuyt et al. "Measurement of unsteady loading and power output variability in a micro wind farm model in a wind tunnel". In: *Experiments in Fluids* 58.1 (2017), pp. 1–17.

[19] G. Box et al. *Time Series Analysis: Forecasting and Control*. Wiley Series in Probability and Statistics. Wiley, 2015.

[20] S. Boyd et al. *Linear Matrix Inequalities in System and Control Theory*. Studies in Applied Mathematics. SIAM, 1994.

[21] S. Boyd et al. *Convex Optimization*. Cambridge University Press, 2004.

[22] S. Boyd et al. "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers". In: *Foundations and Trends® in Machine Learning* 3.1 (2011), pp. 1–122.

[23] E. Bradford and L. Imsland. "Stochastic Nonlinear Model Predictive Control Using Gaussian Processes". In: *Proc. European Control Conf. (ECC)*. 2018, pp. 1027–1034.

[24] E. A. Buehler, J. A. Paulson, and A. Mesbah. "Lyapunov-based stochastic nonlinear model predictive control: Shaping the state probability distribution functions". In: *Proc. American Control Conf. (ACC)*. 2016, pp. 5389–5394.

[25] T. Burton et al. *Wind Energy Handbook*. Wiley, 2011.

[26] G. C. Calafiore and L. E. Ghaoui. "On Distributionally Robust Chance-Constrained Linear Programs". In: *Journal of Optimization Theory and Applications* 130.1 (2006), pp. 1–22.

[27] G. C. Calafiore and L. Fagiano. "Robust Model Predictive Control via Scenario Optimization". In: *IEEE Transactions on Automatic Control* 58.1 (2013), pp. 219–224.

[28] M. C. Campi and S. Garatti. "A Sampling-and-Discarding Approach to Chance-Constrained Optimization: Feasibility and Optimality". In: *Journal of Optimization Theory and Applications* 148.2 (2011), pp. 257–280.

[29] M. Cannon, B. Kouvaritakis, and D. Ng. "Probabilistic tubes in linear stochastic model predictive control". In: *Systems & Control Letters* 58.10 (2009), pp. 747–753.

[30] M. Cannon, B. Kouvaritakis, and X. Wu. "Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints". In: *Automatica* 45.1 (2009), pp. 167–172.

[31] M. Cannon, B. Kouvaritakis, and X. Wu. "Probabilistic Constrained MPC for Multiplicative and Additive Stochastic Uncertainty". In: *IEEE Transactions on Automatic Control* 54.7 (2009), pp. 1626–1632.

[32] M. Cannon et al. "Stochastic Tubes in Model Predictive Control With Probabilistic Constraints". In: *IEEE Transactions on Automatic Control* 56.1 (2011), pp. 194–200.

[33] M. Cannon et al. "Stochastic tube MPC with state estimation". In: *Automatica* 48.3 (2012), pp. 536–541.

[34] X. Chen. "A New Generalization of Chebyshev Inequality for Random Vectors". In: *arXiv preprint arXiv:0707.0805* (2007).

[35] Z. Chen, M. Sim, and H. Xu. "Distributionally Robust Optimization with Infinitely Constrained Ambiguity Sets". In: *Operations Research* 67.5 (2019), pp. 1328–1344.

[36] A. Cherukuri and J. Cortés. "Cooperative Data-Driven Distributionally Robust Optimization". In: *IEEE Transactions on Automatic Control* 65.10 (2020), pp. 4400–4407.

[37] L. Chisci, J. A. Rossiter, and G. Zappa. "Systems with persistent disturbances: predictive control with restricted constraints". In: *Automatica* 37.7 (2001), pp. 1019–1028.

[38] P. D. Christofides et al. "Distributed model predictive control: A tutorial review and future research directions". In: *Computers & Chemical Engineering* 51 (2013), pp. 21–41.

[39] C. Conte et al. "Computational aspects of distributed optimization in model predictive control". In: *Proc. 51st IEEE Conf. on Decision and Control (CDC)*. 2012, pp. 6819–6824.

[40] C. Conte et al. "Cooperative distributed tracking MPC for constrained linear systems: Theory and synthesis". In: *Proc. 52nd IEEE Conf. on Decision and Control (CDC)*. 2013, pp. 3812–3817.

[41] C. Conte et al. “Robust distributed model predictive control of linear systems”. In: *Proc. European Control Conf. (ECC)*. 2013, pp. 2764–2769.

[42] C. Conte et al. “Distributed synthesis and stability of cooperative distributed model predictive control for linear systems”. In: *Automatica* 69 (2016), pp. 117–125.

[43] P. Coppens and P. Patrinos. “Data-Driven Distributionally Robust MPC for Constrained Stochastic Systems”. In: *IEEE Control Systems Letters* 6 (2022), pp. 1274–1279.

[44] P. Coppens, M. Schuurmans, and P. Patrinos. “Data-driven distributionally robust LQR with multiplicative noise”. In: *Proc. 2nd Conf. on Learning for Dynamics and Control*. PMLR, 2020, pp. 521–530.

[45] J. Coulson, J. Lygeros, and F. Dörfler. “Distributionally Robust Chance Constrained Data-Enabled Predictive Control”. In: *IEEE Transactions on Automatic Control* 67.7 (2022), pp. 3289–3304.

[46] L. Dai et al. “Cooperative distributed stochastic MPC for systems with state estimation and coupled probabilistic constraints”. In: *Automatica* 61 (2015), pp. 89–96.

[47] L. Dai et al. “Distributed Stochastic MPC of Linear Systems With Additive Uncertainty and Coupled Probabilistic Constraints”. In: *IEEE Transactions on Automatic Control* 62.7 (2017), pp. 3474–3481.

