Stochastic MPC with robustness to bounded parametric uncertainty

Elena Arcari*, Andrea Iannelli†, Andrea Carron*, Melanie N. Zeilinger*

***Abstract*— In this paper, we present a stochastic model predictive control approach for discrete-time LTI systems subject to bounded parametric uncertainty and potentially unbounded stochastic additive noise. The proposed scheme makes use of homothetic tubes along the prediction horizon for a robust treatment of parametric uncertainty. Stochastic noise is handled by tightening constraints using the concept of probabilistic reachable sets (PRS), which are typically constructed offline by exploiting noise distribution information. In order to address the presence of additional parametric uncertainty, we introduce a strategy for generating "robustified" PRS based only on the first and second moments of the noise sequence. In the case of quadratic cost functions, and under a further i.i.d. assumption on the noise distribution, we also provide an average asymptotic performance bound for the l2-norm of the closed-loop state. Finally, the proposed approach is demonstrated in both an illustrative example, and for a building temperature control problem.**

***Index Terms*— stochastic model predictive control, bounded parametric uncertainty, chance constraints**

I. Introduction

Model predictive control (MPC) has established itself as the state-of-the-art approach for high-performance control of constrained dynamical systems. In the presence of bounded uncertainty, rigorous theoretical guarantees are provided by [1], considering worst-case scenarios of all uncertainties affecting the dynamics. A robust design may result, however, in overly conservative control strategies, particularly when additional information about the uncertainty is available, e.g. in the form of a distribution or a parametric uncertainty. In this case, it can be beneficial to make use of a stochastic MPC approach, where constraints are imposed in probability, i.e. formulated as chance constraints for which a certain amount of violation is permitted, see e.g. [2]. This motivates the developments in this paper, offering an MPC formulation that handles robustly the presence of parametric uncertainty, which is assumed to be contained in a bounded polytopic set, and external disturbances modeled as additive noise, which have potentially unbounded, correlated distributions.

Previous results addressing multiple uncertainty sources can be found, for instance, in [3], [4], addressing a state and input-dependent uncertainty, or in [5]–[7] considering stochastic multiplicative and additive disturbances. While these results provide strong guarantees in terms of closed-loop chance constraint satisfaction and stability, these rely on the assumption of strict boundedness of all sources of uncertainty affecting the system. Under the same assumption, recent efforts have been made to improve control performance by including model learning [8]–[10], parameter adaptation [11], [12], or dual actions [13], [14].

In the presented paper, we relax the boundedness assumption on the additive disturbance, while still allowing for a robust analysis in terms of closed-loop feasibility and performance guarantees in the presence of bounded parametric uncertainty. The main idea builds on tube MPC concepts [5], [15], [16], and combines the use of homothetic tubes [11], [17] for handling parametric uncertainties along the prediction horizon, together with probabilistic reachable sets (PRS) for tightening state and input constraints, building on results in [2], [18]–[20]. We propose a procedure for "robustifying" the PRS design with respect to all parametric uncertainties, by only assuming knowledge of the first and second moments of the noise sequence affecting the system. Feedback is introduced through the cost function (indirect feedback [19]), which at each time step is computed with respect to the latest state measurement and parameter update. The combined use of homothetic tubes, "robustified" PRS (RPRS), and indirect feedback is shown to offer strong closed-loop guarantees and recursive feasibility. Furthermore, for i.i.d noise sequences and quadratic cost functions, we derive an average asymptotic bound on the $l2$-norm of the state.

A similar scenario, with both model and external uncertainties, has been tackled in the context of model-based safety filters [21], [22]. Differently from these results, we focus on exploiting the parametric structure of the uncertainty, enabling an MPC formulation that simultaneously achieves high performance while providing theoretical guarantees. The computational complexity of the proposed overall control scheme is not affected by the RPRS computations since these are constructed offline. Furthermore, the flexibility of the proposed formulation allows for various extensions to accommodate practical specifications [23], [24] while preserving the analysis, e.g. by incorporating ideas from [25], [26], where computational efficiency is improved at the expense of increased conservativeness by simplifying the homothetic tube online optimization.

The structure of the paper is as follows. The receding-

*The authors are with the Institute for Dynamical Systems and Control, ETH Zurich, ZH-8092, Switzerland: [earcari|carrona|mzeilinger]@ethz.ch.

† The author is with the Institute for Systems Theory and Automatic Control, University of Stuttgart, 70569 Stuttgart, Germany: andrea.iannelli@ist.uni-stuttgart.de.

horizon stochastic optimal control problem for discrete-time LTI systems is introduced in Section II. The problem components are defined in Section III, with a particular focus on the RPRS construction procedure. The MPC formulation and its closed-loop analysis are detailed in Section IV. Finally, numerical results in Section V are provided for both an illustrative example, and a building temperature control task.

Notation

Given a matrix X, $X_{i:l,j:m}$ defines a sub-matrix with elements from row i to l and columns j to m, while for a vector x, $[x]_i$ is the i-th entry. $\mathbb{E}[\cdot]$ and $\mathrm{Var}[\cdot]$ denote the expected value and the variance of a random variable, respectively. The convex hull of a set of points is denoted by $\mathrm{co}\{\cdot\}$. $\mathbf{1}_{n,m}$ and $\mathbf{0}_{n,m}$ denote matrices of ones and zeros of dimension $n\times m$, respectively, and I_n is the identity matrix of dimension n. The notation $x_{i|k}$ is used to refer to predicted quantities, where k identifies the time when the prediction is computed and i denotes the time step for which the prediction is made.

II. Problem formulation

We consider the control of uncertain linear time-invariant discrete-time dynamical systems modeled as

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k + w_k, \tag{1}$$

where the state is denoted by $x_k \in \mathbb{R}^n$, the input by $u_k \in \mathbb{R}^m$, and the additive stochastic disturbance by $w_k \in \mathbb{R}^n$. The system matrices depend affinely on an uncertain parameter vector $\theta \in \Theta \subseteq \mathbb{R}^p$

$$A(\theta) = A_0 + \sum_{i=1}^{p} A_i\,[\theta]_i, \tag{2a}$$

$$B(\theta) = B_0 + \sum_{i=1}^{p} B_i\,[\theta]_i, \tag{2b}$$

where $\Theta = \{\theta \mid H_\theta\theta \le h_\theta\}$, with $H_\theta \in \mathbb{R}^{q\times p}$, is assumed to contain the true unknown parameter vector θ^{true}.

We assume access to state measurements of the *true* system dynamics x_k^{true} resulting from θ^{true}, i.e.

$$x_{k+1}^{\text{true}} = A(\theta^{\text{true}})x_k^{\text{true}} + B(\theta^{\text{true}})u_k + w_k. \tag{3}$$

The goal is to control the true system (3)for the duration T of a finite-horizon task, in the face of uncertainty on θ^{true} and the additive uncertainty w_k. The proposed strategy leverages the model in (1) and takes into account both parametric uncertainty (2), and a potentially non i.i.d. stochastic disturbance sequence $W = [w_0^\top, \dots, w_{T-1}^\top]^\top \sim Q_W$, which may have unbounded support.

Assumption 1: We assume to have access to the first and second moments of Q_W.

System (3) is subject to constraints on both states and inputs. These are formulated as chance constraints that are required to be satisfied point-wise at each time-step $k \ge 0$ - and not jointly for all time-steps - with a probability conditioned on the true initial state x_0^{true}, i.e.

$$\Pr(x_k^{\text{true}} \in \mathcal{X} \mid x_0^{\text{true}}) \ge p_x, \quad \Pr(u_k \in \mathcal{U} \mid x_0^{\text{true}}) \ge p_u, \tag{4}$$

where $\mathcal{X} = \{x \mid Fx \le \mathbf{1}_{n_x,1}\}$, $F \in \mathbb{R}^{n_x\times n}$, and $\mathcal{U} = \{u \mid Gu \le \mathbf{1}_{n_u,1}\}$, $G \in \mathbb{R}^{n_u\times m}$. Additionally, $p_x, p_u \ge 0$ are the assigned probability levels.

Using model (1), we formulate the control task subject to (4) as a constrained optimization problem to be solved in a receding horizon fashion. In order to obtain a tractable formulation that can handle the presence of chance constraints, we restrict the class of control policies over which we optimize to the following affine state feedback law

$$u_k = Kx_k + v_k, \tag{5}$$

where K satisfies the following commonly used assumption in robust MPC [11], [27].

Assumption 2: The state feedback gain $K \in \mathbb{R}^{m\times n}$ is chosen such that the closed-loop dynamics $A_{CL}(\theta) = A(\theta) + B(\theta)K$ is asymptotically stable $\forall\theta \in \Theta$, i.e. there exists a positive definite $P \succ 0$ such that

$$A_{CL}(\theta)^\top P A_{CL}(\theta) - P \prec 0,\ \forall\theta \in \Theta.$$

Furthermore, we define the auxiliary variables $z_k \in \mathbb{R}^n$ and $e_k \in \mathbb{R}^n$ as

$$e_k = x_k - z_k, \tag{6}$$

with the aim of separating the effect of the two different uncertainty sources. Using (5), (1) and (6), we obtain

$$\begin{aligned}
&e_{k+1} + z_{k+1} = x_{k+1},\\
&e_{k+1} + z_{k+1} = A(\theta)x_k + B(\theta)u_k + w_k,\\
&e_{k+1} + z_{k+1} = A_{CL}(\theta)z_k + A_{CL}(\theta)e_k + B(\theta)v_k + w_k.
\end{aligned}$$

Finally, the dynamics of z_k and e_k can be split and defined as

$$z_{k+1} = A_{CL}(\theta)z_k + B(\theta)v_k, \tag{7a}$$

$$e_{k+1} = A_{CL}(\theta)e_k + w_k. \tag{7b}$$

A similar split starting from the true dynamics (3) defines

$$z_{k+1}^{\text{true}} = A_{CL}(\theta^{\text{true}})z_k^{\text{true}} + B(\theta^{\text{true}})v_k, \tag{8a}$$

$$e_{k+1}^{\text{true}} = A_{CL}(\theta^{\text{true}})e_k^{\text{true}} + w_k, \tag{8b}$$

where (8a) represents the true nominal dynamics, while (8b) represents the true error dynamics induced by the presence of additive noise. The case in which θ^{true} is known is addressed in [19] using a similar dynamics split, where the formulated control problem optimizes over an error feedback, rather than a state feedback strategy (5).

We build on tube MPC concepts to address the chance constraints (4). As θ^{true} is unknown, we first exploit robust concepts to predict a tube for the dynamics of z_k. In addition, the stochastic noise w_k is handled by constructing another tube for e_k that is used for tightening both state and input constraints such that the probability levels in (4) can be satisfied. Figures 1 and 2 provide an illustrative scalar example of an autonomous system, and highlight the main differences with the case in which there is no parametric uncertainty (8). It is important to note that, while (7b) evolves autonomously, and therefore enables an offline construction of the tube for e_k, the dynamics of z_k in (7a) are controlled by v_k, which does not allow to pre-compute the tube for z_k as for the autonomous system shown in Figure 2. The idea is to make use of an optimization-based approach to compute such a tube and a corresponding constraint tightening that ensures constraint satisfaction (4) for x_k and u_k. We formulate hereafter an associated optimal control problem to be solved over a horizon $N < T$ in a receding horizon fashion

$$\min_{\substack{\{v_{i|k}\}_{i=0}^{N-1},\\ \{\mathbf{Z}_{i|k}\}_{i=0}^{N}}} \mathbb{E}_{W_k}\left[\sum_{i=0}^{N-1} l_i(\bar{x}_{i|k}, \bar{u}_{i|k}) + l_f(\bar{x}_{N|k})\right] \tag{9a}$$

$$\begin{aligned}
\text{s.t.}\quad & \bar{x}_{i+1|k} = A(\bar{\theta}_k)\bar{x}_{i|k} + B(\bar{\theta}_k)\bar{u}_{i|k} + w_{i|k}, && (9b)\\
& \bar{u}_{i|k} = K\bar{x}_{i|k} + v_{i|k}, && (9c)\\
& W_k = [w_{0|k}^\top, \dots, w_{N-1|k}^\top]^\top \sim Q_{W_k}, && (9d)\\
& \mathbf{Z}_{i|k} \subseteq \mathcal{X} \ominus \mathbf{E}_{k+i}, && (9e)\\
& K\mathbf{Z}_{i|k} \oplus v_{i|k} \subseteq \mathcal{U} \ominus \mathbf{E}^u_{k+i}, && (9f)\\
& \mathbf{Z}_{N|k} \subseteq \mathcal{Z}_f, && (9g)\\
& A_{CL}(\theta)\mathbf{Z}_{i|k} + B(\theta)v_{i|k} \subseteq \mathbf{Z}_{i+1|k},\ \forall \theta \in \Theta, && (9h)\\
& \bar{x}_{0|k} = x_k^{\text{true}}, && (9i)\\
& z_k^{\text{true}} \in \mathbf{Z}_{0|k}. && (9j)
\end{aligned}$$

The overall cost function to be optimized is computed as the sum of potentially time-varying stage costs $l_i(\cdot,\cdot)$, $i \in [0, N-1]$, and a terminal cost $l_f(\cdot)$. The cost is evaluated with respect to a point estimate of the uncertain parameter θ that we denote by $\bar{\theta}_k \in \Theta\ \forall k \geq 0$. Additionally, due to the additive stochastic noise, the cost is defined as the expectation with respect to a predicted noise sequence W_k whose distribution Q_{W_k} is defined by the conditional distribution $p([w_k^\top, \dots, w_{k+N-1}^\top]^\top | [w_0^\top, \dots, w_{k-1}^\top]^\top)$.

