Automatica 162 (2024) 111431

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Robust model predictive control with polytopic model uncertainty through System Level Synthesis☆

Shaoru Chen, Victor M. Preciado, Manfred Morari, Nikolai Matni *

Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, 19104, USA

ARTICLE INFO

Article history:
Received 9 March 2023
Received in revised form 18 September 2023
Accepted 26 October 2023
Available online 22 January 2024

Keywords:
Robust model predictive control
System Level Synthesis
Robust control

ABSTRACT

We propose a robust model predictive control (MPC) method for discrete-time linear systems with polytopic model uncertainty and additive disturbances. Optimizing over linear time-varying (LTV) state feedback controllers has been successfully used for robust MPC when only additive disturbances are present. However, it is challenging to design LTV state feedback controllers in the face of model uncertainty whose effects are difficult to bound. To address this issue, we propose a novel approach to over-approximate the effects of both model uncertainty and additive disturbances by a filtered additive disturbance signal. Using the System Level Synthesis framework, we jointly search for robust LTV state feedback controllers and the bounds on the effects of uncertainty online, which allows us to reduce the conservatism and minimize an upper bound on the worst-case cost in robust MPC. We provide a comprehensive numerical comparison of our method and representative robust MPC methods from the literature. Numerical examples demonstrate that our proposed method can significantly reduce the conservatism over a wide range of uncertainty parameters with comparable computational effort as the baseline methods.

© 2023 Published by Elsevier Ltd.

1. Introduction

In model predictive control (MPC), a finite time constrained optimal control problem (OCP) is solved at each time step and the first optimal control input is applied. When the system dynamics is uncertain, robust model predictive control explicitly takes the uncertainty into account by solving a robust OCP at each time step to guarantee that the state and control input constraints are robustly satisfied for the closed-loop system. Although dynamic programming (Borrelli et al., 2017, Chapter 15) can exactly solve the robust OCP, it suffers from prohibitive computational complexity and is impractical to use. This has motivated the development of alternative solutions to the robust OCP which aim to reach a reasonable compromise between conservatism as measured by the size of the set of feasible initial states, and computational complexity.

Robust MPC with additive disturbances: For uncertain linear time-invariant (LTI) systems with the only uncertainty arising from additive disturbances, a variety of closed-loop methods have been proposed to solve the robust OCP with feedback policies as decision variables. The use of feedback policies does not suffer from the conservatism of open-loop strategies in the presence of uncertainty (Mayne, 2014). To maintain computational tractability, closed-loop methods only search over control policies that admit a finite-dimensional parameterization. Tube-based MPC is a widely used class of robust MPC methods where a tube (normally a sequence of polytopes) is constructed such that it contains all possible system trajectories under uncertainty. Different designs of tube MPC such as rigid (Mayne et al., 2005), homothetic (Raković et al., 2012b), and elastic (Raković et al., 2016) tube MPC have been proposed with increasing flexibility in the parameterization of the tube and the associated control policy. One limitation of most tube-based MPC methods (see the introduction of Raković et al. (2012a)) is the offline design of the tube cross-sections and the associated control law, e.g., by prefixing the shape of the cross-section and the feedback gain of the controller (Raković et al., 2012b), which can be conservative.

Compared with tube-based MPC, state feedback MPC (Goulart et al., 2006; Lofberg, 2003) optimizes the state feedback or the equivalent disturbance feedback gains online (Goulart et al., 2006), resulting in reduced conservatism at the cost of increased computational complexity. In fact, the number of decision variables scales linearly in the horizon for tube-based MPC using an offline

☆ Nikolai Matni is supported by NSF awards CPS-2038873, EECS-2231349, SLES-2331880, and NSF CAREER award ECCS-2045834. The material in this paper was partially presented at the 61st IEEE Conference on Decision and Control, Dec. 6–9, 2022, Cancún, Mexico. This paper was recommended for publication in revised form by Associate Editor Marcello Farina under the direction of Editor Ian R. Petersen. Our codes are publicly available at https://github.com/ShaoruChen/Polytopic-SLSMPC.

* Corresponding author.
E-mail addresses: srchen@seas.upenn.edu (S. Chen), preciado@seas.upenn.edu (V.M. Preciado), morari@seas.upenn.edu (M. Morari), nmatni@seas.upenn.edu (N. Matni).

https://doi.org/10.1016/j.automatica.2023.111431
0005-1098/© 2023 Published by Elsevier Ltd.

design, and quadratically in the horizon for state/disturbance feedback MPC. A conservatism improvement together with the quadratic complexity is also achieved by the parameterized tube MPC (Raković et al., 2012a) which adopts online optimization over both the tube and the control policies. Sieber et al. (2021) propose to search over state feedback gains online through the framework of System Level Synthesis (SLS) (Anderson et al., 2019) with structural restrictions such that the number of variables is linear in the horizon.

Robust MPC with model uncertainty: When model uncertainty is present, robust MPC becomes significantly more challenging since the perturbation to the nominal predicted trajectory directly depends on the states and control inputs and hence on both the uncertainty and the controller parameters.

For polytopic model uncertainty, robust MPC methods based on linear matrix inequalities (LMI) are presented in Kothare et al. (1996), Kouvaritakis et al. (2000), Muñoz-Carpintero et al. (2015) and Schuurmans and Rossiter (2000) where affine state feedback policies for a min–max cost function over an infinite horizon are considered. The LMI-based methods do not require much offline parameter design, but their computational complexity tends to scale badly with the system dimension because of the use of semidefinite programming (Gesser et al., 2018). In addition, the ellipsoidal inner approximation of polyhedral state and control input constraints (Kothare et al., 1996) can be conservative.

The application of tube-based MPC to handle model uncertainty can be found in Fleming et al. (2014), Köhler et al. (2019), Langson et al. (2004) and Lu and Cannon (2019). Langson et al. (2004) utilize a homothetic tube to bound the system trajectories and parameterize a feedback control policy associated with each vertex of the tube cross-section. Fleming et al. (2014), Köhler et al. (2019) and Lu and Cannon (2019) apply an offline-designed tube and a control policy $u_t = Kx_t + v_t$ with a pre-stabilizing feedback gain K to guarantee robust constraint satisfaction. Similar to tube MPC in the disturbance-only scenario, the offline design of either the tube or the associated control policy can be conservative.