[48] L. Dai et al. “Distributed stochastic MPC for systems with parameter uncertainty and disturbances”. In: *International Journal of Robust and Nonlinear Control* 28.6 (2018), pp. 2424–2441.

[49] E. Delage and Y. Ye. “Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems”. In: *Operations Research* 58.3 (2010), pp. 595–612.

[50] S. Dharmadhikari and K. Joag-Dev. *Unimodality, Convexity, and Applications*. Elsevier Science, 1988.

[51] S. W. Dharmadhikari and K. Jogdeo. “Multivariate Unimodality”. In: *The Annals of Statistics* 4.3 (1976), pp. 607–613.

[52] A. Dixit, M. Ahmadi, and J. W. Burdick. “Distributionally Robust Model Predictive Control With Total Variation Distance”. In: *IEEE Control Systems Letters* 6 (2022), pp. 3325–3330.

[53] L. El Ghaoui. “State-feedback control of systems with multiplicative noise via linear matrix inequalities”. In: *Systems & Control Letters* 24.3 (1995), pp. 223–228.

[54] M. Ellis, H. Durand, and P. D. Christofides. “A tutorial review of economic model predictive control methods”. In: *Journal of Process Control* 24.8 (2014), pp. 1156–1178.

[55] M. A. Evans, M. Cannon, and B. Kouvaritakis. "Robust MPC Tower Damping for Variable Speed Wind Turbines". In: *IEEE Transactions on Control Systems Technology* 23.1 (2015), pp. 290–296.

[56] M. Farina, L. Giulioni, and R. Scattolini. "Distributed Predictive Control of stochastic linear systems with chance constraints". In: *Proc. American Control Conf. (ACC)*. 2016, pp. 20–25.

[57] M. Farina, L. Giulioni, and R. Scattolini. "Stochastic linear Model Predictive Control with chance constraints – A review". In: *Journal of Process Control* 44 (2016), pp. 53–67.

[58] M. Farina and S. Misiano. "Stochastic Distributed Predictive Tracking Control for Networks of Autonomous Systems With Coupling Constraints". In: *IEEE Transactions on Control of Network Systems* 5.3 (2018), pp. 1412–1423.

[59] M. Farina and R. Scattolini. "Model predictive control of linear systems with multiplicative unbounded uncertainty and chance constraints". In: *Automatica* 70 (2016), pp. 258–265.

[60] M. Farina et al. "A probabilistic approach to Model Predictive Control". In: *Proc. 52nd IEEE Conf. on Decision and Control (CDC)*. 2013, pp. 7734–7739.

[61] M. Farina et al. "An approach to output-feedback MPC of stochastic linear discrete-time systems". In: *Automatica* 55 (2015), pp. 140–149.

[62] A. Ferramosca et al. "Cooperative distributed MPC for tracking". In: *Automatica* 49.4 (2013), pp. 906–914.

[63] P. Florchinger. "Lyapunov-Like Techniques for Stochastic Stability". In: *SIAM Journal on Control and Optimization* 33.4 (1995), pp. 1151–1169.

[64] A. Geletu et al. "Monotony analysis and sparse-grid integration for nonlinear chance constrained process optimization". In: *Engineering Optimization* 43.10 (2011), pp. 1019–1041.

[65] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski. "Optimization over state feedback policies for robust control with constraints". In: *Automatica* 42.4 (2006), pp. 523–533.

[66] S. Grammatico, A. Subbaraman, and A. R. Teel. "Discrete-time stochastic control systems: A continuous Lyapunov function implies robustness to strictly causal perturbations". In: *Automatica* 49.10 (2013), pp. 2939–2952.

[67] M. Grant, S. Boyd, and Y. Ye. *CVX: Matlab software for disciplined convex programming*. 2008.

[68] B. Gravell, P. M. Esfahani, and T. Summers. "Learning Optimal Controllers for Linear Systems With Multiplicative Noise via Policy Gradient". In: *IEEE Transactions on Automatic Control* 66.11 (2021), pp. 5283–5298.

[69] S. Gros and A. Schild. "Real-time economic nonlinear model predictive control for wind turbine control". In: *International Journal of Control* 90.12 (2017), pp. 2799–2812.

[70] L. Grüne and J. Pannek. *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer International Publishing, 2016.

[71] J. D. Grunnet et al. "Aeolus Toolbox for Dynamics Wind Farm Model, Simulation and Control". In: *Proc. European Wind Energy Conf. and Exhibition (EWEC)*. 2010.

[72] Y. Guo et al. "Data-Based Distributionally Robust Stochastic Optimal Power Flow–Part I: Methodologies". In: *IEEE Transactions on Power Systems* 34.2 (2019), pp. 1483–1492.

[73] A. Hakobyan and I. Yang. "Learning-Based Distributionally Robust Motion Control with Gaussian Processes". In: *Proc. 2020 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2020, pp. 7667–7674.

[74] A. Hakobyan and I. Yang. "Toward Improving the Distributional Robustness of Risk-Aware Controllers in Learning-Enabled Environments". In: *Proc. 60th IEEE Conf. on Decision and Control (CDC)*. 2021, pp. 6024–6031.

[75] A. Hakobyan and I. Yang. "Wasserstein Distributionally Robust Motion Control for Collision Avoidance Using Conditional Value-at-Risk". In: *IEEE Transactions on Robotics* 38.2 (2022), pp. 939–957.

[76] L. Hewing, K. P. Wabersich, and M. N. Zeilinger. "Recursively feasible stochastic model predictive control using indirect feedback". In: *Automatica* 119 (2020), p. 109095.

[77] L. Hewing and M. N. Zeilinger. "Stochastic Model Predictive Control for Linear Systems Using Probabilistic Reachable Sets". In: *Proc. 57th IEEE Conf. on Decision and Control (CDC)*. 2018, pp. 5182–5188.