Assumption 3: We assume to have access to either the density function or samples of the distribution Q_{W_k} in order to compute the cost function expectation in (9a).

Problem (9) deals with the presence of parametric model uncertainty by optimizing over a sequence of bounded sets $\{\mathbf{Z}_{i|k}\}_{i=0}^{N}$ along the prediction horizon that we refer to as *nominal tube*, ensuring robust containment of z_k for all $\theta \in \Theta$ (see, e.g., the tube for z_k in Figure 2). Furthermore, we design a sequence of confidence regions $\mathbf{E}_k$ containing e_k, which we use to tighten state constraints. We refer to $\{\mathbf{E}_k\}_{k=1}^{T}$ as the *stochastic error tube*, for which containment holds for all $\theta \in \Theta$ with a probability dictated by the distribution of the sequence W (see, e.g., the tube for e_k in Figure 2). Similarly, we construct sets $\mathbf{E}_k^u$ containing in probability $e_k^u = Ke_k$, which are used to tighten input constraints. The following sections provide details on how to design both the nominal tube and the confidence regions needed for state and input constraint tightening (9e), (9f). Further clarifications are also given regarding the construction of an appropriate terminal set $\mathcal{Z}_f$ (9g), and the reformulation of the nominal tube containment condition (9h), such that ultimately the overall problem is recursively feasible, and guarantees closed-loop chance constraint satisfaction (4). Note that recursive feasibility also requires that condition (9j) is guaranteed despite not having access to the true nominal system dynamics (8a).

Remark 1: Parameter estimate update
We do not make any assumption on the learning scheme

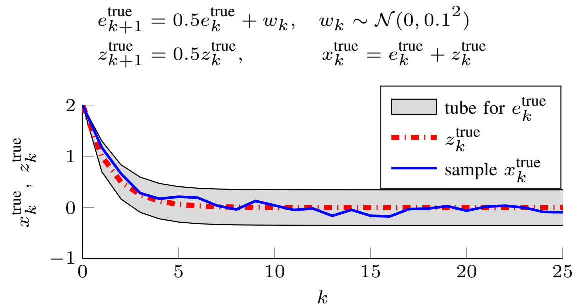

Fig. 1. Without parametric uncertainty, the error between the nominal trajectory z_k^{true} (in red) and any sampled trajectory x_k^{true} (in blue) depends on the particular realization of w_k. The tube for e_k^{true} is depicted in grey, and defined as three standard deviations of its distribution.

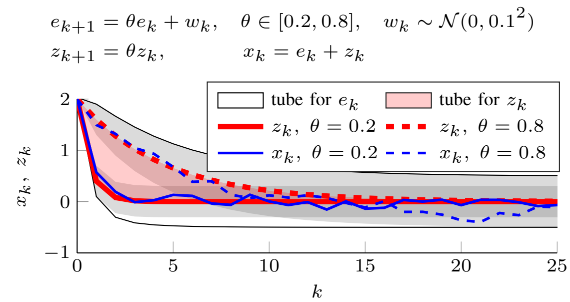

Fig. 2. With parametric uncertainty, all nominal trajectories z_k for some $\theta \in [0.2, 0.8]$ are contained in the tube for z_k (in red) for which we highlight two sample trajectories. In blue, we depict two sample trajectories of x_k. The grey shaded areas represent three standard deviations of e_k, for $\theta = 0.2$, and $\theta = 0.8$. The boundaries (in black) of the tube for e_k are computed with respect to three standard deviations of the worst case error realizations and therefore contain all possible tubes computed for any fixed value of θ.

chosen to update the point estimate $\bar{\theta}_k$. The only condition to be satisfied is containment in the bounded set Θ, which can be always guaranteed by adding a projection step to any update scheme. An example is to use a recursive least squares update [28], with added set projection.

III. Tractable Formulation of Stochastic MPC with Bounded Parametric Uncertainty

In the following section, we provide details regarding the nominal tube, and how its structure can be exploited for reformulating the containment condition along the prediction horizon. Then, we focus on the procedure for constructing confidence regions for any noise sequence affecting the system, and for determining an appropriate constraint tightening despite the presence of parametric uncertainty. Finally, the overall stochastic MPC problem is defined, expanding the formulation provided in (9).

A. Nominal tube

The nominal tube predicted along a horizon of length N is defined as a sequence of sets $\mathbf{Z}_{i|k},\ i \in [0\!:\!N]$. In order to ease computations, these sets are restricted to be translations and scalings of a given convex set $\bar{\mathbf{Z}}$, which are typically referred to as homothetic tubes [27]

$$\mathbf{Z}_{i|k} = \{s_{i|k}\} \oplus \alpha_{i|k}\bar{\mathbf{Z}}, \tag{10}$$

where $s_{i|k} \in \mathbb{R}^n$, $\alpha_{i|k} \in \mathbb{R}$. By choosing the base set $\bar{\mathbf{Z}}$ to be a polytope defined as $\{\bar{z} \mid H_z\bar{z} \leq \mathbf{1}_{r,1}\} = \mathrm{co}\{\bar{z}^1, \dots, \bar{z}^{v_1}\}$, the containment condition (9h) can be reformulated similarly to [11] as

$$\begin{aligned}
&H_z(A_{CL}(\theta)(s_{i|k} + \alpha_{i|k}\bar{z}) + B(\theta)v_{i|k} - s_{i+1|k}) \\
&\qquad -\alpha_{i+1|k}\mathbf{1}_{r,1} \leq \mathbf{0}_{r,1}, \quad \forall \bar{z} \in \bar{\mathbf{Z}},\ \theta \in \Theta, \quad \Leftrightarrow \\
&H_z(A_{CL}(\theta)(s_{i|k} + \alpha_{i|k}\bar{z}^j) + B(\theta)v_{i|k} - s_{i+1|k}) \\
&-\alpha_{i+1|k}\mathbf{1}_{r,1} \leq \mathbf{0}_{r,1}, \quad \forall j \in \{1, \dots, v_1\},\ \theta \in \Theta, \quad \Leftrightarrow \\
&\max_{\theta\in\Theta}\{H_z D^j_{i|k}\theta\} + H_z d^j_{i|k} \leq \alpha_{i+1|k}\mathbf{1}_{r,1}, \quad \forall j \in \{1, \dots, v_1\},
\end{aligned} \tag{11}$$

where the last expression is obtained by using (2), and by introducing the following terms for all $j \in \{1, \dots, v_1\}$

$$\begin{aligned}
p^j_{i|k} &= s_{i|k} + \alpha_{i|k}\bar{z}^j, \\
r^j_{i|k} &= K(s_{i|k} + \alpha_{i|k}\bar{z}^j) + v_{i|k}, \\
D^j_{i|k} &= D(p^j_{i|k}, r^j_{i|k}), \\
d^j_{i|k} &= (A_0 + B_0K)(s_{i|k} + \alpha_{i|k}\bar{z}^j) + B_0v_{i|k} - s_{i+1|k},
\end{aligned}$$

where the function $D(a, b) : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^{n\times p}$ maps vectors a and b to a matrix whose columns are defined with respect to $\{A_i\}_{i=1}^p$ and $\{B_i\}_{i=1}^p$ (2), resulting in

$$D(a, b) = \left[A_1a + B_1b, \dots, A_pa + B_pb\right]. \tag{12}$$

Finally, maximization in (11) can be cast as its corresponding dual problem, i.e. minimization with respect to the dual variables $\{\Lambda^j_{i|k}\}_{j=1}^{v_1}$. We can therefore reformulate (9h) following a procedure similar to [11] as

$$\Lambda^j_{i|k}h_\theta + H_z d^j_{i|k} \leq \alpha_{i+1|k}\mathbf{1}_{r,1}, \tag{13a}$$

$$H_z D^j_{i|k} = \Lambda^j_{i|k}H_\theta, \tag{13b}$$

$$\Lambda^j_{i|k} \in \mathbb{R}^{r\times p}_{\geq 0}, \tag{13c}$$

where we include positivity conditions (13c) and Lagrangian stationarity, thus ensuring optimality.

B. Stochastic error tube

The purpose of the stochastic error tube is to bound in probability e_k, and consequently the true error state e^{true}_k at each time-step k. A procedure that makes use of the first and second moments of e^{true}_k for constructing such confidence regions, i.e. k-step PRS, is given in [19]. Since the matrix $A_{CL}(\theta^{\text{true}})$ determining the dynamics of e^{true}_k is unknown in the considered setup, these sets cannot be directly constructed. The idea is to formally define a "bound" on the moments of e^{true}_k that can be used to construct a sequence of confidence regions, which we refer to as k-step "robustified" PRS (RPRS) $\mathbf{E}_k$, satisfying the following condition for $k \in [1\!:\!T]$

$$\Pr(e_k \in \mathbf{E}_k \mid e^{\text{true}}_0) \geq p,\ \forall \theta \in \Theta. \tag{14}$$

The remainder of this section is devoted to detailing a procedure for synthesizing k-step RPRS both in the case of i.i.d., and correlated noise sequences affecting the system dynamics.

Note that, given $e^{\text{true}}_0 = \mathbf{0}_{n,1}$, and $\mathbb{E}[W] = \mathbf{0}_{nT,1}$, then $\mathbb{E}[e_k] = \mathbf{0}_{n,1}$, $\forall k \geq 0$ and $\forall \theta \in \Theta$. Therefore, each confidence region associated with a particular value of θ remains centered at the origin (see e.g. [29]).

Remark 2: Noise sequence with non-zero mean Disturbance sequences with first moment different from zero, $\overline{W} = \mathbb{E}[W]$, can be considered by defining a stochastic sequence $\widetilde{W} = W - \overline{W}$, with $\mathbb{E}[\widetilde{W}] = \mathbf{0}_{nT,1}$, and $\mathrm{Var}[\widetilde{W}] = \mathrm{Var}[W]$. Then, the sequence $\overline{W}$ can be directly included in the dynamics (7a), and handled by the nominal tube.

In order to compute the stochastic error tube $\{\mathbf{E}_k\}_{k=1}^T$, the aim is to "bound" at each time-step k the marginal variance $\{\mathrm{Var}[e_k]\}_{k=1}^T$, corresponding to the n-dimensional block diagonal entries of $\mathrm{Var}[E] \in \mathbb{R}^{nT\times nT}$, i.e. the variance of the sequence $E = [e_1^\top, \dots, e_T^\top]^\top$ defined as

$$\mathrm{Var}[E] = \overline{A_{CL}}(\theta)\mathrm{Var}[W]\overline{A_{CL}}(\theta)^\top,$$

$$\overline{A_{CL}}(\theta) = \begin{bmatrix} I_n & \mathbf{0}_{n,n} & \dots & \mathbf{0}_{n,n} \\ A_{CL}(\theta) & I_n & \dots & \mathbf{0}_{n,n} \\ A_{CL}(\theta)^2 & A_{CL}(\theta) & \dots & \mathbf{0}_{n,n} \\ \vdots & \vdots & \ddots & \vdots \\ A_{CL}(\theta)^{N-1} & A_{CL}(\theta)^{N-2} & \dots & I_n \end{bmatrix}.$$

Note that for each $\theta \in \Theta$, $\mathrm{Var}[E]$ is a well-defined covariance matrix. We formalize a bound in terms of the Loewner order, which minimizes the variance spread in the direction of its principal components [30], [31]. The associated optimization problem can be formulated as

$$(\overline{\mathrm{Var}}[e_k])^{-1} =$$

$$\underset{X^{-1}}{\arg\min}\ -\log\det X^{-1} \tag{15a}$$

$$\text{s.t.}\quad X - \mathrm{Var}[e_k] \succeq 0, \quad \forall \theta \in \Theta \tag{15b}$$

to be solved for $k \in [1 : T]$. In the following subsections, we describe reformulations of problem (15) determining the bounding sequence $\{\overline{\mathrm{Var}}[e_k]\}_{k=1}^T$, both for i.i.d. and correlated noise sequences.