Recent works by Bujarbaruah et al. (2021, 2022) use LTV state feedback controllers for robust MPC under model uncertainty. Because the LTV state feedback controller has more degrees of freedom, the analysis of the effects of uncertainty is more complex and guaranteeing robust constraint satisfaction is more challenging. Therefore, a simplified, conservative uncertainty over-approximation is needed to solve the robust OCP in a numerically efficient manner. Bujarbaruah et al. (2021) adopt a simple approach that globally over-approximates the effects of model uncertainty by additive disturbances with fixed norm bounds and reduces the robust OCP to the disturbance-only case discussed above. Although easily implementable, this method suffers from the inherent conservatism induced by the global uncertainty over-approximation which fails to utilize the structure of the model uncertainty. In Bujarbaruah et al. (2022), the effects of model uncertainty are first bounded offline as a function of the control inputs, and an LTV state feedback controller is optimized online afterward. We denote this class of methods as *uncertainty over-approximation-based* MPC to highlight the main challenge of bounding the uncertainty when an LTV state feedback controller is considered.

1.1. Contributions

In this work, we propose a robust MPC method which optimizes over LTV state feedback controllers online for systems subject to polytopic model uncertainty and additive disturbances. Inspired by Bujarbaruah et al. (2021), we describe the uncertain system dynamics as the sum of the nominal dynamics and the lumped uncertainty that captures the deviation from the nominal predicted state at each time instant. Unlike (Bujarbaruah et al., 2021) who simply over-approximate the lumped uncertainty uniformly by a sufficiently large additive disturbance signal, our method exactly characterizes how the lumped uncertainty depends on the model uncertainty and controller parameters by leveraging the framework of System Level Synthesis, which establishes the equivalence between the optimization of LTV state feedback controllers and closed-loop system responses. To guarantee robust constraint satisfaction, we propose a novel method that employs a parameterized filtered additive disturbance signal to over-approximate the lumped uncertainty at each time step. Importantly, our method retains the dependence of the uncertainty over-approximation on the controller parameters and thus allows joint optimization of the uncertainty over-approximation parameters and the controller in a convex manner.

We denote the proposed method as *SLS MPC* since our robust OCP formulation is derived in the space of system responses made possible by SLS. By construction, SLS MPC considers a more flexible controller class than the tube-based MPC methods (Fleming et al., 2014; Köhler et al., 2019; Lu & Cannon, 2019) which use a fixed pre-stabilizing controller. Compared with methods (Bujarbaruah et al., 2021, 2022) that also apply LTV state feedback controllers, SLS MPC handles the uncertainty over-approximation problem using a novel approach, the effectiveness of which is demonstrated by numerical examples. Our contributions are summarized as follows.

(1) We propose a novel robust MPC method, SLS MPC, for uncertain LTI systems with polytopic model uncertainty and additive disturbances. The method solves the robust OCP by jointly optimizing LTV state feedback controllers and uncertainty over-approximation parameters through a convex quadratic program. In addition to the typical nominal cost (Bujarbaruah et al., 2022; Fleming et al., 2014; Köhler et al., 2019; Langson et al., 2004), our method allows minimizing upper bounds over the worst-case costs with respect to the model uncertainty in robust MPC.
(2) We provide a comprehensive comparison of tube-based (Köhler et al., 2019; Langson et al., 2004; Lorenzen et al., 2019; Lu & Cannon, 2019) and uncertainty over-approximation-based (Bujarbaruah et al., 2021, 2022) robust MPC methods, including SLS MPC, through numerical examples. We demonstrate that SLS MPC can significantly reduce conservatism compared with all other methods over a wide range of uncertainty parameters.

The rest of the paper is organized as follows. The uncertain LTI system and the problem formulation are introduced in Section 2. In Section 3, we introduce System Level Synthesis and use it to characterize the effects of model uncertainty, which are over-approximated in Section 4 by a filtered disturbance signal. Section 5 presents the proposed SLS MPC method, derives upper bounds on the worst-case costs, and discusses the closed-loop properties of SLS MPC. Extensive simulation is given in Section 6 to compare SLS MPC with representative baseline methods. Section 7 concludes the paper. A preliminary version of this work, which was specialized to time-delay systems, was published in Chen et al. (2022).

Notation: For a dynamical system, we denote the system state at time k by $x(k)$ and the t-step prediction of the state in an MPC loop by x_t. For two vectors x and y, $x \leq y$ denotes element-wise comparison. For a symmetric matrix Q, $Q \succeq 0$ denotes that Q is positive semidefinite. The notation $x_{i:j}$ is shorthand for the set $\{x_i, x_{i+1}, \dots, x_j\}$. For a vector $d \in \mathbb{R}^n$, $S = \text{diag}(d)$ denotes a $n \times n$ dimensional diagonal matrix with d being the diagonal vector. For a sequence of matrices $S_1, \dots, S_N$, $S = \text{blkdiag}(S_1, \dots, S_N)$ denotes that S is a block diagonal matrix whose diagonal blocks

are $S_1, \ldots, S_N$ arranged in the order. We represent a linear, causal operator $\mathbf{R}$ defined over a horizon T by the block-lower-triangular matrix

$$\mathbf{R} = \begin{bmatrix} R^{0,0} & & & \\ R^{1,1} & R^{1,0} & & \\ \vdots & \ddots & \ddots & \\ R^{T,T} & \cdots & R^{T,1} & R^{T,0} \end{bmatrix} \tag{1}$$

where $R^{i,j} \in \mathbb{R}^{p\times q}$ is a matrix of compatible dimensions. The set of such matrices is denoted by $\mathcal{L}_{TV}^{T,p\times q}$ and we will drop the superscript T or $p \times q$ when it is clear from the context. Let $\mathbf{R}(i,:)$ denote the ith block row of $\mathbf{R}$, and $\mathbf{R}(:,j)$ denote the jth block column of $\mathbf{R}$, both indexing from 0 .[8] The support function of a non-empty, closed convex set $\mathcal{X}$ is defined as $h_{\mathcal{X}}(c) = \max\{c^\top x \text{ s.t. } x \in \mathcal{X}\}$.