[78] L. Hewing and M. N. Zeilinger. "Performance Analysis of Stochastic Model Predictive Control with Direct and Indirect Feedback". In: *Proc. 59th IEEE Conf. on Decision and Control (CDC)*. 2020, pp. 672–678.

[79] L. Hewing and M. N. Zeilinger. "Scenario-Based Probabilistic Reachable Sets for Recursively Feasible Stochastic Model Predictive Control". In: *IEEE Control Systems Letters* 4.2 (2020), pp. 450–455.

[80] A. R. Hota, A. Cherukuri, and J. Lygeros. "Data-Driven Chance Constrained Optimization under Wasserstein Ambiguity Sets". In: *Proc. American Control Conf. (ACC)*. 2019, pp. 1501–1506.

[81] D. Hsu, S. Kakade, and T. Zhang. "Tail inequalities for sums of random matrices that depend on the intrinsic dimension". In: *Electronic Communications in Probability* 17 (2012), pp. 1–13.

[82] T. Hu, Z. Lin, and B. M. Chen. "Analysis and design for discrete-time linear systems subject to actuator saturation". In: *Systems & Control Letters* 45.2 (2002), pp. 97–112.

[83] Z. Hu and L. J. Hong. "Kullback-Leibler Divergence Constrained Distributionally Robust Optimization". In: *Available at Optimization Online* 1.2 (2013), p. 9.

[84] K. Ito and K. Kashima. "Sinkhorn MPC: Model predictive optimal transport over dynamical systems". In: *Proc. American Control Conf. (ACC)*. 2022, pp. 2057–2062.

[85] J. Jin and Y. Xu. "Segregated Linear Decision Rules for Distributionally Robust Control With Linear Dynamics and Quadratic Cost". In: *IEEE Systems Journal* 15.1 (2021), pp. 355–364.

[86] K. Johansson and J. Nunes. "A multivariable laboratory process with an adjustable zero". In: *Proc. American Control Conf. (ACC). (Cat. No.98CH36207)*. Vol. 4. 1998, pp. 2045–2049.

[87] J. Jonkman et al. *Definition of a 5-MW Reference Wind Turbine for Offshore System Development*. Tech. rep. NREL/TP-500-38060. National Renewable Energy Lab. (NREL), Golden, CO (United States), 2009.

[88] E. C. Kerrigan and J. M. Maciejowski. "Soft constraints and exact penalty functions in model predictive control". In: *Proc. Control 2000 Conf., Cambridge*. Citeseer. 2000, pp. 2319–2327.

[89] A. Klenke. *Probability Theory: A Comprehensive Course*. Springer, 2013.

[90] A. J. Kleywegt, A. Shapiro, and T. Homem-de-Mello. "The Sample Average Approximation Method for Stochastic Discrete Optimization". In: *SIAM Journal on Optimization* 12.2 (2002), pp. 479–502.

[91] T. Knudsen, T. Bak, and M. Svenstrup. "Survey of wind farm control – power and fatigue optimization". In: *Wind Energy* 18.8 (2015), pp. 1333–1351.

[92] J. Köhler, M. A. Müller, and F. Allgöwer. "Distributed model predictive control–Recursive feasibility under inexact dual optimization". In: *Automatica* 102 (2019), pp. 1–9.

[93] J. Köhler and M. N. Zeilinger. "Recursively Feasible Stochastic Predictive Control Using an Interpolating Initial State Constraint". In: *IEEE Control Systems Letters* 6 (2022), pp. 2743–2748.

[94] J. Köhler et al. "Real time economic dispatch for power networks: A distributed economic model predictive control approach". In: *Proc. 56th IEEE Conf. on Decision and Control (CDC)*. 2017, pp. 6340–6345.

[95] J. Köhler et al. "A Computationally Efficient Robust Model Predictive Control Framework for Uncertain Nonlinear Systems". In: *IEEE Transactions on Automatic Control* 66.2 (2021), pp. 794–801.

[96] M. Korda et al. “Strongly feasible stochastic model predictive control”. In: *Proc. 50th IEEE Conf. on Decision and Control and European Control Conf.* 2011, pp. 1245–1251.

[97] B. Kouvaritakis and M. Cannon. *Model Predictive Control*. Springer International Publishing, 2016.

[98] B. Kouvaritakis et al. “Explicit use of probabilistic distributions in linear predictive control”. In: *Automatica* 46.10 (2010), pp. 1719–1724.

[99] D. Kuhn et al. “Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning”. In: *Operations Research & Management Science in the Age of Analytics*. INFORMS, 2019, pp. 130–166.

[100] H. J. Kushner. *Stochastic Stability and Control*. Elsevier Science, 1967.

[101] N. Lefebure et al. “Distributed model predictive control of buildings and energy hubs”. In: *Energy and Buildings* 259 (2022), p. 111806.

[102] S. Leirens et al. “Coordination in urban water supply networks using distributed model predictive control”. In: *Proc. American Control Conf. (ACC)*. 2010, pp. 3957–3962.

[103] B. Li et al. “A Distributionally Robust Optimization Based Method for Stochastic Model Predictive Control”. In: *IEEE Transactions on Automatic Control* 67.11 (2022), pp. 5762–5776.

[104] D. Limón et al. “MPC for tracking piecewise constant references for constrained linear systems”. In: *Automatica* 44.9 (2008), pp. 2382–2387.

[105] J. Lofberg. “YALMIP : a toolbox for modeling and optimization in MATLAB”. In: *Proc. IEEE Int. Conf. on Robotics and Automation (IEEE Cat. No.04CH37508)*. 2004, pp. 284–289.