1) I.i.d. noise sequences: We consider the particular case

$$\mathrm{Var}[W] = \begin{bmatrix} \Sigma_w & \mathbf{0}_{n,n} & \dots & \mathbf{0}_{n,n} \\ \mathbf{0}_{n,n} & \Sigma_w & \dots & \mathbf{0}_{n,n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{n,n} & \mathbf{0}_{n,n} & \dots & \Sigma_w \end{bmatrix},$$

i.e. the only non-zero entries of $\mathrm{Var}[W]$ are its identical block-diagonal entries $\Sigma_w \in \mathbb{R}^{n\times n}$. This means that the marginal variances can be iteratively computed for $k = 1, \dots, T$ as

$$\mathrm{Var}[e_{k+1}] = A_{CL}(\theta)\mathrm{Var}[e_k]A_{CL}(\theta)^\top + \Sigma_w,\ \forall \theta \in \Theta. \tag{16}$$

Since e^{true}_0 is known, we can directly infer that $\overline{\mathrm{Var}}[e_1] = \Sigma_w$. For $k = 2, \dots, T$, we can use (16) in problem (15), to obtain

$$\begin{aligned}
&(\overline{\mathrm{Var}}[e_{k+1}])^{-1} = \\
&\underset{X^{-1}}{\arg\min}\ -\log\det X^{-1} \\
&\text{s.t.}\quad X - A_{CL}(\theta)\overline{\mathrm{Var}}[e_k]A_{CL}(\theta)^\top - \Sigma_w \succeq 0, \\
&\qquad \forall \theta \in \Theta
\end{aligned} \tag{17}$$

where we iteratively use the solution at time-step k to obtain the solution at $k+1$. Since by construction $\Sigma_w \succ 0$, and $\overline{\mathrm{Var}}[e_k] \succ 0, \forall k > 0$, problem (17) admits a convex reformulation following Lemma 3 in Appendix VII-B as

$$\begin{aligned}
&(\overline{\mathrm{Var}}[e_{k+1}])^{-1} = \\
&\underset{X^{-1}}{\arg\min} \; -\log\det X^{-1} \\
&\text{s.t.} \quad \begin{bmatrix} X^{-1} & X^{-1}\Sigma_w & X^{-1}A_{CL}(\theta^j) \\ \Sigma_w X^{-1} & \Sigma_w & 0 \\ A_{CL}(\theta^j)^\top X^{-1} & 0 & (\overline{\mathrm{Var}}[e_k])^{-1} \end{bmatrix} \succeq 0, \\
&\qquad \forall j \in \{1,\dots,v_2\},
\end{aligned} \tag{18}$$

where $\{\theta^j\}_{j=1}^{v_2}$ are the vertices of Θ.

2) Correlated noise sequences: For correlated noise sequences with full covariance matrix $\mathrm{Var}[W]$, we cannot sequentially compute the bounding matrix sequence as in the i.i.d. case. The marginal variance $\mathrm{Var}[e_k]$ depends on all previous time steps and therefore contains a series of nonlinear terms, i.e. powers of $A_{CL}(\theta)$:

$$\begin{aligned}
&\mathrm{Var}[e_k] = \\
&\left[A_{CL}^{k-1}(\theta) \dots I_n\right] \mathrm{Var}[W]_{1:kn,1:kn} \left[A_{CL}^{k-1}(\theta)^\top \dots I_n\right]^\top,
\end{aligned} \tag{19}$$

which cause problem (15) to be intractable. In the following, we propose a procedure summarized in Algorithm 1 where (15) is broken down into a sequence of tractable sub-problems that admit convex reformulations similar to problem (18). As for the i.i.d. case, $\overline{\mathrm{Var}}[e_1]$ is initialized to $\mathrm{Var}[W]_{1:n,1:n}$. For each time-step $k \in [2:T]$, the idea is to sequentially factorize the matrix product (19) such that at each iteration $i \in [1:k-1]$ two sub-problems, defined as *bound1* and *bound2* problems, provide tractable intermediate solutions $\bar{Y}(i),\ i = 1,\dots,k$, with $\bar{Y}(k)$ corresponding to the final bound $\overline{\mathrm{Var}}[e_k]$. Further details regarding these sub-problems are provided in the rest of this section. First, note that the following relation holds for each $k \in [2:T]$

$$\left[A_{CL}^{k-1}(\theta) \quad A_{CL}^{k-2}(\theta) \ \dots\right] = \left[A_{CL}^{k-2}(\theta)\left[A_{CL}(\theta) \quad I_n\right] \ \dots\right],$$

and that by setting $A_I(\theta) = \left[A_{CL}(\theta) \quad I_n\right]$, we obtain the reformulation in (22) that determines the first sub-problem at iteration $i=1$, i.e. the *bound1* problem defined as:

$$\begin{aligned}
&\min_{D(1)^{-1}} \; -\log\det D(1)^{-1} \\
&\text{s.t.} \quad D(1) - A_I(\theta)\bar{Y}(1)_{1:2n,1:2n}A_I(\theta)^\top \succeq 0, \ \forall\theta \in \Theta.
\end{aligned} \tag{20}$$

Using $D(1)$, we construct the matrix $Y_1(\theta)$ defined in (22), which depends affinely on the parameter θ. Therefore, this can be again bounded as

$$\begin{aligned}
&\min_{\bar{Y}^{-1}(2)} \; -\log\det \bar{Y}^{-1}(2) \\
&\text{s.t.} \quad \bar{Y}(2) - Y_1(\theta) \succeq 0, \quad \forall\theta\in\Theta,
\end{aligned} \tag{21}$$

which we refer to as the *bound2* problem. Matrix $\bar{Y}(2)$ can now be used to proceed with the recursion, i.e. we again isolate a block $A_I(\theta)$ from $\left[A_{CL}^{k-2}(\theta) \quad \dots \quad I_n\right]$, and solve the associated problem (20) to obtain $D(2)$, and (21) to obtain $\bar{Y}(3)$. This factorization is repeated for all i until we reach the final step returning $\bar{Y}(k)$, which provides a bound for $\mathrm{Var}[e_k]$ (see (22)).

Note that *bound1* problem in (20) admits a convex reformulation, provided that the following matrix

$$\begin{aligned}
\widetilde{X}(\theta) = {} & A_{CL}(\theta)\bar{Y}(1)_{1:n,n+1:2n} \\
& + \bar{Y}(1)_{n+1:2n,1:n}A_{CL}(\theta)^\top + \bar{Y}(1)_{n+1:2n,n+1:2n}.
\end{aligned}$$

is positive definite. Then, Lemma 3 in Appendix VII-B can be applied, and (20) becomes for each *bound1* problem at step $i \in [1:k-1]$:

$$\begin{aligned}
&\min_{D(1)^{-1}} \; -\log\det D(1)^{-1} \\
&\text{s.t.} \\
&\begin{bmatrix} D(1)^{-1} & D(1)^{-1}\widetilde{X}(\theta^j) & D(1)^{-1}A_{CL}(\theta^j) \\ \widetilde{X}(\theta^j)D(1)^{-1} & \widetilde{X}(\theta^j) & 0 \\ A_{CL}(\theta^j)^\top D(1)^{-1} & 0 & (\bar{Y}(1)_{1:n,1:n})^{-1} \end{bmatrix} \succeq 0, \\
&\forall j \in \{1,\dots,v_2\}.
\end{aligned} \tag{23}$$

Finally, by pre- and post-multiplying by $\bar{Y}^{-1}(2)$ in (21), and making use of the Schur complement together with Lemma 1, we obtain the following convex reformulation

$$\begin{aligned}
&\min_{\bar{Y}^{-1}(2)} \; -\log\det \bar{Y}^{-1}(2) \\
&\text{s.t.} \quad \begin{bmatrix} \bar{Y}^{-1}(2) & \bar{Y}^{-1}(2)Y(\theta^j) \\ Y(\theta^j)\bar{Y}^{-1}(2) & Y(\theta^j) \end{bmatrix} \succ 0, \\
&\qquad \forall j \in \{1,\dots,v_2\}.
\end{aligned} \tag{24}$$

which is used for each *bound2* problem at step $i \in [1 : k-1]$.

Remark 3: Positive definitiveness requirement
Satisfaction of the requirement $\widetilde{X}(\theta^j) \succ 0$ for each $i \in [1 : k-1]$, $k \geq 2,\ j \in \{1,\dots,v_2\}$, needed for applying Lemma 3, will typically depend on how strong correlations are in the noise sequence W that affects the evolution of the system dynamics. Alternatively, one can compute a positive definite upper bound for $\widetilde{X}(\theta^j)$, which may ultimately generate more conservative k-step RPRS. Note that a similar condition can be found in [32], referred to as correlation bound.

Remark 4: Scalability of Algorithm 1
While all k-step RPRS are pre-computed offline and therefore do not increase the complexity of the associated control problem, the procedure outlined in Algorithm 1 can become computationally expensive for high dimensional systems, and for long noise sequences, as it requires to solve $2(k-2)+1$ semi-definite programs for each time-step $k \geq 2$. One way to improve scalability is to replace the linear matrix inequalities with diagonal dominance constraints, i.e. a sufficient condition that allows for reformulating all optimization problems involved as linear programs (see e.g. Theorem 6.1.10 in [33]).

Remark 5: While in this paper we do not consider parametric uncertainty in B_w, the computation of k-step RPRS can similarly address the case in which the matrix B_w depends affinely on an uncertain parameter θ, as for the dynamics matrices in (2), since Lemma 1 can be applied.

$$
\begin{aligned}
&\text{Var}[e_k] = \left[A_{CL}^{k-2}(\theta)A_I(\theta) \ \dots \ I_n\right] \bar{Y}(1) \left[A_I(\theta)^\top A_{CL}^{k-2}(\theta)^\top \dots I_n\right]^\top = \\
&\left[A_{CL}^{k-2}(\theta) \ \dots \ I_n\right] \begin{bmatrix} \overbrace{A_I(\theta)\bar{Y}(1)_{1:2n,1:2n}A_I(\theta)^\top}^{\preceq D(1),\ \text{problem (20)}} & A_I(\theta)\bar{Y}(1)_{1:2n,2n+1:kn} \\ \bar{Y}(1)_{2n+1:kn,1:2n}A_I(\theta)^\top & \bar{Y}(1)_{2n+1:kn,2n+1:kn} \end{bmatrix} \left[A_{CL}^{k-2}(\theta)^\top \dots I_n\right]^\top \preceq \\
&\left[A_{CL}^{k-2}(\theta) \ \dots \ I_n\right] \underbrace{\begin{bmatrix} D(1) & A_I(\theta)\bar{Y}(1))_{1:2n,2n+1:kn} \\ \bar{Y}(1)_{2n+1:kn,1:2n}A_I(\theta)^\top & \bar{Y}(1)_{2n+1:kn,2n+1:kn} \end{bmatrix}}_{=Y_1(\theta)} \left[A_{CL}^{k-2}(\theta)^\top \dots I_n\right]^\top \underbrace{\preceq}_{(21)} \left[A_{CL}^{k-2}(\theta) \ \dots \ I_n\right] \bar{Y}(2) \left[A_{CL}^{k-2}(\theta)^\top \dots I_n\right]^\top \preceq \\
&\dots \preceq \left[A_{CL}(\theta) \quad I_n\right] \bar{Y}(k-1) \left[A_{CL}(\theta)^\top \quad I_n\right]^\top \preceq D(k-1) = Y_{k-1}(\theta) = \bar{Y}(k) = \overline{\text{Var}}[e_k]
\end{aligned} \tag{22}
$$

Algorithm 1 Marginal variance bound for correlated noise

Require: $A_{CL}(\theta)$, $\text{Var}[W]$, Θ
  $\overline{\text{Var}}[e_1] = \text{Var}[W]_{1:n,1:n}$
  for $k \in \{2, \dots, T\}$ **do**
    $\bar{Y}(1) = \text{Var}[W]_{1:kn,1:kn}$
    for $i \in \{1, \dots, k-1\}$ **do**
      $D(i)$ computed with *bound1* problem in (20)
      Define $Y_i(\theta)$ as in (22)
      if $Y_i(\theta) \neq D(i)$ **then**
        $\bar{Y}(i+1)$ computed with *bound2* problem in (21)
      else if $Y_i(\theta) == D(i)$ **then**
        $\bar{Y}(i+1) = D(i)$
      end if
    end for
    return $\overline{\text{Var}}[e_k] = \bar{Y}(k)$
  end for
  return $\left[\overline{\text{Var}}[e_1] \quad \dots \quad \overline{\text{Var}}[e_T]\right]$

C. Variance-based k-step RPRS

Once the sequence $[\overline{\text{Var}}[e_1], \dots, \overline{\text{Var}}[e_T]]$ is available, the uncertainty of e_k at each time-step k is fully specified for all $\theta \in \Theta$. We can then construct different types of confidence regions based on Chebychev's bound: one option is to generate ellipsoidal k-step RPRS as

$$\mathbf{E}_k^{\text{ell}} = \{e \mid e^\top (\overline{\text{Var}}[e_k])^{-1} e \leq \tilde{p}\}, \tag{25}$$

where $\tilde{p} = \frac{n}{1-p}$ with p being the probability level, and n is the dimension of e_k. If the distribution of the error sequence is Gaussian, then we can set $\tilde{p} = \chi_n^2(p)$, i.e. the quantile function of the chi-squared distribution with n degrees of freedom. Alternatively, one can consider half-spaces:

$$\mathbf{E}_k^{\text{hs}} = \{e \mid h^\top e \leq \sqrt{\tilde{p} h^\top \overline{\text{Var}}[e_k] h}\}, \tag{26}$$

which is a k-step RPRS of probability level p with $\tilde{p} = \frac{1}{1-p}$, or $\tilde{p} = \chi_1^2(2p-1)$ for Gaussian distributions (further details can be found in [19], [34]).