2. Problem formulation

Consider a discrete-time linear system with uncertain dynamics

$$x(k+1) = (\hat{A} + \Delta_A)x(k) + (\hat{B} + \Delta_B)u(k) + w(k) \tag{2}$$

for $k \geq 0$ where $x(k) \in \mathbb{R}^{n_x}$ is the state, $u(k) \in \mathbb{R}^{n_u}$ is the control input, $w(k) \in \mathbb{R}^{n_x}$ is the additive disturbance, and $(\hat{A}, \hat{B})$ are the known nominal dynamics. The matrices (Δ_A, Δ_B) denote additive model uncertainty and belong to a polytopic set comprising M vertices

$$(\Delta_A, \Delta_B) \in \text{Conv}\{(\Delta_{A,1}, \Delta_{B,1}), \ldots, (\Delta_{A,M}, \Delta_{B,M})\} \tag{3}$$

with Conv denoting the convex hull. For simplicity, we denote the polytopic model uncertainty set defined in (3) as $\mathcal{P}$.[9] The additive disturbance $w(k)$ is assumed to lie in a polytopic set, i.e, $w(k) \in \mathcal{W}$ where $\mathcal{W}$ is a polytope.

In robust MPC, at each time instant k, we solve a finite time robust OCP which aims to synthesize a feedback controller that guarantees the robust satisfaction of state and input constraints. The first control input in the optimal solution is applied to drive the system to the next state, and this process is repeated. In this paper, we focus on solving the following robust OCP with polytopic model uncertainty and additive disturbances.

Problem 1. Solve the following finite time constrained robust OCP with horizon T:

$$\begin{aligned} J_T^* = \min_{f\pi} \quad & J_T(\pi) \\ \text{s.t.} \quad & x_{t+1} = (\hat{A} + \Delta_A)x_t + (\hat{B} + \Delta_B)u_t + w_t \\ & u_t = \pi_t(x_{0:t}) \\ & x_t \in \mathcal{X}, u_t \in \mathcal{U}, x_T \in \mathcal{X}_T, t = 0, 1, \ldots, T-1 \\ & \forall (\Delta_A, \Delta_B) \in \mathcal{P}, \forall w_t \in \mathcal{W}, t = 0, 1, \ldots, T-1 \\ & x_0 = x(k) \end{aligned} \tag{4}$$

where the search is over causal LTV state feedback control policies $\pi = \pi_{0:T-1}$ and $J_T(\pi)$ denotes the nominal cost function as is typical in the robust MPC literature[10]:

$$\begin{aligned} J_T(\pi) = & \textstyle\sum_{t=0}^{T-1}(\hat{x}_t^\top Q \hat{x}_t + u_t^\top R u_t) + \hat{x}_T^\top Q_T \hat{x}_T \\ \text{s.t. } & \hat{x}_{t+1} = \hat{A}\hat{x}_t + \hat{B}u_t, u_t = \pi_t(\hat{x}_{0:t}) \\ & \hat{x}_0 = x(k), \quad \forall t = 0, 1, \ldots, T-1, \end{aligned} \tag{5}$$

where $\hat{x}_{0:T}$ denotes the nominal trajectory, and $Q \succeq 0, R \succ 0$, $Q_T \succeq 0$ denote the state, input, and terminal weight matrices, respectively. The sets $\mathcal{X}, \mathcal{U}$, and $\mathcal{X}_T$ are the polytopic state, input, and terminal constraints, defined as

$$\mathcal{X} = \{x \in \mathbb{R}^{n_x} \mid F_x x \leq b_x\}, \ \mathcal{U} = \{u \in \mathbb{R}^{n_u} \mid F_u u \leq b_u\},$$
$$\mathcal{X}_T = \{x \in \mathbb{R}^{n_x} \mid F_T x \leq b_T\}.$$

We assume that the sets $\mathcal{X}, \mathcal{U}$, and $\mathcal{X}_T$ are compact and contain the origin in their interior.

As shown by Borrelli et al. (2017, Chapter 15), the robust OCP (4) can be solved exactly by dynamic programming which has a high computational cost. For numerical tractability, a finite-dimensional, parameterized controller π_t has to be considered in (4) and the state and input constraints are tightened to guarantee robust constraint satisfaction. In this work, we consider searching over causal LTV state feedback controllers represented by $u_t = \sum_{i=0}^{t} K^{t,t-i}x_i$ for $t = 0, \ldots, T-1$ with $K^{t,t-i}$ being our design parameters. Our approach to Problem (4) relies on over-approximating the effects of uncertainties in the space of closed-loop system responses as shown in the following sections.

3. Characterization of effects of uncertainty

In the robust OCP, the system dynamics can be decomposed as the sum of nominal dynamics and uncertainty-related terms:

$$\begin{aligned} x_{t+1} &= \hat{A}x_t + \hat{B}u_t + \underbrace{\Delta_A x_t + \Delta_B u_t + w_t}_{\eta_t} \\ &= \hat{A}x_t + \hat{B}u_t + \eta_t \end{aligned} \tag{6}$$

where we define $\eta_t := \Delta_A x_t + \Delta_B u_t + w_t$ as the *lumped uncertainty* which models the perturbation to the nominal dynamics at each time step. To find a feasible solution to (4) with minimal conservatism, a proper characterization of η_t is required. However, this is challenging since the lumped uncertainty depends on both the uncertainty parameters and the feedback controller to be designed. In this section, we show that SLS allows us to *exactly* characterize the dynamics of the lumped uncertainty η_t as a function of both the uncertainty and the controller parameters.