[106] M. Lorenzen et al. “Constraint-Tightening and Stability in Stochastic Model Predictive Control”. In: *IEEE Transactions on Automatic Control* 62.7 (2017), pp. 3165–3177.

[107] M. Lorenzen et al. “Stochastic MPC with offline uncertainty sampling”. In: *Automatica* 81 (2017), pp. 176–183.

[108] S. Lu, J. H. Lee, and F. You. “Soft-constrained model predictive control based on data-driven distributionally robust optimization”. In: *AIChE Journal* 66.10 (2020), e16546.

[109] C. Mark and S. Liu. “Distributed Stochastic Model Predictive Control for dynamically coupled Linear Systems using Probabilistic Reachable Sets”. In: *Proc. European Control Conf. (ECC)*. 2019, pp. 1362–1367.

[110] C. Mark and S. Liu. "A stochastic output-feedback MPC scheme for distributed systems". In: *Proc. American Control Conf. (ACC)*. extended version: arXiv:2001.10838. 2020, pp. 1937–1942.

[111] C. Mark and S. Liu. "Stochastic MPC with Distributionally Robust Chance Constraints". In: *Proc. 21st IFAC World Congress*. extended version: arXiv:2005.00313. 2020, pp. 7136–7141.

[112] C. Mark and S. Liu. "Data-driven distributionally robust model predictive control: An indirect feedback approach". In: *arXiv preprint arXiv:2109.09558* (2021). Submitted to International Journal of Robust and Nonlinear Control (2022).

[113] C. Mark and S. Liu. "Stochastic Distributed Predictive Tracking Control Under Inexact Minimization". In: *IEEE Transactions on Control of Network Systems* 8.4 (2021), pp. 1892–1904.

[114] C. Mark and S. Liu. "Stochastic Model Predictive Control for tracking of distributed linear systems with additive uncertainty". In: *Proc. European Control Conf. (ECC)*. extended version: arXiv:2103.01087. 2021, pp. 216–221.

[115] C. Mark and S. Liu. "A stochastic MPC scheme for distributed systems with multiplicative uncertainty". In: *Automatica* 140 (2022), p. 110208.

[116] C. Mark and S. Liu. "Distributionally robust model predictive control for wind farms". In: *arXiv preprint arXiv:2303.03276* (2023). Accepted for presentation at the 22nd IFAC World Congress.

[117] C. Mark and S. Liu. "Recursively Feasible Data-Driven Distributionally Robust Model Predictive Control With Additive Disturbances". In: *IEEE Control Systems Letters* 7 (2023), pp. 526–531.

[118] D. Q. Mayne, M. M. Seron, and S. V. Raković. "Robust model predictive control of constrained linear systems with bounded disturbances". In: *Automatica* 41.2 (2005), pp. 219–224.

[119] D. Q. Mayne et al. "Constrained model predictive control: Stability and optimality". In: *Automatica* 36.6 (2000), pp. 789–814.

[120] R. D. McAllister and J. B. Rawlings. "Nonlinear Stochastic Model Predictive Control: Existence, Measurability, and Stochastic Asymptotic Stability". In: *IEEE Transactions on Automatic Control* 68.3 (2023), pp. 1524–1536.

[121] A. Mesbah. "Stochastic Model Predictive Control: An Overview and Perspectives for Future Research". In: *IEEE Control Systems Magazine* 36.6 (2016), pp. 30–44.

[122] A. Mesbah et al. "Stochastic nonlinear model predictive control with probabilistic constraints". In: *Proc. American Control Conf. (ACC)*. 2014, pp. 2413–2419.

[123] F. Micheli and J. Lygeros. "Scenario-based Stochastic MPC for systems with uncertain dynamics". In: *Proc. European Control Conf. (ECC)*. 2022, pp. 833–838.

[124] F. Micheli, T. Summers, and J. Lygeros. "Data-driven distributionally robust MPC for systems with uncertain dynamics". In: *Proc. 61st IEEE Conf. on Decision and Control (CDC)*. 2022, pp. 4788–4793.

[125] P. Milan, M. Wächter, and J. Peinke. "Stochastic modeling and performance monitoring of wind farm power production". In: *Journal of Renewable and Sustainable Energy* 6.3 (2014), p. 033119.

[126] P. Mohajerin Esfahani and D. Kuhn. "Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations". In: *Mathematical Programming* 171.1 (2018), pp. 115–166.

[127] E. C. Morgan et al. "Probability distributions for offshore wind speeds". In: *Energy Conversion and Management* 52.1 (2011), pp. 15–26.

[128] S. Muntwiler et al. "Data-Driven Distributed Stochastic Model Predictive Control with Closed-Loop Chance Constraint Satisfaction". In: *Proc. European Control Conf. (ECC)*. 2021, pp. 210–215.

[129] K. Murphy. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.

[130] R. Negenborn and J. Maestre. "Distributed Model Predictive Control: An Overview and Roadmap of Future Research Opportunities". In: *IEEE Control Systems Magazine* 34.4 (2014), pp. 87–97.

[131] R. R. Negenborn et al. "Distributed model predictive control of irrigation canals". In: *Networks and Heterogeneous Media* 4.2 (2009), pp. 359–380.

[132] A. Nemirovski and A. Shapiro. "Convex Approximations of Chance Constrained Programs". In: *SIAM Journal on Optimization* 17.4 (2007), pp. 969–996.

[133] C. Ning and F. You. "Online learning based risk-averse stochastic MPC of constrained linear uncertain systems". In: *Automatica* 125 (2021), p. 109402.

[134] C. Ning and F. You. "Deep Learning Based Distributionally Robust Joint Chance Constrained Economic Dispatch Under Wind Power Uncertainty". In: *IEEE Transactions on Power Systems* 37.1 (2022), pp. 191–203.