D. Chance constraint reformulation

Using the stochastic error tube, we can now define a time-varying state constraint tightening $\mathcal{Z}_k = \mathcal{X} \ominus \mathbf{E}_k$, and input constraint tightening $\mathcal{V}_k = \mathcal{U} \ominus \mathbf{E}_k^u$. A k-step RPRS $\mathbf{E}_k^u$ for the input can be easily obtained based on the variance propagation of $e_k^u = K e_k$, and re-using the computations from the procedure outlined in section III-B such that $\overline{\text{Var}}[e_k^u] = K \overline{\text{Var}}[e_k] K^\top$ for each time-step k.

Starting from the options provided in subsection III-C, there are two possibilities: one, is to first generate ellipsoidal sets $\mathbf{E}_k^{\text{ell}}$ and $\mathbf{E}_k^{\text{ell},u}$ (25), and tighten the constraints using the support function of an ellipsoid. An alternative approach is to make use of Boole's inequality for defining the single half-spaces $\mathbf{E}_k^{\text{hs}}$ and $\mathbf{E}_k^{\text{hs},u}$ (26), whose intersection determines the overall polytopic k-step RPRS (see e.g. [34]). Either option provides tightened sets of the form $\mathcal{Z}_k = \{z \mid Fz \leq \mathbf{1}_{n_x,1} - f_k\}$, and $\mathcal{V}_k = \{u \mid Gu \leq \mathbf{1}_{n_u,1} - g_k\}$, which enable the following state and input containment conditions

$$
\begin{aligned}
&\mathbf{Z}_{i|k} \subseteq \mathcal{Z}_{k+i} \Leftrightarrow F s_{i|k} \leq \mathbf{1}_{n_x,1} - f_{k+i} - \alpha_{i|k} \max_{\bar{z} \in \bar{\mathbf{Z}}} F\bar{z} \\
&K\mathbf{Z}_{i|k} \oplus v_{i|k} \subseteq \mathcal{V}_{k+i} \Leftrightarrow \\
&G(K s_{i|k} + v_{i|k}) \leq \mathbf{1}_{n_u,1} - g_{k+i} - \alpha_{i|k} \max_{\bar{z} \in \bar{\mathbf{Z}}} GK\bar{z},
\end{aligned}
$$

where we define $\bar{f} = \max_{\bar{z} \in \bar{\mathbf{Z}}} F\bar{z}$, and $\bar{g} = \max_{\bar{z} \in \bar{\mathbf{Z}}} GK\bar{z}$.

Remark 6: Non-conservative constraint tightening
Since constraint tightening guarantees should ideally hold jointly for the entire $\mathcal{X}$ and $\mathcal{U}$, either approach for generating k-step RPRS presented above can introduce conservatism. However, in case either the state constraints are aligned with the nominal tube $\mathbf{Z}_{i|k}$, or the input constraints are aligned with $K\mathbf{Z}_{i|k}$, due to the particular choice of the base set $\bar{\mathbf{Z}}$, one can construct - for either state or input - a non-conservative tightening for each half-space independently (see e.g. [19]).

E. Final problem

Before stating the final MPC problem, we provide conditions for an appropriate terminal set design analogously to [11], [27].

Assumption 4: (Terminal set for nominal tube)
There exists a non-empty terminal set $\mathcal{Z}_f = \{(s, \alpha) \in \mathbb{R}^{n+1} \mid H_T s + h_T \alpha \leq \mathbf{1}_{n_f,1}\}$, with $s \oplus \alpha \bar{\mathbf{Z}} \subseteq \mathcal{Z}_\infty = \bigcap_{k=1}^T \mathcal{Z}_k, \ \forall (s, \alpha) \in \mathcal{Z}_f$, that is robust positively invariant for the set dynamics (9h) under the zero terminal control law contained in $\mathcal{V}_\infty = \bigcap_{k=1}^T \mathcal{V}_k$, i.e. we have $\forall \theta \in \Theta$

$$
\begin{aligned}
&(s, \alpha) \in \mathcal{Z}_f \Rightarrow \exists (s^+, \alpha^+) \in \mathcal{Z}_f \text{ s.t.} \\
&\qquad A_{CL}(\theta)(\{s\} \oplus \alpha \bar{\mathbf{Z}}) \subseteq \{s^+\} \oplus \alpha^+ \bar{\mathbf{Z}}
\end{aligned}
$$

Introducing the components derived in sections III-A, III-B, III-D in problem (9), and following Assumption 4, we state the overall problem to be solved in a receding horizon fashion, assuming to start from a known initial condition, i.e. $x_0^{\text{true}} = z_0^{\text{true}} = s_{0|0}$, and $\alpha_{0|0} = 0$

$$\min_{\mathbf{v}, \mathbf{s}, \boldsymbol{\alpha}, \boldsymbol{\Lambda}} \quad \mathbb{E}_{W_k}\left[\sum_{i=0}^{N-1} l_k(\bar{x}_{i|k}, \bar{u}_{i|k}) + l_f(\bar{x}_{N|k})\right] \tag{27a}$$

s.t.

$$\bar{x}_{i+1|k} = A(\bar{\theta}_k)\bar{x}_{i|k} + B(\bar{\theta}_k)\bar{u}_{i|k} + w_{i|k}, \tag{27b}$$

$$\bar{u}_{i|k} = K\bar{x}_{i|k} + v_{i|k}, \tag{27c}$$

$$W_k = [w_{0|k}^\top, \dots, w_{N-1|k}^\top]^\top \sim Q_{W_k}, \tag{27d}$$

$$F s_{i|k} \le \mathbf{1}_{n_x,1} - f_{k+i} - \alpha_{i|k}\bar{f}, \tag{27e}$$

$$G(K s_{i|k} + v_{i|k}) \le \mathbf{1}_{n_u,1} - g_{k+i} - \alpha_{i|k}\bar{g}, \tag{27f}$$

$$H_T s_{N|k} \le \mathbf{1}_{n_f,1} - \alpha_{N|k} h_T, \tag{27g}$$

$$\Lambda_{i|k}^j h_\theta + H_z d_{i|k}^j - \alpha_{i+1|k}\mathbf{1}_{r,1} \le 0, \quad j = 1, \dots, v_1, \tag{27h}$$

$$H_z D_{i|k}^j = \Lambda_{i|k}^j H_\theta, \quad j = 1, \dots, v_1, \tag{27i}$$

$$\Lambda_{i|k}^j \in \mathbb{R}_{\ge 0}^{r \times q}, \quad j = 1, \dots, v_1, \tag{27j}$$

$$\alpha_{i+1|k} \ge 0, \tag{27k}$$

$$\bar{x}_{0|k} = x_k^{\text{true}}, \; s_{0|k} = s_{1|k-1}, \; \alpha_{0|k} = \alpha_{1|k-1}, \tag{27l}$$

where we optimize over the input sequence $\mathbf{v} = \{v_{0|k}, \dots, v_{N-1|k}\}$, and the variables determining the nominal tube $\{\mathbf{Z}_{i|k}\}_{i=1}^N$, i.e. $\mathbf{s} = \{s_{1|k}, \dots, s_{N|k}\}$, and $\boldsymbol{\alpha} = \{\alpha_{1|k}, \dots, \alpha_{N|k}\}$. Additionally, we optimize over the dual variables $\boldsymbol{\Lambda} = \{\Lambda_{0|k}^j, \dots, \Lambda_{N-1|k}^j\}_{j=1}^{v_1}$ needed for nominal tube containment (9h), which is expressed by constraints (27h), (27i), and (27j) (see Section III-A). The cost function expectation (27a) is taken with respect to a predicted noise sequence W_k whose distribution Q_{W_k} is defined by the conditional distribution $p([w_k^\top, \dots, w_{k+N-1}^\top]^\top | [w_0^\top, \dots, w_{k-1}^\top]^\top)$ (see Assumption 3). Conditions (9e), (9f) are expressed via (27e), (27f), ensuring that the sequences $\mathbf{s}$ and $\boldsymbol{\alpha}$ are constrained to lie within the tightened state and input constraints. At each time step the first element of the sequence is initialized at the shifted solution from the previously optimized predicted trajectory (27l). This ensures that initial containment of the true unknown nominal state z_k^{true} in (9j) is always guaranteed by construction. The measured state x_k^{true} only enters constraints (27b) and (27c), which are the predicted state and input sequences computed with the current parameter estimate $\bar{\theta}_k$ at time k used to evaluate the cost function along the horizon, and therefore introduce *indirect* feedback in the MPC optimization problem [19]. Condition (27g) imposes containment in the terminal set (as in (9g)) satisfying Assumption 4. Finally, the nonnegativity constraint in (27k) guarantees a sequence of well-defined sets determining the nominal tube.

Note that the computational complexity of the proposed formulation is similar to [11] since the RPRS computations are all carried out offline. While the choice of homothetic tubes increases the number of optimization variables with respect to [19], the proposed approach can handle the presence of model mismatch and therefore guarantees constraints to be fulfilled within the prescribed probability level.

IV. Analysis of Closed-Loop Properties

The theorems presented in this section establish recursive feasibility of the control scheme based on (27), and closed-loop chance constraint satisfaction of the true unknown system thanks to a combined use of homothetic tubes for handling parametric uncertainty and of indirect feedback. Furthermore, we derive an average asymptotic performance bound on the $l2$-norm of the state in the case of quadratic cost functions and i.i.d. noise sequences.

A. Recursive feasibility and closed-loop properties

Theorem 1: (Recursive feasibility and closed-loop chance constraint satisfaction)
Consider system (1) under the control law (5) using the optimal input sequence $\mathbf{v}^*$ resulting from (27). If Assumptions 2 and 4 hold, $\theta^{\text{true}} \in \Theta$, and the optimization problem (27) is feasible for $x_0^{\text{true}} = z_0^{\text{true}} = s_{0|0}$ and $\alpha_{0|0} = 0$, then:

(i) Problem (27) is recursively feasible.

(ii) The true state x_k^{true} and input u_k^{true} satisfy the closed-loop chance constraints (4).

Proof:

(i) Let $\mathbf{v}^* = \{v_{0|k}^*, \dots, v_{N-1|k}^*\}$ be the optimal solution of optimization problem (27) at time-step k, with $\mathbf{s}^* = \{s_{0|k}^*, \dots, s_{N|k}^*\}$ and $\boldsymbol{\alpha}^* = \{\alpha_{0|k}^*, \dots, \alpha_{N|k}^*\}$ satisfying stage-wise constraints (27e), (27f), nominal tube containment (27h)-(27j) and terminal condition (27g). We can construct the following admissible nominal tube $\{s_{i|k}^*\} \oplus \alpha_{i|k}^* \bar{\mathbf{Z}}$, $i = 0, \dots, N$. The goal is to find a candidate solution $\tilde{\mathbf{v}} = \{\tilde{v}_{0|k+1}, \dots, \tilde{v}_{N-1|k+1}\}$ which similarly satisfies stage-wise and terminal constraints for the next time-step $k+1$. We choose the following candidate solution by shifting $\mathbf{v}^*$, and applying the terminal admissible control input $\tilde{v}_{N-1|k+1} = 0$, obtaining $\tilde{\mathbf{v}} = \{v_{1|k}^*, \dots, v_{N-1|k}^*, 0\}$. Then, the resulting candidate nominal tube at time-step $k+1$ is admissible since the first $N-1$ steps are the shifted solution $\{s_{i|k}^*\} \oplus \alpha_{i|k}^* \bar{\mathbf{Z}}$, $i = 1, \dots, N$, and the N-th step $A_{CL}(\{s_{N|k}^*\} \oplus \alpha_{N|k}^* \bar{\mathbf{Z}})$ satisfies constraint (27g) due to Assumption 4.

(ii) Due to probabilistic containment ensured by the sets in the stochastic error tube, we have that $\Pr(e_k \in \mathbf{E}_k \mid e_0^{\text{true}}) \ge p_x$, $\forall \theta \in \Theta$, $\forall k \in [1:T]$, and therefore this condition holds also for the true error state e_k^{true} evolving with respect to the unknown true parameter θ^{true}. Then, due to feasibility of problem (27), the nominal state tube $\mathbf{Z}_{0|k}$ contains the true nominal state z_k^{true}, i.e. $z_k^{\text{true}} \in \mathbf{Z}_{0|k} \subseteq \mathcal{X} \ominus \mathbf{E}_k$, and therefore the true state $x_k^{\text{true}} = e_k^{\text{true}} + z_k^{\text{true}}$ satisfies $\Pr(x_k^{\text{true}} \in \mathcal{X} \mid x_0^{\text{true}}) \ge p_x$. The same result can be derived for the input.

■

B. Average asymptotic cost bound

We now consider the particular case in which the cost function (27a) is quadratic, i.e.

$$l_k(x, u) = ||x||_Q^2 + ||u||_R^2, \tag{28a}$$

$$l_f(x) = ||x||_P^2, \tag{28b}$$

where $Q \succeq 0$, $R \succ 0$, and P satisfies the following condition

$$A_{CL}(\theta)^\top P A_{CL}(\theta) - P \preceq -Q - K^\top RK, \quad \forall\theta \in \Theta. \quad (29)$$

Furthermore, we assume that the system is affected by zero-mean i.i.d. noise sequences, i.e. $\mathbb{E}[w_k] = 0, \mathrm{Var}[w_k] = \Sigma_w$, $\forall k \geq 0$, and therefore the expected value of the overall objective can be explicitly computed in closed-form. This enables an analysis of the closed-loop true state x_k^{true} in terms of its $l2$-norm, which reflects its energy, and for which we provide an average performance asymptotic bound.