3.1. Finite-horizon system level synthesis

To apply SLS, we first stack all relevant state, control input, and uncertainty variables over horizon T as

$$\begin{aligned} &\mathbf{x} = [x_0^\top \ \cdots \ x_T^\top]^\top, \quad \mathbf{u} = [u_0^\top \ \cdots \ u_T^\top]^\top, \\ &\boldsymbol{\eta} = [x_0^\top \ \eta_0^\top \ \cdots \ \eta_{T-1}^\top]^\top, \quad \mathbf{w} = [x_0^\top \ w_0^\top \ \cdots \ w_{T-1}^\top]^\top. \end{aligned} \tag{7}$$

Note that the initial state x_0 is set as the first component in $\boldsymbol{\eta}$ and $\mathbf{w}$, and x_0 can be interpreted as a known disturbance from the origin in this case. The vectors in (7) can be interpreted as finite horizon signals. The parameterization of the LTV state feedback controller $\mathbf{K} \in \mathcal{L}_{TV}^{T,n_u\times n_x}$ is represented by the block-lower-triangular matrix (1) with entries $K^{t,t-i}$. The controller $\mathbf{K}$ can be interpreted as a time-varying linear operator and the state feedback controller is given by $\mathbf{u} = \mathbf{K}\mathbf{x}$. Similarly, we stack the dynamics matrices and uncertainty matrices as

$$\begin{aligned} &\hat{\mathbf{A}} = \text{blkdiag}(\hat{A}, \ldots, \hat{A}), \quad \hat{\mathbf{B}} = \text{blkdiag}(\hat{B}, \ldots, \hat{B}), \\ &\boldsymbol{\Delta}_A = \text{blkdiag}(\Delta_A, \ldots, \Delta_A), \boldsymbol{\Delta}_B = \text{blkdiag}(\Delta_B, \ldots, \Delta_B), \end{aligned} \tag{8}$$

[8] In this paper, we refer to a block matrix in a block-lower-triangular matrix $\mathbf{R}$ using its superscripts shown in Eq. (1). If we let $\mathbf{R}(i,j)$ denote the block matrix in the ith row and jth column, then we have $\mathbf{R}(i,j) = R^{i,i-j}$ with (i,j) indexing from 0.

[9] While in this work we focus on solving the robust MPC problem, our proposed method can be easily incorporated into the adaptive MPC framework where the uncertainty parameters (Δ_A, Δ_B) are assumed time-invariant and the polytopic model uncertainty set $\mathcal{P}$ is updated online as shown by Köhler et al. (2019) and Lu and Cannon (2019).

[10] The worst-case cost function is considered in Section 5.2.

which all belong to $\mathcal{L}_{TV}^{T}$. With this compact notation, the open-loop dynamics of the system (6) can be written as

$$\mathbf{x} = Z\hat{\mathbf{A}}\mathbf{x} + Z\hat{\mathbf{B}}\mathbf{u} + \boldsymbol{\eta}, \tag{9}$$

and the closed-loop dynamics under $\mathbf{u} = \mathbf{K}\mathbf{x}$ follows as

$$\mathbf{x} = Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K})\mathbf{x} + \boldsymbol{\eta}, \tag{10}$$

where Z is a block-downshift operator with the first block sub-diagonal filled with identity matrices and zeros everywhere else. Note that the lumped uncertainty $\boldsymbol{\eta}$ depends on $\mathbf{K}$ and $\mathbf{x}$ as will be shown in the next subsection.

From (10), the mapping from the lumped uncertainty to the state and control input under the feedback controller is given by

$$\begin{bmatrix}\mathbf{x}\\ \mathbf{u}\end{bmatrix} = \begin{bmatrix}(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\\ \mathbf{K}(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\end{bmatrix}\boldsymbol{\eta}. \tag{11}$$

Because of the block-downshift operator Z, we know that the matrix inverse in (11) exists and has a block-lower-triangular structure. The maps from $\boldsymbol{\eta}$ to $(\mathbf{x}, \mathbf{u})$ in (11) are called *system responses*, and we denote them by $\boldsymbol{\Phi}_x \in \mathcal{L}_{TV}^{T,n_x\times n_x}$, $\boldsymbol{\Phi}_u \in \mathcal{L}_{TV}^{T,n_u\times n_x}$ such that

$$\begin{bmatrix}\mathbf{x}\\ \mathbf{u}\end{bmatrix} = \begin{bmatrix}\boldsymbol{\Phi}_x\\ \boldsymbol{\Phi}_u\end{bmatrix}\boldsymbol{\eta}. \tag{12}$$

The relationship between the LTV state feedback controller $\mathbf{K}$ and the system responses $(\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u)$ can be written explicitly as:

$$\begin{aligned}\boldsymbol{\Phi}_x &= (I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1},\\ \boldsymbol{\Phi}_u &= \mathbf{K}(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}.\end{aligned} \tag{13}$$

The following theorem allows us to transform the design of the feedback controller $\mathbf{K}$ into the design of the system responses $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ without explicitly using the nonlinear map (11).

Theorem 1 (*Anderson et al., 2019, Theorem 2.1*)**.** *Over the horizon $t = 0, 1, \ldots, T$, for the system dynamics (6) with the block-lower-triangular state feedback control law $\mathbf{K} \in \mathcal{L}_{TV}^{T,n_u\times n_x}$ defining the control action as $\mathbf{u} = \mathbf{K}\mathbf{x}$, we have:*

(1) The affine subspace defined by

$$\begin{bmatrix}I - Z\hat{\mathbf{A}} & -Z\hat{\mathbf{B}}\end{bmatrix}\begin{bmatrix}\boldsymbol{\Phi}_x\\ \boldsymbol{\Phi}_u\end{bmatrix} = I,\ \ \boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u \in \mathcal{L}_{TV}^{T} \tag{14}$$

parameterizes all possible system responses (12).

(2) For any block-lower-triangular matrices $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\} \in \mathcal{L}_{TV}^{T}$ satisfying (14), the controller $\mathbf{K} = \boldsymbol{\Phi}_u\boldsymbol{\Phi}_x^{-1} \in \mathcal{L}_{TV}^{T}$ achieves the desired responses (12).