[135] F. Oldewurtel et al. "Use of model predictive control and weather forecasts for energy efficient building climate control". In: *Energy and Buildings* 45 (2012), pp. 15–27.

[136] M. Ono et al. "Risk-limiting power grid control with an ARMA-based prediction model". In: *Proc. 52nd IEEE Conf. on Decision and Control*. 2013, pp. 4949–4956.

[137] J. A. Paulson et al. "Stochastic model predictive control with joint chance constraints". In: *International Journal of Control* 93.1 (2020), pp. 126–139.

[138] A. Pichler and A. Shapiro. "Mathematical Foundations of Distributionally Robust Multistage Optimization". In: *SIAM Journal on Optimization* 31.4 (2021), pp. 3044–3067.

[139] M. Prandini, S. Garatti, and J. Lygeros. "A randomized approach to Stochastic Model Predictive Control". In: *Proc. 51st IEEE Conf. on Decision and Control (CDC)*. 2012, pp. 7315–7320.

[140] J. A. Primbs and C. H. Sung. "Stochastic Receding Horizon Control of Constrained Linear Systems With State and Control Multiplicative Noise". In: *IEEE Transactions on Automatic Control* 54.2 (2009), pp. 221–230.

[141] S. A. P. Quintero, D. A. Copp, and J. P. Hespanha. "Robust UAV coordination for target tracking using output-feedback model predictive control with moving horizon estimation". In: *Proc. American Control Conf. (ACC)*. 2015, pp. 3758–3764.

[142] H. Rahimian. "Risk-Averse and Distributionally Robust Optimization: Methodology and Applications". PhD thesis. The Ohio State University, 2018.

[143] H. Rahimian and S. Mehrotra. "Frameworks and Results in Distributionally Robust Optimization". In: *Open Journal of Mathematical Optimization* 3 (2022), pp. 1–85.

[144] J. Rawlings and D. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Pub., 2009.

[145] J. B. Rawlings, D. Angeli, and C. N. Bates. "Fundamentals of economic model predictive control". In: *Proc. 51st IEEE Conf. on Decision and Control (CDC)*. 2012, pp. 3851–3861.

[146] S. Riverso et al. "Model Predictive Controllers for Reduction of Mechanical Fatigue in Wind Farms". In: *IEEE Transactions on Control Systems Technology* 25.2 (2017), pp. 535–549.

[147] R. T. Rockafellar, S. Uryasev, et al. "Optimization of conditional value-at-risk". In: *Journal of Risk* 2.3 (2000), pp. 21–41.

[148] V. Rostampour and T. Keviczky. "Distributed Stochastic Model Predictive Control Synthesis for Large-Scale Uncertain Linear Systems". In: *Proc. American Control Conf. (ACC)*. 2018, pp. 2071–2077.

[149] T. L. Santos et al. "A Constraint-Tightening Approach to Nonlinear Model Predictive Control with Chance Constraints for Stochastic Systems". In: *Proc. American Control Conf. (ACC)*. 2019, pp. 1641–1647.

[150] G. Schildbach et al. "Randomized Model Predictive Control for stochastic linear systems". In: *Proc. American Control Conf. (ACC)*. 2012, pp. 417–422.

[151] G. Schildbach et al. "The scenario approach for Stochastic Model Predictive Control with bounds on closed-loop constraint violations". In: *Automatica* 50.12 (2014), pp. 3009–3018.

[152] H. Schlüter and F. Allgöwer. "A Constraint-Tightening Approach to Nonlinear Stochastic Model Predictive Control under General Bounded Disturbances". In: *Proc. 21st IFAC World Congress*. 2020, pp. 7130–7135.

[153] H. Schlüter and F. Allgöwer. "Stochastic Model Predictive Control using Initial State Optimization". In: *Proc. 25th International Symposium on Mathematical Theory of Networks and Systems (MTNS)*. 2022, pp. 454–459.

[154] M. Schuurmans and P. Patrinos. "Learning-Based Distributionally Robust Model Predictive Control of Markovian Switching Systems with Guaranteed Stability and Recursive Feasibility". In: *Proc. 59th IEEE Conf. on Decision and Control (CDC)*. 2020, pp. 4287–4292.

[155] M. Schuurmans et al. "Safe, Learning-Based MPC for Highway Driving under Lane-Change Uncertainty: A Distributionally Robust Approach". In: *arXiv preprint arXiv:2206.13319* (2022).

[156] A. Shapiro. "Tutorial on risk neutral, distributionally robust and risk averse multistage stochastic programming". In: *European Journal of Operational Research* 288.1 (2021), pp. 1–13.

[157] A. Shapiro, D. Dentcheva, and A. Ruszczynski. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2021.

[158] V. Spudić et al. *Aeolus project-Deliverable 3.3: Reconfigurable control extension*. Tech. rep. Technical report, University of Zagreb, 2010.

[159] V. Spudić. "Coordinated optimal control of wind farm active power". PhD thesis. Dept. Control Comput. Eng., University of Zagreb, Zagreb, Croatia, 2012.

[160] V. Spudić, M. Jelavić, and M. Baotić. "Wind Turbine Power References in Coordinated Control of Wind Farms". In: *Automatika* 52.2 (2011), pp. 82–94.

[161] V. Spudić et al. "Cooperative distributed model predictive control for wind farms". In: *Optimal Control Applications and Methods* 36.3 (2015), pp. 333–352.

[162] Y. Tan et al. "A distributionally robust optimization approach to two-sided chance constrained stochastic model predictive control with unknown noise distribution". In: *arXiv preprint arXiv:2203.08457* (2022).

[163] C. Tang and T. Basar. "Stochastic stability of singularly perturbed nonlinear systems". In: *Proc. 40th IEEE Conf. on Decision and Control (Cat. No.01CH37228)*. Vol. 1. 2001, pp. 399–404.