Theorem 2: (Average asymptotic $l2$-norm bound) Consider system (1) subject to i.i.d. disturbances under the control law (5) resulting from problem (27) using cost function (28), (29). There exist constants $c_0, c_1 \in \mathbb{R}_{>0}$ such that for given $\epsilon_0, \epsilon_1 \in \mathbb{R}_{>0}$:

$$\lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T} ||x_k^{\text{true}}||_2^2\right] \leq \frac{\frac{(1+\epsilon_1)(1+\frac{1}{\epsilon_0})c_1}{\mu}||\Delta\theta^{\max}||_2^2 + \mathrm{tr}(\Sigma)}{\lambda_{min}(\bar{Q}) - \epsilon_0 c_0},$$

where $\frac{1}{\mu} > \sup_{(z,Kz+v)\in\mathcal{Z}_\infty\times\mathcal{V}_\infty} ||D(z, Kz + v)||_2^2$, with $D(\cdot,\cdot)$ defined in (12), and $\Delta\theta^{\max}$ is the diameter of the set Θ. The term Σ is defined as $\Sigma = P\Sigma_w + \epsilon_0\Sigma_0 + (1 + \frac{1}{\epsilon_0})\Sigma_1$, where $P\Sigma_w$ is the cost incurred under no model mismatch. Conversely, $\Sigma_0, \Sigma_1 \succeq 0$ arise due to parametric uncertainty and are functions of the variance matrix Σ_w (see equations (33), and (34) in the Appendix for definitions). Finally, $\lambda_{min}(\bar{Q})$ denotes the maximum eigenvalue of $\bar{Q} = Q + KRK^T$, and ϵ_0 is chosen such that $\lambda_{min}(\bar{Q}) - \epsilon_0 c_0 > 0$, while ϵ_1 can be chosen to be arbitrarily small.

Proof: Proof details are given in Appendix VII-A. ■

The two terms in the performance bound provide an explicit characterization of the unavoidable cost incurred due to model mismatch, and due to the presence of stochastic noise.

Remark 7: (*Case with no model mismatch*) The development in Appendix VII-A shows that in the absence of model mismatch, we recover the same expected cost decrease bound shown in [19]. As a consequence, we also obtain the same average asymptotic cost bound with cost matrices (28).

Remark 8: (*Effect of parameter learning scheme*) Note that in the proof details in Appendix VII-A, we construct based on the properties of Θ a worst-case bound for the term

$$\lim_{T\to\infty} \frac{1}{T}\sum_{k=0}^{T} ||\theta^{\text{true}} - \bar{\theta}_k||_2^2.$$

In doing so, we do not leverage the properties of the learning scheme chosen to update the point estimate of θ that can potentially provide conditions for convergence, and therefore improve the performance bound.

V. Numerical results

A. Illustrative example

We first make use of an illustrative example for demonstrating the properties of the presented control approach. The considered model is of the form

$$x_{k+1} = Ax_k + B(\theta)u_k + w_k, \quad (30)$$

with $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$. Uncertainty affects only the input matrix, i.e. we consider the case of misspecified actuator gains, where $B = B_0 + \theta B_1$, and $B_0 = B_1 = \begin{bmatrix} 0.5 & 1 \end{bmatrix}^\top$. The additive stochastic disturbance affecting the system is i.i.d. Gaussian distributed as $w_k \sim \mathcal{N}\left(0, \begin{bmatrix} 0.3 & 0.5 \\ 0.5 & 1 \end{bmatrix}\right)$. We study the behavior of our proposed approach in terms of constraint violation by varying both the level of chance constraint satisfaction and the amount of model mismatch. The system is subject to state chance constraints on the second dimension $\Pr(|[x_k]_2| \leq 3) \geq p_x$, such that the probability level p_x belongs to the set $\{0.85, 0.9, 0.95\}$. The model mismatch is bounded and contained in the interval $\Theta_\alpha := [-\alpha, 0]$, where $\alpha \in \{0.04, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4\}$ is a parameter which we use to vary the magnitude of the considered model mismatch.

The associated MPC problem in (27) is solved in a receding horizon fashion, where we choose the prediction horizon to have length $N = 30$. The cost function is chosen to be quadratic in the state and in the input as in (28), and the weights are set as $Q = I_2$, and $R = 1$. The objective is computed by fixing the parameter estimate to $\bar{\theta}_\alpha = -\alpha$, while the unknown true parameter is $\theta^{\text{true}} = 0$. Constraints are tightened, by constructing k-step RPRS based on marginal distributions following the procedure outlined in Section III-B.1, and with respect to which we compute a polytopic terminal set as in [11], satisfying Assumption 4.

We conduct numerical simulations comparing our approach (RSMPC) with the nominal stochastic MPC (SMPC) scheme in [19] that is not designed to handle the presence of a model mismatch. In this framework, we define the nominal model with respect to $\bar{\theta}_\alpha$, and therefore as α increases so does the unaccounted amount of model mismatch. For each pair α and p_x, we run simulations for $N_s = 1000$ noise sequence realizations over a time horizon of length $T = 100$, and we compute the empirical constraint satisfaction $N_c(k)$ for each time-step $k \in [0, T]$ as

$$N_c(k) = \#(|[x_k^{\text{true}}]_2| \leq 3) \;/\; N_s \cdot 100 \;[\%],$$

i.e. the percentage of times the true simulated trajectory satisfies the constraint at time-step k divided by the total number of simulations. Then, N_c is obtained as the minimum over all time-steps $N_c = \min_{k\in(0:T)} N_c(k)$.

Figure 3 depicts N_c against the percentage of model mismatch, defined as 100α [%]. Each plot shows the behavior of the proposed scheme compared with the SMPC scheme for a fixed value of p_x. We observe that as model mismatch increases, the ability of SMPC to satisfy the imposed probability level decreases until it falls below the satisfaction threshold. On the other hand, the proposed approach has similar behavior to the SMPC scheme for small mismatch and becomes slightly more conservative only for larger values of α.

Additionally, we investigate the case in which we consider the system in (30) to be affected by correlated noise $\tilde{w}_k$ that we obtain via the following linear dynamics

$$\tilde{w}_{k+1} = A_w\tilde{w}_k + w_k,$$

where w_k is i.i.d. Gaussian noise as defined for system (30),

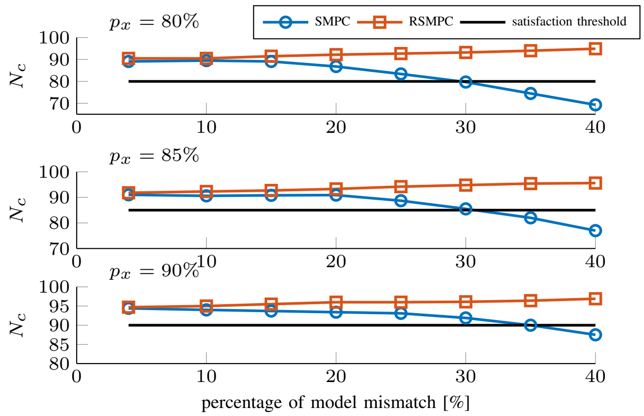

Fig. 3. Minimum empirical constraint satisfaction N_c computed for $N_s = 1000$ *i.i.d.* noise realizations, for different levels of model mismatch and imposed chance constraint probability level.

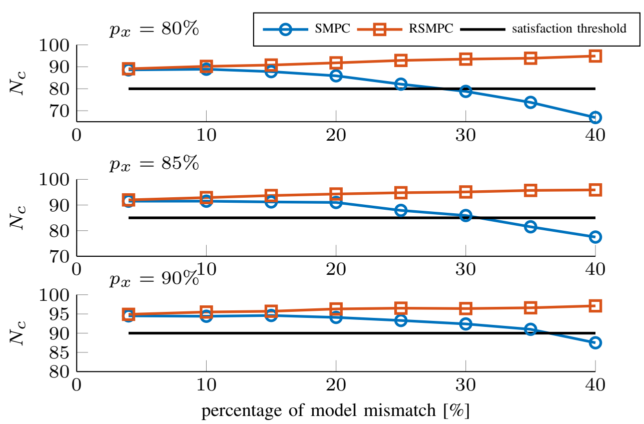

Fig. 4. Minimum empirical constraint satisfaction N_c computed for $N_s = 1000$ *correlated* noise realizations, for different levels of model mismatch and imposed chance constraint probability level.

and $A_w = \begin{bmatrix} 0.3 & 0.001 \\ 0 & 0.5 \end{bmatrix}$. In Figure 4, we show again a comparison with respect to SMPC where the k-step PRS take into account the correlation, but whose construction is again unaware of parametric uncertainty. For this example, the introduced correlation via the dynamics A_w did not introduce any further conservatism in the construction of the robustified k-step PRS, and therefore the behavior of RSMPC is similar to the i.i.d case.

B. Building temperature control

Motivated by the increasing interest as an application for MPC [35], [36], the case of a building temperature control problem is considered in the following. The goal is to maintain a predefined temperature in four adjacent rooms, for which fluctuations are controlled by heating/cooling units and vary according to the interaction between rooms and the outside. The system dynamics depends on physical parameters that are often not precisely known, and therefore we conduct robustness tests with respect to parameters of interest, e.g. thermal conductance. The dynamics is also subject to the effect of the uncertain outside temperature, which we model as an additive disturbance sequence, correlated in time. The system has the following form:

$$x_{k+1} = A(\theta)x_k + Bu_k + B_w w_k,$$

where the state $x_k \in \mathbb{R}^4$ captures the room temperatures, the input $u_k \in \mathbb{R}^4$ controls each room, and the disturbance sequence representing outdoor temperature fluctuations $W = [w_0, \dots, w_T]^\top \sim \mathcal{N}(\mu_W, \Sigma_W)$ is Gaussian distributed. Following Remark 2, we split it into a deterministic sequence $\overline{W} = \mu_W$, and a zero-mean stochastic sequence $\widetilde{W} = W - \overline{W}$.

Model uncertainty is represented by the parameter $\theta = [\theta_1, \theta_2]^\top \in \Theta_{\text{rooms}} \subseteq \mathbb{R}^2$, with $\Theta_{\text{rooms}} = \{\theta \mid ||\theta||_\infty \leq 1\}$. By considering an error on the thermal conductance between rooms 1-2, and 1-3, uncertainty only affects matrix A, while B and B_w are assumed to be known. Therefore, we have that $A(\theta) = A_0 + \theta_1 A_1 + \theta_2 A_2$, where A_0 is the nominal matrix, and A_1, A_2 are computed by perturbing A_0 by $\pm 10\%$. Nominal matrices A_0, B, and B_w are defined as in [19].

We choose the temperature to be tracked as $T_{ref} = 21°$ for each room, and define the cost function as $l(x, u) = ||x - T_{ref}\mathbf{1}_{4,1}||_Q^2 + ||u||_1$, where $Q = 50I_4$. The system is subject to the following state chance constraints for dimensions $\{j\}_{j=1}^4$:

$$\Pr([x_k]_j \geq 20°) \geq p_x, \quad \Pr([x_k]_j \leq 22°) \geq p_x,$$

where $p_x = 0.9$. By choosing base set $\bar{\mathbf{Z}}$ aligned with the state constraints, we can design a non-conservative constraint tightening by constructing a half-space k-step RPRS for each dimension $j = 1, \dots, 4$ at probability level p_x (see Remark 6). The system is also subject to input chance constraints for dimensions $j = 1, \dots, 4$:

$$\Pr([u_k]_j \geq -4.5kW) \geq p_u, \quad \Pr([u_k]_j \leq 4.5kW) \geq p_u,$$

where $p_u = 0.99$. In this case, the half-space input constraint tightening of probability level p_u will determine, according to Boole's inequality, a joint chance constraint satisfaction level of at least 0.96. Furthermore, the polytopic terminal set is computed with respect to the tightened state and input constraints similar to the illustrative example in section V-A.

Simulations depicted in Figures 5 and 6 are carried out over a period of $T = 29$ hours, for which we show the closed-loop behavior and the prediction at the last time-step. We average results over 10000 outdoor temperature sequence realizations that are shown in terms of mean and 2 standard deviations in the top subplot of Figure 5. A similar representation of the state corresponding to room 4 and the input is given in the two subplots below, where the 100%-quantiles contain all closed-loop and predicted trajectories. Constraint violations are visible particularly when the input action tries to counteract low outdoor temperature fluctuations. We compute the minimum empirical constraint satisfaction out of the 10000 simulation scenarios and obtain 91.01% for the closed-loop state and 100% for the closed-loop input, falling close to the predefined levels p_x and p_u. In Figure 6 we observe the behavior of the state and input nominal tube boundaries corresponding to room 4, for which we plot the median behavior and the 100%-quantiles to show that for all simulated noise realizations, the tightened constraints are always satisfied. The closed-loop nominal tube boundaries are in median non-conservative and tend to grow in prediction when approaching the terminal set, which is constructed to guarantee containment of the reference temperature.