Theorem 1 shows the equivalence between system responses and LTV state feedback controllers through the affine constraint (14). Therefore, we can refer to either $\mathbf{K}$ or $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ as the controller parameters. With constraint (14), an optimization problem originally in $\mathbf{K}$ can be transformed into one on $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$. Such a transformation may result in a convex problem in $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ while the original one is not. It provides a direct description of the effects of the lumped uncertainty $\boldsymbol{\eta}$ on the states and control inputs. More importantly, as will be shown next, this transformation can reveal additional structural properties of the robust OCP in the space of system responses that can be exploited to reduce the conservatism of the solution.

3.2. Dynamics of lumped uncertainty

By the definition of lumped uncertainty (6) and the compact notations from (7), (8), we have

$$\boldsymbol{\eta} = Z\begin{bmatrix}\boldsymbol{\Delta}_A & \boldsymbol{\Delta}_B\end{bmatrix}\begin{bmatrix}\mathbf{x}\\ \mathbf{u}\end{bmatrix} + \mathbf{w}. \tag{15}$$

The system responses (12) allow an explicit characterization of the dynamics of η_t under the controller $\mathbf{K}$ as

$$\boldsymbol{\eta} = Z\begin{bmatrix}\boldsymbol{\Delta}_A & \boldsymbol{\Delta}_B\end{bmatrix}\begin{bmatrix}\boldsymbol{\Phi}_x\\ \boldsymbol{\Phi}_u\end{bmatrix}\boldsymbol{\eta} + \mathbf{w} \tag{16}$$

which can be decomposed into the following set of equations[11]

$$\begin{aligned}\eta_t = {} & \Delta_A(\Phi_x^{t,t}x_0 + \sum_{i=1}^{t}\Phi_x^{t,t-i}\eta_{i-1}) +\\ & \Delta_B(\Phi_u^{t,t}x_0 + \sum_{i=1}^{t}\Phi_u^{t,t-i}\eta_{i-1}) + w_t\end{aligned} \tag{17}$$

for $t = 0, \ldots, T-1$. Since η_t only depends on η_i with $i \le t-1$, i.e., the lumped uncertainty before time t, it follows that the values of η_t are *uniquely* determined by the uncertainty parameters $\{\boldsymbol{\Delta}_A, \boldsymbol{\Delta}_B, \mathbf{w}\}$ and the closed-loop system responses $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$. Therefore, we can treat $\boldsymbol{\eta}$ as a function of the uncertainty and controller parameters.

When (Δ_A, Δ_B) and w_t are unknown, the values of η_t become uncertain. Under the uncertainty assumptions in Section 2, we denote by $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ the set of all possible values of $\boldsymbol{\eta}$ under a given controller $\mathbf{K} = \boldsymbol{\Phi}_u\boldsymbol{\Phi}_x^{-1}$:

$$\begin{aligned}\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}) := \{\boldsymbol{\eta} \mid {} & \exists w_t \in \mathcal{W}, 0 \le t \le T-1,\\ & \text{and } (\Delta_A, \Delta_B) \in \mathcal{P},\ \text{s.t. (17) holds}\}.\end{aligned} \tag{18}$$

We call $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ the reachable set of $\boldsymbol{\eta}$ under the controller $\mathbf{K} = \boldsymbol{\Phi}_u\boldsymbol{\Phi}_x^{-1}$. The second argument $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ highlights the dependence of the reachable set on the controller parameters. Despite being exact, Eq. (16) or the derived reachable set $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ is too complex to use for directly solving the robust OCP (4). To address this issue, in the next section, we over-approximate $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ by an uncertainty set with a simpler structure while retaining the dependence on the system responses $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$.

4. Uncertainty over-approximation

To over-approximate $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$, we design a disturbance signal represented by $\boldsymbol{\Sigma}\widetilde{\mathbf{w}}$ where $\boldsymbol{\Sigma} \in \mathcal{L}_{TV}^{T,n_x\times n_x}$ is a filter operating on a normalized unit norm-bounded virtual disturbance signal

$$\widetilde{\mathbf{w}} = [x_0^\top\ \tilde{w}_0^\top\ \cdots\ \tilde{w}_{T-1}^\top]^\top \text{ with } \|\tilde{w}_t\|_\infty \le 1. \tag{19}$$

We denote the set of $\widetilde{\mathbf{w}}$ satisfying the unit norm bound constraint (19) as $\mathcal{W}_{\widetilde{\mathbf{w}}}$, and define the reachable set of the filtered disturbance $\boldsymbol{\Sigma}\widetilde{\mathbf{w}}$ as

$$\mathcal{R}(\boldsymbol{\Sigma}\widetilde{\mathbf{w}}) := \{\boldsymbol{\Sigma}\widetilde{\mathbf{w}} \mid \widetilde{\mathbf{w}} \in \mathcal{W}_{\widetilde{\mathbf{w}}}\} \tag{20}$$

When parameterizing $\boldsymbol{\Sigma}$, we require $\Sigma^{0,0} = I$ such that the first component of $\boldsymbol{\Sigma}\widetilde{\mathbf{w}}$ is x_0. Our goal is to over-approximate the reachable set of the lumped uncertainty $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ by that of the filtered disturbances $\mathcal{R}(\boldsymbol{\Sigma}\widetilde{\mathbf{w}})$, i.e., $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}) \subseteq \mathcal{R}(\boldsymbol{\Sigma}\widetilde{\mathbf{w}})$, such that it suffices to consider the dynamical system

$$\mathbf{x} = Z\hat{\mathbf{A}}\mathbf{x} + Z\hat{\mathbf{B}}\mathbf{u} + \boldsymbol{\Sigma}\widetilde{\mathbf{w}} \tag{21}$$

with the surrogate disturbances $\boldsymbol{\Sigma}\widetilde{\mathbf{w}}$ for solving the robust OCP. The unit norm-bounded assumption on the virtual disturbances $\tilde{w}_t$ simplifies the constraint tightening of the robust OCP (4), while the filter $\boldsymbol{\Sigma}$, which is our design parameter, controls the complexity of $\mathcal{R}(\boldsymbol{\Sigma}\widetilde{\mathbf{w}})$ such that it can over-approximate $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\})$ with minimal conservatism. Next, we discuss the parameterization of the filter $\boldsymbol{\Sigma}$ and formulate a set of linear constraints on $\boldsymbol{\Sigma}$ and $\{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}$ such that $\mathcal{R}(\boldsymbol{\eta}; \{\boldsymbol{\Phi}_x, \boldsymbol{\Phi}_u\}) \subseteq \mathcal{R}(\boldsymbol{\Sigma}\widetilde{\mathbf{w}})$ holds.