[164] B. P. G. Van Parys et al. "Distributionally Robust Control of Constrained Stochastic Systems". In: *IEEE Transactions on Automatic Control* 61.2 (2016), pp. 430–442.

[165] A. N. Venkat et al. "Distributed Output Feedback MPC for Power System Control". In: *Proc. 45th IEEE Conf. on Decision and Control*. 2006, pp. 4038–4045.

[166] A. N. Venkat et al. "Distributed MPC Strategies With Application to Power System Automatic Generation Control". In: *IEEE Transactions on Control Systems Technology* 16.6 (2008), pp. 1192–1206.

[167] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018.

[168] J. Wang, R. Gao, and Y. Xie. "Sinkhorn Distributionally Robust Optimization". In: *arXiv preprint arXiv:2109.11926* (2021).

[169] M. N. Zeilinger, M. Morari, and C. N. Jones. "Soft Constrained Model Predictive Control With Robust Stability Guarantees". In: *IEEE Transactions on Automatic Control* 59.5 (2014), pp. 1190–1202.

[170] J. Zhai et al. "Distributionally Robust Joint Chance-Constrained Dispatch for Integrated Transmission-Distribution Systems via Distributed Optimization". In: *IEEE Transactions on Smart Grid* 13.3 (2022), pp. 2132–2147.

[171] J. Zhang and T. Ohtsuka. "Stochastic Model Predictive Control Using Simplified Affine Disturbance Feedback for Chance-Constrained Systems". In: *IEEE Control Systems Letters* 5.5 (2021), pp. 1633–1638.

[172] Y. Zheng et al. "Distributed Model Predictive Control for Heterogeneous Vehicle Platoons Under Unidirectional Topologies". In: *IEEE Transactions on Control Systems Technology* 25.3 (2017), pp. 899–910.

[173] Z. Zhong, E. A. del Rio-Chanona, and P. Petsagkourakis. "Data-driven distributionally robust MPC using the Wasserstein metric". In: *arXiv preprint arXiv:2105.08414* (2021).

[174] Z. Zhong, E. A. del Rio-Chanona, and P. Petsagkourakis. "Distributionally Robust MPC for Nonlinear Systems". In: *Proc. 13th IFAC Symposium on Dynamics and Control of Process Systems, including Biosystems DYCOPS*. 2022, pp. 606–613.

[175] Q. Zhu et al. "Wind Speed Prediction with Spatio-Temporal Correlation: A Deep Learning Approach". In: *Energies* 11.4 (2018), p. 705.

[176] A. Zolanvari and A. Cherukuri. "Data-driven distributionally robust iterative risk-constrained model predictive control". In: *Proc. European Control Conf. (ECC)*. 2022, pp. 1578–1583.

[177] S. Zymler, D. Kuhn, and B. Rustem. "Distributionally robust joint chance constraints with second-order moment information". In: *Mathematical Programming* 137.1 (2013), pp. 167–198.

Deutsche Kurzfassung

Teil I

Im ersten Teil dieser Dissertation wird eine Klasse von verteilten Systemen betrachtet, die auf einem Graphen darstellbar sind, d.h., das Gesamtsystem besteht aus mehreren nicht überlappenden Teilsystemen mit einer von Nachbar-zu-Nachbar gekoppelten Dynamik. Es werden ausschließlich iterative, kooperative und parallele verteilte modellprädiktive Regler (engl.: model predictive control, MPC) untersucht, wobei die folgenden Forschungsziele berücksichtigt werden:

(i) Quantifizierung der Auswirkungen von additiver und multiplikativer Unsicherheit auf verteilte Weise.

(ii) Der Regler sollte verteilt synthetisierbar und im Online-Betrieb sollte kein zentraler Koordinationsknoten erforderlich sein.

Punkt (i) wird in den Kapiteln 3 und 4 mit dem Konzept der verteilten probabilistisch erreichbaren Mengen (engl.: probabilistic reachable set, PRS) umgesetzt. In Kapitel 5 wird eine ähnliche Methodik verwendet, wobei die verteilte Kovarianz Matrix des prädizierten Zustandsvektors mittels linearen Matrixungleichungen (engl.: linear matrix inequality, LMI) propagiert wird. Um den verteilten online Betrieb nach Punkt (ii) zu gewährleisten, wird in jedem vorgeschlagenen MPC Regler ein verteilter Konsensus ADMM Algorithmus verwendet.

Kapitel 3 In diesem Kapitel wird eine verteilte stochastische MPC zur Verfolgung von stückweise konstanten Ausgangsreferenzsignalen entwickelt, die gekoppelten Zustands- und lokalen Zufallsbeschränkungen (engl.: chance constraints) unterliegen. Zuerst werden zwei praktische Entwurfsverfahren für verteilte PRS vorgestellt, so dass die gekoppelten Zustandsbeschränkungen mittels eines verteilten Verfahrens verschärft werden können (engl. constraint tightening). Zweitens wird ein analytischer Ansatz für die verteilte stochastische MPC mit verteilten PRS entwickelt, während der Online MPC Algorithmus auf der ADMM Methode basiert und aus praktischen Gründen eine einfache Stoppbedingung enthält. Dies führt jedoch zu Ungenauigkeiten in der globalen optimalen Lösung, was die rekursive Machbarkeit und die Erfüllung der Zufallsbedingungen im geschlossenen Regelkreis beeinträchtigt. Dieses Problem wird durch die explizite Einbeziehung der Ungenauigkeit in die Entwurfsphase des MPC Optimierungsproblems über eine robuste Verschärfung der nominellen Beschränkungen gelöst. Das MPC Optimierungsproblem ist nachweislich rekur-

siv durchführbar, die Kosten konvergieren zu einer asymptotischen durchschnittlichen Leistungsschranke und das geschlossene Regelsystem verifiziert die Zufallsbedingungen zu jedem Zeitpunkt. Die Eigenschaften des Reglers werden anhand von zwei numerischen Beispielen demonstriert.