Finally, we provide a study of the behavior of the RSMPC scheme in terms of closed-loop cost for different combinations of parametric uncertainty and of probability level of chance constraint satisfaction. In Figure 7, we depict the percentage of cost increase with respect to a nominal SMPC scheme simulated with no model mismatch. For each pair of $\alpha\Theta_{\text{rooms}}$, $\alpha \in \{0.2, 0.6, 1\}$ and $p_x = p_u \in \{0.8, 0.85, 0.9, 0.92, 0.95, 0.97\}$, we show mean and 2 standard deviations computed with respect to 1000 simulations. While the influence of parametric uncertainty is particularly noticeable for large values of α, the relative cost increase is generally very small.

VI. Conclusions

A model predictive control scheme for the control of systems affected by bounded parametric uncertainty and additive stochastic noise - with potentially unbounded support - was presented in this paper. The effects of the sources of uncertainty are separated by splitting the dynamics into two components: the first is only affected by bounded parametric uncertainty, dealt with by constructing a homothetic tube along the MPC prediction horizon, which we refer to as the nominal tube. The second evolves autonomously with uncertain dynamics and is perturbed by additive stochastic noise that is handled by means of the stochastic error tube, i.e. a sequence of k-step RPRS for which we present a synthesis procedure both for i.i.d. and correlated noise sequences. The tubes, and the additional use of indirect feedback, provide recursive feasibility and closed-loop chance constraint satisfaction of the proposed control scheme while allowing for using pointwise estimate updates of the uncertain parameters to compute the cost function. Potential future research involves a formal integration of an online learning scheme while maintaining probabilistic constraint satisfaction guarantees. Finally, we compute a bound for the average asymptotic $l2$-norm of the state, under the assumption of i.i.d. additive noise sequences affecting the system, and quadratic cost functions. Results are demonstrated on both an illustrative example, and on a building temperature control problem.

Acknowledgement

The authors would like to thank Dr. Johannes Koehler for the technical discussions and for providing valuable feedback.

VII. Appendix

A. Proof of Theorem 2

In the following, a performance analysis of the $l2$-norm of the closed-loop state x_k^{true} is carried out by means of an asymptotic analysis of its average behavior. The idea is to first quantify the expected cost difference between two consecutive time steps k and $k+1$ by providing a bound in expectation, which in turn is used to show that the average asymptotic $l2$-norm is bounded.

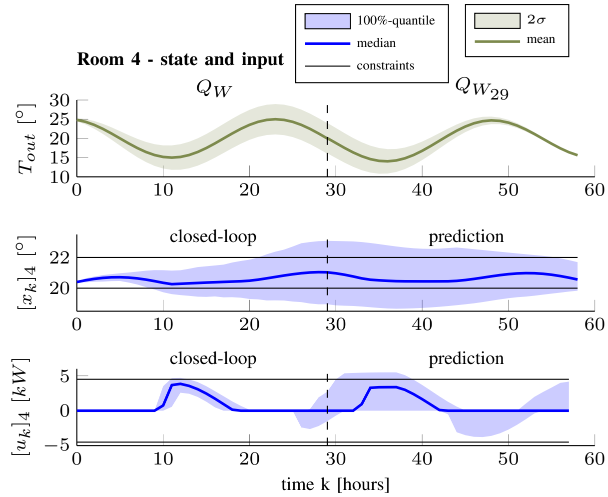

Fig. 5. Mean and 2 standard deviations computed over 10000 realizations of outdoor temperature sequences depicted in the top subplot. The middle and bottom subplots show the evolution of median and 100%-quantiles of the state and input of room 4, respectively. The black horizontal lines represent the constraints, while the vertical dashed line separates closed-loop behavior from prediction.

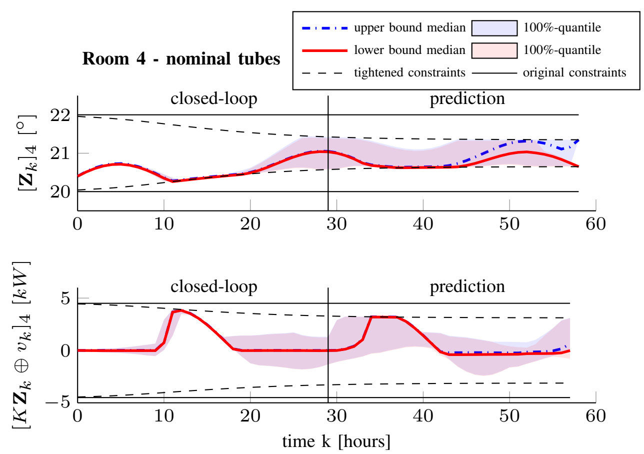

Fig. 6. Median and 100%-quantiles of the state and input nominal tube boundaries for 10000 realizations of outdoor temperature sequences. In blue we depict the nominal tube upper bound and in red the lower bound. The black continuous lines represent the original constraints and the dashed lines are the tightened constraints with respect to the k-step RPRS.

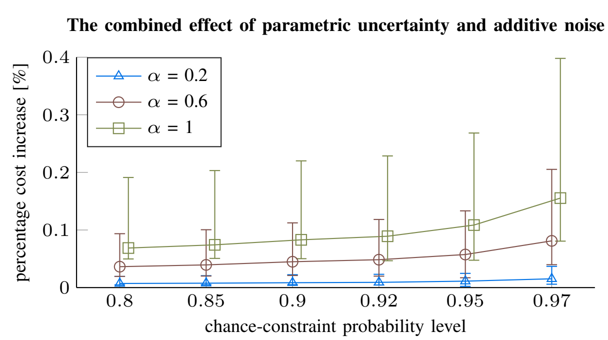

Fig. 7. Percentage of overall closed-loop cost increase expressed in terms of mean and 2 standard deviations computed over 1000 outdoor temperature realizations for each pair of $\alpha\Theta$ and $p_x = p_u$.

We assume that at time-step $k+1$ the system evolves under a shifted sequence $\bar{V} = \{\bar{v}_0, \dots, \bar{v}_{N-1}\} = \{\hat{v}_1, \dots, \hat{v}_{N-1}, 0\}$, where $\{\hat{v}_i\}_{i=0}^{N-1}$ is the optimal control sequence at time-step k, obtained by solving problem (27). We then define the predicted state sequence at time k as $\{\hat{x}_0, \dots, \hat{x}_N\}$, and at time $k+1$ as $\{\bar{x}_0, \dots, \bar{x}_N\}$, evolving under $\bar{V}$. Furthermore, it holds that $w_k \overset{\text{d}}{=} \hat{w}_0$, i.e. the closed and open-loop disturbance realizations are drawn from the same distribution, therefore

$$\begin{aligned}\bar{x}_0 &\overset{\text{d}}{=} \hat{x}_1 + (A(\theta^{\text{true}}) - A(\bar{\theta}_k))x_k^{\text{true}} + (B(\theta^{\text{true}}) - B(\bar{\theta}_k))u_k^{\text{true}} \\ &= \hat{x}_1 + \tilde{x}_k.\end{aligned}$$

The predicted states sequences at, respectively, time-step k and $k+1$ have the following relation for $i = 0, \dots, N-1$

$$\bar{x}_i \overset{\text{d}}{=} \hat{x}_{i+1} + \delta\hat{x}_i,$$

and at the last predicted time-step $i = N$ we have $\bar{x}_N \overset{\text{d}}{=} A_{CL}(\bar{\theta}_k)\hat{x}_N + \hat{w}_N + \delta\hat{x}_N$. The terms $\{\delta\hat{x}_i\}_{i=0}^N$ represent the cumulated prediction error due to model mismatch, and evolve according to the following dynamics

$$\delta\hat{x}_0 = \tilde{x}_k \tag{31a}$$

$$\begin{aligned}\delta\hat{x}_{i+1} = A_{CL}(\bar{\theta}_{k+1})\delta\hat{x}_i + (A_{CL}(\bar{\theta}_{k+1}) - A_{CL}(\bar{\theta}_k))\hat{x}_{i+1} \\ + (B(\bar{\theta}_{k+1}) - B(\bar{\theta}_k))\hat{v}_{i+1}.\end{aligned} \tag{31b}$$

We observe two sources of model mismatch propagated along the horizon, as the model parameters are updated from time-step k to $k+1$. The initial prediction error $\tilde{x}_k$, depending on the difference between the true unknown model parameter θ^{true} and the previous estimate $\bar{\theta}_k$, and the propagation error in (31b), depending on the difference between the updated estimate $\bar{\theta}_{k+1}$ and the previous $\bar{\theta}_k$. The procedure for bounding the expected cost difference is outlined in (32), in which conditioning is omitted for the sake of readability, and $\bar{Q} = Q + KRK^\top$. All inequalities make use of the i.i.d. assumption on the additive disturbance sequence, which also entails that $\hat{x}_{i+1}$ and δx_i are independent given the initial condition x_k^{true}, and therefore allows for applying Lemma 2. We now analyse each term independently, starting from (1)

$$\begin{aligned}&\mathbb{E}\left[||A_{CL}(\hat{\theta})\hat{x}_N||_P^2 + ||\hat{x}_N||_Q^2 + ||\hat{u}_N||_R^2 - ||\hat{x}_N||_P^2\right] \\ &= \mathbb{E}\left[||\hat{x}_N||^2_{\underbrace{A_{CL}(\hat{\theta})^\top P A_{CL}(\hat{\theta}) + Q + KRK^\top - P}_{=\tilde{P}\preceq 0}}\right] \leq 0,\end{aligned}$$

which holds thanks to (29). Then, the second term (2) vanishes since we are subtracting the shifted sequence

$$\mathbb{E}\left[\sum_{i=0}^{N-2} ||\hat{x}_{i+1}||_Q^2 + ||\hat{u}_{i+1}||_R^2 - \left(\sum_{i=1}^{N-1} ||\hat{x}_i||_Q^2 + ||\hat{u}_i||_R^2\right)\right] = 0.$$

For the third term (3), we can explicitly evaluate the expected value since x_k, u_k are given and the disturbance distribution is known

$$\begin{aligned}&\mathbb{E}\left[-||x_k^{\text{true}}||_Q^2 - ||u_k^{\text{true}}||_R^2 + ||\hat{w}_N||_P^2\right] \\ &= \text{tr}(P\Sigma_w) - ||x_k^{\text{true}}||_Q^2 - ||u_k^{\text{true}}||_R^2.\end{aligned}$$

The last two terms are costs incurred due to the presence of a model mismatch. The expected value of the quadratic form (4) is explicitly evaluated

$$\begin{aligned}&\mathbb{E}\left[||A_{CL}(\hat{\theta})\hat{x}_N||_P^2 + \sum_{i=0}^{N-1} ||\hat{x}_{i+1}||_Q^2 + ||\hat{u}_{i+1}||_R^2\right] \\ &= \mathbb{E}\left[||\hat{x}_N||_{\tilde{P}}^2 + ||\hat{x}_N||_P^2 + \sum_{i=1}^{N-1} ||\hat{x}_i||_Q^2 + ||\hat{u}_i||_R^2\right] \\ &\leq ||\mathbb{E}[\hat{x}_N]||_P^2 + \sum_{i=1}^{N-1} ||\mathbb{E}[\hat{x}_i]||_Q^2 + ||\mathbb{E}[\hat{u}_i]||_R^2 \\ &\quad + \text{tr}(P\text{Var}(\hat{x}_N)) + \sum_{i=1}^{N-1} \text{tr}(Q\text{Var}(\hat{x}_i)) + \text{tr}(R\text{Var}(\hat{u}_i)) \\ &\leq c_0||x_k^{\text{true}}||_2^2 + \text{tr}(P\bar{\Sigma}_w^N) + \text{tr}(Q\sum_{i=1}^{N-1}\bar{\Sigma}_w^i) + \text{tr}(RK\sum_{i=1}^{N-1}\bar{\Sigma}_w^i K^\top) \\ &= c_0||x_k^{\text{true}}||_2^2 + \text{tr}(\Sigma_0),\end{aligned}$$

where the first inequality uses $\tilde{P} \preceq 0$. The second inequality uses the same argument used in [11], i.e. the cost associated with the expected values of the predicted states is a continuous, piecewise quadratic function in x_0, $\forall \theta \in \Theta$ [37]. Therefore, it can be upper bounded with a quadratic function of the initial condition for some $c_0 > 0$. The variances can be expressed exactly as a function of $\bar{\Sigma}_w^i = \sum_{l=0}^{i-1} A_{CL}(\bar{\theta}_k)^l \Sigma_w A_{CL}^\top(\bar{\theta}_k)^l$, and we define

$$\Sigma_0 = P\bar{\Sigma}_w^N + Q\sum_{i=1}^{N-1}\bar{\Sigma}_w^i + RK\sum_{i=1}^{N-1}\bar{\Sigma}_w^i K^\top. \tag{33}$$