[11] We adopt the convention that when $t = 0$, the summation terms in (17) vanish.

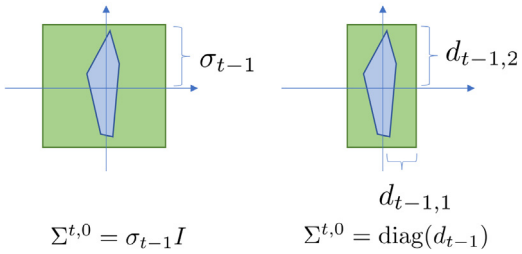

Fig. 1. Blue: reachable set of the lumped uncertainty η_t. Green: uncertainty over-approximation by the filtered disturbance signal $\sigma_t \tilde{w}_t$ with the block diagonal parameterization of the filter $\mathbf{\Sigma}$. The norm ball (left) and hyperrectangle (right) over-approximations are demonstrated.

4.1. Parameterization of the filter

To motivate the use of the filtered disturbance signal $\mathbf{\Sigma}\widetilde{\mathbf{w}}$ for lumped uncertainty over-approximation, we first consider two special block diagonal parameterizations of $\mathbf{\Sigma}$, i.e., only the matrices $\Sigma^{t,0}$ on the diagonal of $\mathbf{\Sigma} \in \mathcal{L}_{TV}$ are non-zero. Note that we always have $\Sigma^{0,0} = I$. For $t = 1, \dots, T$, the first parameterization is given by $\Sigma^{t,0} = \sigma_{t-1} I$ with $\sigma_{t-1} > 0$, while the second parameterization is given by $\Sigma^{t,0} = \mathrm{diag}(d_{t-1})$ with $d_{t-1} \in \mathbb{R}^{n_x}$ and $d_{t-1} > 0$. In the first case, $\mathbf{\Sigma}\widetilde{\mathbf{w}}$ represents the Cartesian product of ℓ_∞ norm balls with radii $\sigma_{t-1} > 0$; in the second case, it represents the Cartesian product of hyperrectangles whose edge lengths are given by the entries in d_{t-1} times 2. We want to search the filter parameters σ_t or d_t such that at time t, the reachable set of the lumped uncertainty η_t is bounded by the simple geometric sets as shown in Fig. 1. Then, we can tighten the constraints of the robust OCP (4) efficiently using the surrogate uncertain linear dynamics (21) which only has additive disturbances.

As shown in Fig. 1, bounding the reachable set of the lumped uncertainty $\boldsymbol{\eta}$ by the reachable set of the additive disturbances $\mathbf{\Sigma}\widetilde{\mathbf{w}}$ necessarily introduces conservatism (indicated by the gap between the blue and green regions in Fig. 1). To reduce the conservatism in uncertainty over-approximation while maintaining numerical tractability, in this work, we parameterize the filter $\mathbf{\Sigma}$ as follows: the sub-diagonal blocks $\Sigma^{t,t-i}$ are non-zero, and the diagonal blocks are parameterized by $\Sigma^{0,0} = I$, $\Sigma^{t,0} = \mathrm{diag}(d_{t-1})$ for $1 \le t \le T$ where $d_{t-1} \in \mathbb{R}^{n_x}$ and $d_{t-1} > 0$. We refer to this parameterization as the full parameterization of the filter. It has the following features: (a) By construction, $\mathbf{\Sigma}$ is invertible. (b) It contains the aforementioned diagonal parameterization as special cases. (c) It utilizes most of the degrees of freedom offered by the block-lower-triangular parameterization of $\mathbf{\Sigma}$ since all off-diagonal matrices are free variables. A numerical comparison of the full and diagonal parameterizations of the filter is given in Section 6.3.4. With the full parameterization of the filter, we are now ready to find sufficient conditions for $\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}})$.

4.2. Reachable set over-approximation problem

The reachable set over-approximation problem can be stated as finding $\{\mathbf{\Phi}_x, \mathbf{\Phi}_u, \mathbf{\Sigma}\}$ such that

$$\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}}) \tag{22}$$

for all $(\Delta_A, \Delta_B) \in \mathcal{P}$ and $w_t \in \mathcal{W}$. The following lemma gives a rigorous formulation of constraint (22), and its proof is given in Appendix A.1.

Lemma 1. *Let the uncertainty assumption $(\Delta_A, \Delta_B) \in \mathcal{P}$ and $w_t \in \mathcal{W}$ hold for all $0 \le t \le T-1$. For a filter $\mathbf{\Sigma}$ and LTV state feedback controller $\mathbf{K}$ with induced system responses $\{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}$, the following two conditions are equivalent:*

(1) $\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}})$.
(2) The following system of equations

$$Z \begin{bmatrix} \mathbf{\Delta}_A & \mathbf{\Delta}_B \end{bmatrix} \begin{bmatrix} \mathbf{\Phi}_x \\ \mathbf{\Phi}_u \end{bmatrix} \mathbf{\Sigma}\widetilde{\mathbf{w}} + \mathbf{w} = \mathbf{\Sigma}\widetilde{\mathbf{w}} \tag{23}$$

has a solution $\widetilde{\mathbf{w}}^* \in \mathcal{W}_{\widetilde{\mathbf{w}}}$ for all possible realizations of the uncertainty parameters $\{\mathbf{\Delta}_A, \mathbf{\Delta}_B, \mathbf{w}\}$.