Kapitel 4 In diesem Kapitel wird die Zustandsrückführung auf den Fall der Ausgangsrückführung erweitert, wobei zwei verteilte stochastische MPC Methoden zur Stabilisierung verteilter Systeme vorgeschlagen werden. Der erste Ansatz nutzt die direkte Rückführung (Abschnitt 4.2) und der zweite Ansatz basiert auf der indirekten Rückführung (Abschnitt 4.3).

Der erste Beitrag ist die Erweiterung des analytischen verteilten PRS Entwurfs aus Kapitel 3 auf den Fall der Ausgangsrückkopplung, welcher in Abschnitt 4.2 zur Herleitung eines analytischen verteilten stochastischer MPC Algorithmus mit direkter Ausgangsrückkopplung genutzt wird. Das MPC Optimierungsproblem erweist sich als rekursiv durchführbar, wobei die Kosten zu einer asymptotischen durchschnittlichen Leistungsschranke konvergieren, während das geschlossene System die Zufallsbeschränkungen zu jeden Zeitpunkt einhält. Der analytische Entwurf des verteilten PRS ist leider mit einem großen Konservatismus verbunden, den wir in Abschnitt 4.3 mit einem szenariobasierten verteilten PRS adressieren.

Hierzu werden die Zufallsbedingungen lediglich mit Garantien aus der Szenario-Optimierung verifiziert, d.h., der geschlossene Regelkreis verifiziert die Zufallsbedingungen mit hoher Wahrscheinlichkeit anstatt mit Wahrscheinlichkeit 1. Ein weiterer Beitrag ist ein verteiltes szenariobasiertes Entwurfsverfahren für das verteilte PRS. Der stochastische MPC mit indirekter Rückkopplung ist nachweislich rekursiv durchführbar, während die Kosten unter einer zusätzlichen Annahme von mittelwertfreiem unabhängig und identisch verteiltem (engl.: independent and identically distributed, i.i.d.) Rauschen zu einer asymptotischen durchschnittlichen Leistungsschranke konvergieren, die nicht schlechter ist als bei einem linearen Regler.

Kapitel 5 In diesem Kapitel wird eine Klasse von verteilten Systemen untersucht, die lokalen individuellen Zufallsbeschränkungen und multiplikativem Rauschen unterliegen. Der Hauptbeitrag ist eine auf analytischer Näherung basierende verteilte stochastische MPC, welche die Erwartungswert-Kostenfunktion analytisch über Mittelwert-Varianz Dynamik approximiert und lokale Zufallsbeschränkungen mittels der Cantelli-Ungleichung annähert. Das MPC Optimierungsproblem verwendet LMIs, um die verteilten Kovarianzmatrizen zu propagieren, welches wir anschließend mit verteiltem Konsensus ADMM lösen. Darüber hinaus wird ein verteiltes Entwurfsverfahren für die MPC Bestandteile bereitgestellt, so dass der Regler vollständig verteilt synthetisiert werden kann und kein zentraler Koordinationsknoten erforderlich ist. Das MPC Optimierungsproblem erweist sich als rekursiv durchführbar, der geschlossene Regelkreis als punktweise konvergent, und die Zufallsbeschränkungen werden zu jedem Zeitpunkt verifiziert.

Teil II

Im zweiten Teil dieser Arbeit untersuchen wir verteilungsrobuste MPC (engl.: distributionally robust MPC, DR-MPC) Verfahren für eine Klasse beschränkter stochastischer Systeme. DR-MPC ist ein relativ neues Forschungsgebiet, das sich aus einem praktischen Aspekt im Hinblick auf die Anwendbarkeit der stochastischen MPC ergibt. Insbesondere setzt die stochastische MPC voraus, dass die zugrundeliegende Wahrscheinlichkeitsverteilung oder die Momente exakt bekannt sind, was in der Praxis eine eher restriktive Annahme ist, d.h., die wahre Verteilung ist selten bekannt und muss aus begrenzten Daten geschätzt werden [126]. Dies ist besonders problematisch, wenn der Prozess der Datengenerierung kostspielig oder zeitaufwändig ist. Bei DR-MPC wird die Annahme der exakten Kenntnis der Verteilung (oder Momente) aufgehoben, indem das stochastische Optimalsteuerungsproblem über eine Klasse von Wahrscheinlichkeitsverteilungen optimiert wird, die in einer so genannten Ambiguitätsmenge (engl. ambiguity set) enthalten sind. In Kapitel 6 nutzen wir die Wasserstein Metrik als Diskrepanzmaß, welches zu einer Wasserstein Ambiguitätsmenge führt, wobei in Kapitel 7 eine Momenten-basierte Ambiguitätsmenge verwendet wird.

Kapitel 6 In diesem Kapitel werden zwei DR-MPC Verfahren mit Wasserstein Ambiguitätsmengen untersucht, bei denen wir Methoden aus der verteilungsrobusten Optimierung verwenden, um die stochastischen Optimalsteuerungsprobleme in Bezug auf die verteilungsbedingte Unsicherheit robust zu gestalten.