We then obtain a bound for (5)

$$\begin{aligned}&\mathbb{E}\left[||\delta\hat{x}_N||_P^2 + \sum_{i=0}^{N-1}||\delta\hat{x}_i||_{\bar{Q}}^2\right] = ||\mathbb{E}[\delta\hat{x}_N]||_P^2 + \sum_{i=0}^{N-1}||\mathbb{E}[\delta\hat{x}_i]||_{\bar{Q}}^2 \\ &\quad + \text{tr}(P\text{Var}(\delta\hat{x}_N)) + \sum_{i=0}^{N-1}\text{tr}(\bar{Q}\text{Var}(\delta\hat{x}_i)) \\ &\leq c_1||\tilde{x}_k||_2^2 + \text{tr}(P\text{Var}(\delta\hat{x}_N)) + \sum_{i=0}^{N-1}\text{tr}(\bar{Q}\text{Var}(\delta\hat{x}_i)) \\ &= c_1||\tilde{x}_k||_2^2 + \text{tr}(P\bar{\bar{\Sigma}}_w^N) + \text{tr}(\bar{Q}\sum_{i=0}^{N-1}\bar{\bar{\Sigma}}_w^i) = c_1||\tilde{x}_k||_2^2 + \text{tr}(\Sigma_1),\end{aligned}$$

where again we make use of the bound on the cost of the expected values of $\delta\hat{x}_i$, which is a function of the initial condition $\tilde{x}_k$ (31a). The variances can be expressed as a function of $\bar{\Sigma}_w^i$, and therefore depend on the known noise variance Σ_w, resulting in the following relation: $\bar{\bar{\Sigma}}_w^i = \sum_{j=1}^i A_{CL}(\bar{\theta}_{k+1})^{i-j}(A_{CL}(\bar{\theta}_{k+1}) - A_{CL}(\bar{\theta}_k))\bar{\Sigma}_w^j(A_{CL}(\bar{\theta}_{k+1}) - A_{CL}(\bar{\theta}_k))^\top A_{CL}^\top(\bar{\theta}_{k+1})^{i-j}$. Finally, we define

$$\Sigma_1 = P\bar{\bar{\Sigma}}_w^N + \bar{Q}\sum_{i=0}^{N-1}\bar{\bar{\Sigma}}_w^i. \tag{34}$$

Putting everything together and rearranging terms, we obtain that there exists an $\epsilon_0 > 0$ such that

$$\mathbb{E}\left[J^*(x_{k+1}^{\text{true}}, \bar{\theta}_{k+1}) \mid x_k^{\text{true}}\right] - J^*(x_k^{\text{true}}, \bar{\theta}_k)$$

$$
\begin{aligned}
&\mathbb{E}\left[J^*(x_{k+1}^{\text{true}},\bar\theta_{k+1}) \mid x_k^{\text{true}}\right] - J^*(x_k^{\text{true}},\bar\theta_k) \leq \mathbb{E}\left[J(x_{k+1}^{\text{true}},\bar\theta_{k+1},\bar V) \mid x_k^{\text{true}}\right] - J^*(x_k^{\text{true}},\bar\theta_k)\\
&= \mathbb{E}\left[||\bar x_N||_P^2 + \sum_{i=0}^{N-1}||\bar x_i||_Q^2 + ||\bar u_i||_R^2\right] - \mathbb{E}\left[\left(||\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\hat x_i||_Q^2 + ||\hat u_i||_R^2\right)\right]\\
&= \mathbb{E}\left[||A_{CL}(\hat\theta)\hat x_N + \hat w_N + \delta\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\hat x_{i+1} + \delta\hat x_i||_Q^2 + ||K(\hat x_{i+1}+\delta\hat x_i) + \hat v_{i+1}||_R^2\right] - \mathbb{E}\left[||\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\hat x_i||_Q^2 + ||\hat u_i||_R^2\right]\\
&\leq \mathbb{E}\left[(1+\epsilon_0)||A_{CL}(\hat\theta)\hat x_N||_P^2 + ||\hat w_N||_P^2 + (1+\frac{1}{\epsilon_0})||\delta\hat x_N||_P^2\right] + \mathbb{E}\left[\sum_{i=0}^{N-1}(1+\epsilon_0)(||\hat x_{i+1}||_Q^2 + ||\hat u_{i+1}||_R^2) + (1+\frac{1}{\epsilon_0})||\delta\hat x_i||_Q^2\right]\\
&-\mathbb{E}\left[||\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\hat x_i||_Q^2 + ||\hat u_i||_R^2\right]\\
&= \underbrace{\mathbb{E}\left[||A_{CL}(\hat\theta)\hat x_N||_P^2 + ||\hat x_N||_Q^2 + ||\hat u_N||_R^2 - ||\hat x_N||_P^2\right]}_{(1)} + \underbrace{\mathbb{E}\left[\sum_{i=0}^{N-2}||\hat x_{i+1}||_Q^2 + ||\hat u_{i+1}||_R^2 - \left(\sum_{i=1}^{N-1}||\hat x_i||_Q^2 + ||\hat u_i||_R^2\right)\right]}_{(2)}\\
&+\underbrace{\mathbb{E}\left[-||x_k^{\text{true}}||_Q^2 - ||u_k^{\text{true}}||_R^2 + ||\hat w_N||_P^2\right]}_{(3)} + \epsilon_0\underbrace{\mathbb{E}\left[||A_{CL}(\hat\theta)\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\hat x_{i+1}||_Q^2 + ||\hat u_{i+1}||_R^2\right]}_{(4)} + (1+\frac{1}{\epsilon_0})\underbrace{\mathbb{E}\left[||\delta\hat x_N||_P^2 + \sum_{i=0}^{N-1}||\delta\hat x_i||_Q^2\right]}_{(5)}
\end{aligned}
\tag{32}
$$

$$
\begin{aligned}
&\leq \text{tr}(P\Sigma_w) - ||x_k^{\text{true}}||_Q^2 - ||u_k^{\text{true}}||_R^2\\
&\quad + \epsilon_0(c_0||x_k^{\text{true}}||_2^2 + \text{tr}(\Sigma_0)) + (1+\frac{1}{\epsilon_0})(c_1||\tilde x_k||_2^2 + \text{tr}(\Sigma_1))\\
&\leq -(\lambda_{min}(\bar Q) - \epsilon_0 c_0)||x_k^{\text{true}}||_2^2 + (1+\frac{1}{\epsilon_0})c_1||\tilde x_k||_2^2 + \text{tr}(\Sigma),
\end{aligned}
$$

with $\lambda_{min}(\bar Q) - \epsilon_0 c_0 > 0$. We can now proceed with an analysis of the asymptotic behavior of the $l2$-norm of x_k^{true}. We start by using a standard argument in stochastic MPC, i.e. repeatedly applying the law of iterated expectations, and using the expected cost difference bound:

$$
\begin{aligned}
&\mathbb{E}\left[J^*(x_T^{\text{true}},\bar\theta_T) \mid x_0^{\text{true}}\right] - J^*(x_0^{\text{true}},\bar\theta_0) \leq\\
&\mathbb{E}\left[\sum_{k=0}^{T} -C||x_k^{\text{true}}||_2^2 + (1+\frac{1}{\epsilon_0})c_1||\tilde x_k||_2^2 + \text{tr}(\Sigma) \,\middle|\, x_0^{\text{true}}\right],
\end{aligned}
$$

where $C = (\lambda_{min}(\bar Q) - \epsilon_0 c_0)$. Taking the limit for $T \to \infty$

$$
\begin{aligned}
&0 \leq \lim_{T\to\infty}\frac{1}{T}\left(\mathbb{E}\left[J^*(x_T^{\text{true}},\bar\theta_T) \mid x_0^{\text{true}}\right] - J^*(x_0^{\text{true}},\bar\theta_0)\right) \leq\\
&\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T} -C||x_k^{\text{true}}||_2^2 + (1+\frac{1}{\epsilon_0})c_1||\tilde x_k||_2^2 + \text{tr}(\Sigma) \,\middle|\, x_0^{\text{true}}\right]
\end{aligned}
$$

Then, we can derive the $l2$-norm bound on x_k^{true}

$$
\begin{aligned}
&\lim_{T\to\infty}\frac{1}{T}\left(\mathbb{E}\left[\sum_{k=0}^{T}||x_k^{\text{true}}||_2^2 \,\middle|\, x_0^{\text{true}}\right]\right)\\
&\leq \frac{\text{tr}(\Sigma)}{C} + \frac{(1+\frac{1}{\epsilon_0})c_1}{C}\lim_{T\to\infty}\frac{1}{T}\mathbb{E}[\textstyle\sum_{k=0}^{T}||\tilde x_k||_2^2 \mid x_0^{\text{true}}].
\end{aligned}
$$

This limit can be explicitly computed: in (35) we express the state as $x^{\text{true}} = z^{\text{true}} + e^{\text{true}}$, and apply Lemma 2. The split allows for bounding $D(z^{\text{true}}, Kz^{\text{true}} + v^{\text{true}})$ (12) using its supremum, which exists since $\mathcal{Z}_\infty, \mathcal{V}_\infty$ are compact sets. We then exploit the boundedness of Θ so that we can isolate the expected norm of the error state e_k^{true}, for which we know that $\mathbb{E}[e_k^{\text{true}} \mid x_0^{\text{true}}] = 0$, $\forall k \geq 0$ due to the assumption on the additive noise and since $e_0^{\text{true}} = 0$. Furthermore, since $A_{CL}(\theta^{\text{true}})$ is Hurwitz, we know that under the i.i.d. noise distribution assumption, the variance $\text{Var}(e_k \mid x_0^{\text{true}})$ converges to some matrix Σ_∞ that satisfies the following Lyapunov equation $\Sigma_\infty = A_{CL}(\theta^{\text{true}})\Sigma_\infty A_{CL}^\top(\theta^{\text{true}}) + \Sigma_w$.

B. Lemmas

Lemma 1: Vertex property [38]
Let $F(\theta, x) > 0$ be an inequality of the form

$$F(\theta, x) = F_0(\theta) + \sum_{j=0}^{m} x_j F_j(\theta) > 0,$$

where the functions $F_j(\theta)$ are affine in $\theta \in \Theta$ and Θ is a convex polytope of r vertices defined as $\Theta = \text{co}\{\theta^1, ..., \theta^r\}$. Then, the infinite set of LMIs $F(\theta, x) \succ 0$ holds $\forall \theta \in \Theta$ if and only if $F(\theta, x) \succ 0$ holds at each vertex of Θ, i.e.,

$$F(\theta, x) \succ 0, \forall \theta \in \Theta \Leftrightarrow F(\theta^i, x) \succ 0, i = 1, \dots, r.$$

Lemma 2: Fenchel-Young inequality in expectation
Consider x, y independent random variables, then for all matrices $R = R^\top \succeq 0$, and $\forall\ \epsilon > 0$

$$\mathbb{E}\left[||x+y||_R^2\right] \leq (1+\epsilon)\mathbb{E}\left[||x||_R^2\right] + (1+\frac{1}{\epsilon})\mathbb{E}\left[||y||_R^2\right]$$

Proof:

$$
\begin{aligned}
\mathbb{E}\left[||x+y||_R^2\right] &= ||\mathbb{E}[x] + \mathbb{E}[y]||_R^2 + \text{tr}(R\text{Var}(x)) + \text{tr}(R\text{Var}(y))\\
&\leq (1+\epsilon)||\mathbb{E}[x]||_R^2 + (1+\frac{1}{\epsilon})||\mathbb{E}[y]||_R^2\\
&\quad + \text{tr}(R\text{Var}(x)) + \text{tr}(R\text{Var}(y))\\
&= (1+\epsilon)\left(||\mathbb{E}[x]||_R^2 + \text{tr}(R\text{Var}(x))\right)\\
&\quad + (1+\frac{1}{\epsilon})\left(||\mathbb{E}[y]||_R^2 + \text{tr}(R\text{Var}(y))\right)\\
&\quad - \epsilon\text{tr}(R\text{Var}(x)) - \frac{1}{\epsilon}\text{tr}(R\text{Var}(y))\\
&\leq (1+\epsilon)\mathbb{E}\left[||x||_R^2\right] + (1+\frac{1}{\epsilon})\mathbb{E}\left[||y||_R^2\right],
\end{aligned}
$$