With the help of Lemma 1, our goal is to find sufficient conditions on $\{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}$ and $\mathbf{\Sigma}$ such that (23) is robustly feasible. However, the presence of bilinear terms $\mathbf{\Phi}_x\mathbf{\Sigma}$ and $\mathbf{\Phi}_u\mathbf{\Sigma}$ in (23) makes it challenging. To resolve this issue, we apply the change of variables

$$\widetilde{\mathbf{\Phi}}_x = \mathbf{\Phi}_x\mathbf{\Sigma}, \quad \widetilde{\mathbf{\Phi}}_u = \mathbf{\Phi}_u\mathbf{\Sigma} \tag{24}$$

where $\mathbf{\Sigma}$ is invertible and $\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u \in \mathcal{L}_{TV}$. In this case, guaranteeing robust feasibility of (23) is equivalent to finding $(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma})$ such that

$$Z \begin{bmatrix} \mathbf{\Delta}_A & \mathbf{\Delta}_B \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x \\ \widetilde{\mathbf{\Phi}}_u \end{bmatrix} \widetilde{\mathbf{w}} + \mathbf{w} = \mathbf{\Sigma}\widetilde{\mathbf{w}}, \ \widetilde{\mathbf{w}} \in \mathcal{W}_{\widetilde{\mathbf{w}}} \tag{25}$$

is feasible for all $(\Delta_A, \Delta_B) \in \mathcal{P}$ and $w_t \in \mathcal{W}$. It can be easily verified that all achievable $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u\}$ are parameterized by the following affine constraint

$$\begin{bmatrix} I - Z\hat{\mathbf{A}} & -Z\hat{\mathbf{B}} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x \\ \widetilde{\mathbf{\Phi}}_u \end{bmatrix} = \mathbf{\Sigma}, \ \widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u \in \mathcal{L}_{TV}, \tag{26}$$

which directly follows from the affine constraint (14) that parameterizes all achievable system responses $\{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}$ for system (6). Note that constraint (26) is jointly affine in $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$. The invertibility of the filter guarantees the equivalence of searching $\{\mathbf{\Phi}_x, \mathbf{\Phi}_u, \mathbf{\Sigma}\}$ under constraint (14) and searching $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$ under constraint (26) for $\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}})$ to hold.

As shown in the following corollary, $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u\}$ can be interpreted as system responses mapping $\widetilde{\mathbf{w}}$ to $(\mathbf{x}, \mathbf{u})$ for the system $\mathbf{x} = Z\hat{\mathbf{A}}\mathbf{x} + Z\hat{\mathbf{B}}\mathbf{u} + \mathbf{\Sigma}\widetilde{\mathbf{w}}$ in closed-loop with $\mathbf{u} = \mathbf{K}\mathbf{x}$. The proof is given in Appendix A.2.

Corollary 1. *Let $\mathbf{\Sigma} \in \mathcal{L}_{TV}^{T, n_x \times n_x}$ be invertible and $\mathbf{K} \in \mathcal{L}_{TV}^{T, n_u \times n_x}$ be a state feedback controller. Then, for the closed-loop dynamics*

$$\mathbf{x} = Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K})\mathbf{x} + \mathbf{\Sigma}\widetilde{\mathbf{w}} \tag{27}$$

over the horizon $t = 0, \dots, T$, we have

(1) The affine subspace defined by (26) parameterizes all achievable system responses $\mathbf{x} = \widetilde{\mathbf{\Phi}}_x\widetilde{\mathbf{w}}$, $\mathbf{u} = \widetilde{\mathbf{\Phi}}_u\widetilde{\mathbf{w}}$ for system (27).
(2) For any block-lower-triangular matrices $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u\}$ satisfying (26), the controller $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u\widetilde{\mathbf{\Phi}}_x^{-1}$ achieves the desired response.

4.3. Convex over-approximation constraints

We now present sufficient conditions on $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$ such that the robust feasibility of (25) is guaranteed. The block-downshift operator Z makes it possible to decompose and analyze the equality constraints in (25) for $t = 0, \dots, T-1$ sequentially.

4.3.1. Case $t = 0$

For bounding η_0 at $t = 0$, we have the following constraint from (25):

$$\Delta_A x_0 + \Delta_B \tilde{\Phi}_u^{0,0} x_0 + w_0 = \Sigma^{1,1} x_0 + \Sigma^{1,0} \tilde{w}_0 \tag{28}$$

where we have used $\tilde{\Phi}_x^{0,0} = I$ as a result of (26) and $\Sigma^{0,0} = I$. The above constraint is robustly feasible with a solution $\|\tilde{w}_0\|_\infty \leq 1$ if and only if

$$\|\Sigma^{1,0-1}(\Delta_A x_0 + \Delta_B \tilde{\Phi}_u^{0,0} x_0 - \Sigma^{1,1} x_0 + w_0)\|_\infty \leq 1 \tag{29}$$

for all $(\Delta_A, \Delta_B) \in \mathcal{P}$ and $w_0 \in \mathcal{W}$. Constraint (29) is non-convex in $\Sigma^{1,0}$, but with our diagonal matrix parameterization $\Sigma^{1,0} = \text{diag}(d_0)$ (see Section 4.1) it can be rewritten as

$$|e_i^\top(\Delta_A x_0 + \Delta_B \tilde{\Phi}_u^{0,0} x_0 - \Sigma^{1,1} x_0 + w_0)| \leq d_{0,i} \tag{30}$$

for $i = 1, \ldots, n_x$ where $d_{0,i}$ denotes the ith entry of d_0 and e_i is the ith standard basis. We compute offline[12]

$$\sigma_{w,i} = \max\{h_{\mathcal{W}}(e_i), h_{\mathcal{W}}(-e_i)\}, i = 1, \ldots, n_x \tag{31}$$

where $h_{\mathcal{W}}(\cdot)$ is the support function of $\mathcal{W}$. Then, the following constraints