In Abschnitt 6.2 wird eine szenariobasierte DR-MPC Methode mit indirekter Rückkopplung behandelt, bei der nur eine potenziell kleine Menge von historischen Störungstrajektorien über einen Aufgabenhorizont angenommen wird. Die DR-MPC Methode erweist sich unabhängig von der Wahl des Hilfsreglers (engl.: tube controller) als rekursiv durchführbar, indem die Zustandsbeschränkungen relaxiert werden. Darüber hinaus können durch die indirekte Rückkopplung harte Begrenzung der Eingangssignale auch bei unbeschränkten Unsicherheiten auferlegt werden, was bei herkömmlichen MPC Regelungen mit direkter Rückkopplung nur für beschränkte stochastische Unsicherheiten möglich ist. Als weiteren Beitrag untersuchen wir für nichtlineare und lineare Hilfsregler, ob Garantien für die Kosten und/oder für die Einhaltung der Zufallsbedingungen bezüglich der wahren Verteilungsfunktion gegeben werden können. Anhand eines numerischen Beispiels einer Temperaturregelungsaufgabe eines Gebäudes mit vier Räumen wird der Einfluss des Wasserstein-Radius auf die Erfüllung der Zufallsbedingungen im geschlossenen Regelkreis veranschaulicht.

In Abschnitt 6.3 wird eine analytische DR-MPC Methode mit indirekter Rückkopplung entwickelt, welche die Zufallsbedingungen mit verteilungsrobusten PRS behandelt. Der erste Beitrag dieses Abschnitts ist die Erweiterung von PRS auf den verteilungsrobusten Fall, wobei wir für Wasserstein Ambiguitätsmengen ein szenariobasiertes Entwurfsverfahren vorschlagen. Unter der Annahme von mittelwertfreiem i.i.d. Rauschen, quadratischen Kosten und konstanter Hilfsreglerverstärkung kann gezeigt werden, dass die verteilungsrobusten Varianzkosten unabhängig von den MPC Entscheidungsvariablen sind, was es uns ermöglicht, diesen Teil bei der MPC Implementierung zu vernachlässigen. Das MPC Optimierungsprob-

lem ist aufgrund der indirekten Rückkopplung trivial rekursiv lösbar, während das geschlossene System aufgrund des verteilungsrobusten PRS die Zufallsbedingungen mit hoher Wahrscheinlichkeit verifiziert.

Kapitel 7 Im vorigen Kapitel haben wir Wasserstein Ambiguitätsmengen betrachtet, die Werkzeuge des maschinellen Lernens zur Kalibrierung des Ambiguitätsradius verwenden. In diesem Kapitel stellen wir momentbasierte Ambiguitätsmengen vor, die zwei wesentliche Vorteile gegenüber Wasserstein für eine Klasse von i.i.d. sub-Gauß'schen Zufallsvariablen haben. (i) Die Komplexität des resultierenden Optimierungsproblems nimmt nicht mit der Datenanzahl zu und (ii) der Ambiguitätsradius wird mit angemessener Genauigkeit auf der Grundlage bekannter Informationen der Verteilung über Konzentrationsungleichungen abgeschätzt. Daher sind die sich ergebenden Garantien im Vergleich zu Wasserstein stärker.

Der Hauptbeitrag dieses Kapitels ist eine neuartige DR-MPC Methode für eine Klasse von eingeschränkten stochastischen Systemen, die individuellen Zufallsbeschränkungen für die Zustände und Eingänge unterliegen. Im Gegensatz zu den vorangegangenen Kapiteln wird eine vereinfachte affine Rückkopplungsparametrisierung verwendet, um die Kostenfunktion und Zufallsbedingungen als deterministische Surrogate umzuformulieren. Die rekursive Machbarkeit wird durch die Beschränkung des Anfangszustands auf einer Linie zwischen dem gemessenen Zustandsvektor und einer garantiert machbaren Ersatzlösung hergestellt. Der geschlossene Regelkreis konvergiert nachweislich zu einer asymptotischen durchschnittlichen Leistungsschranke, die nicht schlechter ist als die des äquivalenten linearen quadratischen Reglers. Ein numerisches Beispiel demonstriert die Leistungsverbesserungen für eine zunehmende Anzahl von Daten.

Kapitel 8 In diesem Kapitel wird eine momentbasierte DR-MPC zur koordinierten Regelung von Windparks ab, welche die DR-MPC Formulierung aus Kapitel 7 erweitert, um Ausgangsvariablen in der Kostenfunktion zu verwenden. Wir integrieren ein ARMA Modell, das als optimaler Prädiktor für den turbulenten Wind dient und zusätzlich den stochastischen Vorhersagefehler als weißes i.i.d. Rauschen approximiert. Wir definieren eine parametrisierbare momentbasierte Ambiguitätsmenge für die Wertepaare von mittlerer Windgeschwindigkeit und Turbulenzintensität, wobei der zugehörige Ambiguitätsradius anhand der theoretischen Ergebnisse aus Kapitel 7 bestimmt wird. In einer numerischen Studie untersuchen wir die Vorteile von DR-MPC gegenüber einem klassischen Scheduler-Ansatz und vergleichen dabei für verschiedene Windszenarien den effektiven Leistungsgewinn/-verlust.

Curriculum Vitae

Personal Data

Name	Christoph Mark
E-Mail	mark@eit.uni-kl.de

Education

2015 - 2017	**University of Stuttgart, Stuttgart, Germany** M.Sc., Engineering Cybernetics Specialization: Autonomous Systems and Control Theory Thesis: *Nonlinear Model Predictive Control in application to collision avoidant Robotic Manipulation*
2011 - 2015	**Darmstadt University of Applied Sciences, Darmstadt, Germany** B.Eng., Electrical Engineering and Information Technology Specialization: Automation and Information Technology Thesis: *Synchrone Bewegungssteuerung für mobile Manipulatoren*

Professional Experience

2018 - 2022	**University of Kaiserslautern, Kaiserslautern, Germany** Department of Electrical and Computer Engineering, Institute of Control Systems Research associate
2017 - 2018	**University of Heidelberg, Heidelberg, Germany** Institute of Computer Engineering, Automation Laboratory Research associate

March 2023