$$\begin{aligned}
&\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T}||\tilde{x}_k||_2^2 \,\middle|\, x_0^{\text{true}}\right] = \lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T}||(A(\theta^{\text{true}})-A(\bar{\theta}_k))x_k^{\text{true}}+(B(\theta^{\text{true}})-B(\bar{\theta}_k))u_k^{\text{true}}||_2^2 \,\middle|\, x_0^{\text{true}}\right]\\
&=\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T}||D(z_k^{\text{true}},Kz_k^{\text{true}}+v_k^{\text{true}})(\theta^{\text{true}}-\bar{\theta}_k)+(A_{CL}(\theta^{\text{true}})-A_{CL}(\bar{\theta}_k))e_k^{\text{true}}||_2^2 \,\middle|\, x_0^{\text{true}}\right]\\
&\leq\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T}(1+\epsilon_1)||D(z_k^{\text{true}},Kz_k^{\text{true}}+v_k^{\text{true}})(\theta^{\text{true}}-\bar{\theta}_k)||_2^2+(1+\frac{1}{\epsilon_1})||(A_{CL}(\theta^{\text{true}})-A_{CL}(\bar{\theta}_k))e_k^{\text{true}}||_2^2 \,\middle|\, x_0^{\text{true}}\right]\\
&\leq(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu}+\lim_{T\to\infty}(1+\frac{1}{\epsilon_1})\frac{1}{T}\mathbb{E}\left[\sum_{k=0}^{T}||(A_{CL}(\theta^{\text{true}})-A_{CL}(\bar{\theta}_k))e_k^{\text{true}}||_2^2 \mid x_0^{\text{true}}\right]\\
&\leq(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu}+(1+\frac{1}{\epsilon_1})\max_{\theta\in\Theta}||A_{CL}(\theta^{\text{true}})-A_{CL}(\theta)||_2^2\lim_{T\to\infty}\frac{1}{T}\sum_{k=0}^{T}\mathbb{E}\left[||e_k^{\text{true}}||_2^2 \mid x_0^{\text{true}}\right]\\
&=(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu}+(1+\frac{1}{\epsilon_1})\max_{\theta\in\Theta}||A_{CL}(\theta^{\text{true}})-A_{CL}(\theta)||_2^2\lim_{T\to\infty}\frac{1}{T}\sum_{k=0}^{T}||\mathbb{E}\left[e_k^{\text{true}} \mid x_0^{\text{true}}\right]||_2^2+\text{tr}(\text{Var}(e_k^{\text{true}} \mid x_0^{\text{true}}))\\
&=(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu}+(1+\frac{1}{\epsilon_1})\max_{\theta\in\Theta}||A_{CL}(\theta^{\text{true}})-A_{CL}(\theta)||_2^2\lim_{T\to\infty}\frac{1}{T}\sum_{k=0}^{T}\text{tr}(A_{CL}(\theta^{\text{true}})^k\Sigma_w A_{CL}^\top(\theta^{\text{true}})^k)\\
&=(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu}+(1+\frac{1}{\epsilon_1})\max_{\theta\in\Theta}||A_{CL}(\theta^{\text{true}})-A_{CL}(\theta)||_2^2\lim_{T\to\infty}\frac{\text{tr}(\Sigma_\infty)}{T}=(1+\epsilon_1)\frac{||\Delta\theta^{\max}||_2^2}{\mu},
\end{aligned}\tag{35}$$

where the first inequality makes use of a combination of the Cauchy-Schwarz inequality and of Fenchel-Young's inequality on the norm of the expected values of x and y (see Section 3.3.2 [31]). The second inequality uses the fact that for all X, Y symmetric, and positive (semi-)definite, tr $(XY) \geq 0$. ■

Lemma 3: Convex reformulation
An optimization problem of the form

$$\begin{aligned}\min_{X^{-1}}\ & -\log\det X^{-1}\\ \text{s.t.}\ \ & X - Z - A(\theta)YA(\theta)^\top \succeq 0,\ \forall\theta\in\Theta\end{aligned}$$

is equivalent to the following convex reformulation

$$\begin{aligned}\min_{X^{-1}}\ & -\log\det X^{-1}\\ \text{s.t.}\ & \begin{bmatrix} X^{-1} & X^{-1}Z & X^{-1}A(\theta^j)\\ ZX^{-1} & Z & 0\\ A(\theta^j)^\top X^{-1} & 0 & Y^{-1}\end{bmatrix}\succeq 0,\\ & \forall j\in\{1,\dots,r\},\end{aligned}$$

provided that $Y, Z \succ 0$, $Z = Z^\top$, and $A(\theta)$ is of the form $A(\theta) = A_0 + \sum_{i=1}^{p} A_i \quad [\theta]_i$, with $\theta \in \Theta := \text{co}\{\theta^1, ..., \theta^r\}$.

Proof: Pre- and post-multiply the matrix inequality by X^{-1} to obtain

$$X^{-1} - X^{-1}ZX^{-1} - X^{-1}A(\theta)YA(\theta)^\top X^{-1} \succeq 0,$$

and use the following condition for positive semi-definite matrices based on the Schur complement, i.e. if

$$\begin{aligned}&Y\succ 0,\ \text{then}\ \begin{bmatrix} X^{-1}-X^{-1}ZX^{-1} & X^{-1}A(\theta)\\ A(\theta)^\top X^{-1} & Y^{-1}\end{bmatrix}\succeq 0\\ &\Leftrightarrow\\ &X^{-1}-X^{-1}ZX^{-1}-X^{-1}A(\theta)YA(\theta)^\top X^{-1}\succeq 0.\end{aligned}$$

Applying again the Schur complement to the first diagonal block

$$Z\succ 0, X^{-1}-X^{-1}ZX^{-1}\succ 0 \Leftrightarrow \begin{bmatrix} X^{-1} & X^{-1}Z\\ ZX^{-1} & Z\end{bmatrix}\succ 0,$$

the optimization problem can be reformulated as:

$$\begin{aligned}\min_{X^{-1}}\ & -\log\det X^{-1}\\ \text{s.t.}\ & \begin{bmatrix} X^{-1} & X^{-1}Z & X^{-1}A(\theta)\\ ZX^{-1} & Z & 0\\ A(\theta)^\top X^{-1} & 0 & Y^{-1}\end{bmatrix}\succeq 0,\ \forall\theta\in\Theta\end{aligned}$$

to which we can apply Lemma 1 since the linear matrix inequality is affine in θ, and θ belongs to a convex set Θ. ■

References

[1] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*, 1999.
[2] A. Mesbah, "Stochastic Model Predictive Control: An Overview and Perspectives for Future Research," *IEEE Control Systems Magazine*, 2016.
[3] D. M. Raimondo, D. Limon, M. Lazar, L. Magni, and E. F. ndez Camacho, "Min-max Model Predictive Control of Nonlinear Systems: A Unifying Overview on Stability," *European Journal of Control*, 2009.
[4] G. Pin, D. M. Raimondo, L. Magni, and T. Parisini, "Robust Model Predictive Control of Nonlinear Systems With Bounded and State-Dependent Uncertainties," *IEEE Transactions on Automatic Control*, 2009.
[5] B. Kouvaritakis and M. Cannon, *Model Predictive Control: Classical, Robust and Stochastic*. Springer International Publishing, 2015.
[6] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained mpc for multiplicative and additive stochastic uncertainty," *IEEE Transactions on Automatic Control*, 2009.
[7] J. Fleming and M. Cannon, "Stochastic mpc for additive and multiplicative uncertainty using sample approximations," *IEEE Transactions on Automatic Control*, 2019.
[8] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, 2013.
[9] R. Soloperto, M. A. Müller, S. Trimpe, and F. Allgöwer, "Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty," *IFAC-PapersOnLine*, 2018.

[10] J. M. Manzano, D. Limon, D. Muñoz de la Peña, and J.-P. Calliess, "Robust learning-based mpc for nonlinear constrained systems," *Automatica*, 2020.
[11] M. Lorenzen, M. Cannon, and F. Allgöwer, "Robust MPC with recursive model update," *Automatica*, 2019.
[12] J. Koehler, P. Kötting, R. Soloperto, F. Allgöwer, and M. Müller, "A robust adaptive model predictive control framework for nonlinear uncertain systems," *International Journal of Robust and Nonlinear Control*, 2020.
[13] A. Parsi, A. Iannelli, and R. S. Smith, "An explicit dual control approach for constrained reference tracking of uncertain linear systems," *IEEE Transactions on Automatic Control*, 2022.
[14] R. Soloperto, J. Köhler, and F. Allgöwer, "Augmenting mpc schemes with active learning: Intuitive tuning and guaranteed performance," *IEEE Control Systems Letters*, 2020.
[15] L. Chisci, J. Rossiter, and G. Zappa, "Systems with persistent disturbances: predictive control with restricted constraints," *Automatica*, 2001.
[16] D. Mayne, M. Seron, and S. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, 2005.
[17] S. V. Raković, B. Kouvaritakis, R. Findeisen, and M. Cannon, "Homothetic tube model predictive control," *Automatica*, 2012.
[18] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Constraint-Tightening and Stability in Stochastic Model Predictive Control," *IEEE Transactions on Automatic Control*, 2017.
[19] L. Hewing, K. P. Wabersich, and M. N. Zeilinger, "Recursively feasible stochastic model predictive control using indirect feedback," *Automatica*, 2020.
[20] M. Lorenzen, M. A. Müller, and F. Allgöwer, "Stochastic model predictive control without terminal constraints," *International Journal of Robust and Nonlinear Control*, 2019.
[21] K. P. Wabersich, L. Hewing, A. Carron, and M. N. Zeilinger, "Probabilistic model predictive safety certification for learning-based control," *IEEE Transactions on Automatic Control*, 2022.
[22] A. Didier, K. P. Wabersich, and M. N. Zeilinger, "Adaptive model predictive safety certification for learning-based control," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021.
[23] D. Mayne, "Robust and stochastic model predictive control: Are we going in the right direction?" *Annual Reviews in Control*, 2016.
[24] M. B. Saltık, L. Özkan, J. H. Ludlage, S. Weiland, and P. M. Van den Hof, "An outlook on robust model predictive control algorithms: Reflections on performance and computational aspects," *Journal of Process Control*, 2018.
[25] J. Köhler, R. Soloperto, M. A. Müller, and F. Allgöwer, "A Computationally Efficient Robust Model Predictive Control Framework for Uncertain Nonlinear Systems," *IEEE Transactions on Automatic Control*, 2021.
[26] J. Köhler, E. Andina, R. Soloperto, M. A. Müller, and F. Allgöwer, "Linear robust adaptive model predictive control: Computational complexity and conservatism," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019.
[27] S. V. Raković and Q. Cheng, "Homothetic tube mpc for constrained linear difference inclusions," in *2013 25th Chinese Control and Decision Conference (CCDC)*, 2013.
[28] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, 2006.
[29] L. Hewing, A. Carron, K. P. Wabersich, and M. N. Zeilinger, "On a correspondence between probabilistic and robust invariant sets for linear systems," in *2018 European Control Conference (ECC)*, 2018.
[30] K. Löwner, "Über monotone matrixfunktionen," *Mathematische Zeitschrift*, 1934.
[31] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
[32] M. Fiacchini and T. Alamo, "Probabilistic reachable and invariant sets for linear systems with correlated disturbance," *Automatica*, 2021.
[33] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge University Press, 2012.
[34] L. Hewing and M. Zeilinger, "Cautious model predictive control using gaussian process regression," *IEEE Transactions on Control Systems Technology*, 2020.
[35] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari, "Stochastic model predictive control for building climate control," *IEEE Transactions on Control Systems Technology*, 2014.
[36] J. Drgoňa, J. Arroyo, I. C. Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," *Annual Reviews in Control*, 2020.
[37] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, 2002.
[38] P. Apkarian and H. D. Tuan, "Parameterized LMIs in control theory," *SIAM Journal on Control and Optimization*, 2000.

Elena Arcari received the bachelor's degree from University of Rome Tor Vergata, and her master's degree from ETH Zurich. She is currently a doctoral student at the Institute for Dynamic Systems and Control, ETH Zurich. Her research interests include active learning for model-based control and data-driven control.

Andrea Iannelli is an Assistant Professor at the University of Stuttgart, Germany. He completed his B.Sc. and M.Sc. degrees in Aerospace Engineering at the University of Pisa and received his PhD from the University of Bristol. In his doctoral studies, he worked on the reconciliation between robust control and dynamical systems theory, with application to linear and nonlinear aerospace systems. He was a postdoctoral researcher in the Automatic Control Laboratory at ETH Zürich, where he investigated topics at the intersection between control theory and learning, such as system identification, adaptive optimization-based control, and online learning. His research interests are centered around principled techniques for modeling, analysis, and control of uncertain dynamical systems.

Andrea Carron is a Senior Lecturer at ETH Zürich. He received bachelor's, master's, and Ph.D. degrees in control engineering from the University of Padova in Italy. During his master's and Ph.D. studies, he spent three stays abroad as a Visiting Researcher: the first at the University of California at Riverside, the second at the Max Planck Institute in Tübingen, and the third at the University of California at Santa Barbara. From 2016 to 2019, he was a Post-Doctoral Fellow with the Intelligent Control Systems Group, ETH Zurich. His research interests include safe learning, learning-based control, multi-agent systems, and robotics.

Melanie N. Zeilinger is an Associate Professor at ETH Zurich, Switzerland. She received the Diploma degree in engineering cybernetics from the University of Stuttgart, Germany, in 2006, and the Ph.D. degree with honors in electrical engineering from ETH Zurich, Switzerland, in 2011. From 2011 to 2012 she was a Postdoctoral Fellow with the Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland. She was a Marie Curie Fellow and Postdoctoral Researcher with the Max Planck Institute for Intelligent Systems, Tübingen, Germany until 2015 and with the Department of Electrical Engineering and Computer Sciences at the University of California at Berkeley, CA, USA, from 2012 to 2014. From 2018 to 2019 she was a professor at the University of Freiburg, Germany. Her current research interests include safe learning-based control, as well as optimization-based control under uncertainties, with applications to robotics and human-in-the-loop control.