$$\begin{aligned} &|e_i^\top(\Delta_A x_0 + \Delta_B \tilde{\Phi}_u^{0,0} x_0 - \Sigma^{1,1} x_0)| + \sigma_{w,i} \leq d_{0,i}, \\ &\forall (\Delta_A, \Delta_B) \in \text{Vert}(\mathcal{P}), \quad i = 1, \ldots, n_x \end{aligned} \tag{32}$$

guarantee that (30) is robustly feasible, where $\text{Vert}(\mathcal{P}) := \{(\Delta_{A,j}, \Delta_{B,j}), j = 1, \ldots, M\}$ denotes the set of vertices of $\mathcal{P}$. To prove this, we can show that the left-hand-side (LHS) of (32) is an upper bound on the LHS of (30) by using the triangle inequality of the absolute value $|\cdot|$, the definition of $\sigma_{w,i}$ in (31), and the fact that the LHS of (30) is convex in (Δ_A, Δ_B) and the maximum of a convex function over a convex polytope is achieved at the vertices (Boyd et al., 2004).

Constraint (32) is convex in the design parameters $\tilde{\Phi}_u^{0,0}$, $\Sigma^{1,1}$ and d_1. It guarantees that for all possible realizations of $(\Delta_A, \Delta_B, w_0)$ and the generated lumped uncertainty η_0, we can always find $\tilde{w}_0^*$ such that $\eta_0 = \Sigma^{1,1} x_0 + \Sigma^{1,0} \tilde{w}_0^*$ with $\|\tilde{w}_0^*\|_\infty \leq 1$. We fix $\tilde{w}_0^*$ as the $\tilde{w}_0$-component of the solution $\widetilde{\mathbf{w}}^*$ to (25).

4.3.2. Case: $t = 1$

Similarly, to bound η_1, we can write the relevant equality constraints from (25) as

$$\begin{aligned} &\Delta_A(\tilde{\Phi}_x^{1,1} x_0 + \tilde{\Phi}_x^{1,0} \tilde{w}_0^*) + \Delta_B(\tilde{\Phi}_u^{1,1} x_0 + \tilde{\Phi}_u^{1,0} \tilde{w}_0^*) + w_1 \\ &= \Sigma^{2,2} x_0 + \Sigma^{2,1} \tilde{w}_0^* + \Sigma^{2,0} \tilde{w}_1 \end{aligned} \tag{33}$$

where $\tilde{w}_0^*$ is the solution from the previous time step and captures the effects of η_0 on future perturbations η_t for $t \geq 1$. Since $\Sigma^{2,0} = \text{diag}(d_1)$ is a diagonal matrix, following the same steps as in the case $t = 0$ and grouping the terms in (33) by x_0, $\tilde{w}_0^*$, and $\tilde{w}_1$, we conclude if the inequalities

$$\begin{aligned} &|e_i^\top(\Delta_A \tilde{\Phi}_x^{1,1} + \Delta_B \tilde{\Phi}_u^{1,1} - \Sigma^{2,2}) x_0| + \\ &|e_i^\top(\Delta_A \tilde{\Phi}_x^{1,0} + \Delta_B \tilde{\Phi}_u^{1,0} - \Sigma^{2,1}) \tilde{w}_0^*| + |e_i^\top w_1| \leq d_{1,i} \end{aligned} \tag{34}$$

for $1 \leq i \leq n_x$ hold robustly, then the robust feasibility of (33) is guaranteed. However, the exact value of $\tilde{w}_0^*$ is unknown to us since it depends on $(\Delta_A, \Delta_B, w_0)$. To address this issue, we treat $\tilde{w}_0^*$ as uncertainty satisfying $\|\tilde{w}_0^*\|_\infty \leq 1$ and further tighten the constraint (34) as

$$\begin{aligned} &|e_i^\top(\Delta_A \tilde{\Phi}_x^{1,1} + \Delta_B \tilde{\Phi}_u^{1,1} - \Sigma^{2,2}) x_0| + \\ &\|e_i^\top(\Delta_A \tilde{\Phi}_x^{1,0} + \Delta_B \tilde{\Phi}_u^{1,0} - \Sigma^{2,1})\|_1 + \sigma_{w,i} \leq d_{1,i}, \\ &\forall (\Delta_A, \Delta_B) \in \text{Vert}(\mathcal{P}), \quad i = 1, \ldots, n_x \end{aligned} \tag{35}$$

by applying the Hölder's inequality $|a^\top \tilde{w}_0^*| \leq \|a\|_1 \|\tilde{w}_0^*\|_\infty \leq \|a\|_1$ and using the fact that $|e_i^\top w_1| \leq \sigma_{w,i}$.

[12] The parameters $\sigma_{w,i}$ for $i = 1, \ldots, n_x$ only need to be computed once by solving $2n_x$ simple linear programs. Afterward they are used for solving the robust OCP (4) for all time steps.

4.3.3. General case

We repeat this process from $t = 0$ to $t = T - 1$ to obtain a set of convex constraints on $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$:

$$\begin{aligned} &|e_i^\top(\Delta_A \tilde{\Phi}_x^{t,t} + \Delta_B \tilde{\Phi}_u^{t,t} - \Sigma^{t+1,t+1}) x_0| + \sigma_{w,i} + \\ &\sum_{i=1}^{t} \|e_i^\top(\Delta_A \tilde{\Phi}_x^{t,t-i} + \Delta_B \tilde{\Phi}_u^{t,t-i} - \Sigma^{t+1,t+1-i})\|_1 \leq d_{t,i}, \\ &\forall (\Delta_A, \Delta_B) \in \text{Vert}(\mathcal{P}), i = 1, \ldots, n_x, t = 0, \ldots, T-1. \end{aligned} \tag{36}$$

In fact, the constraints in (36) can be translated into an equivalent set of linear constraints on $(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma})$.

By the derivation of (36), any feasible solution $(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma})$ to (36) guarantees the robust feasibility of (25). By Lemma 1, they also guarantee that $\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}})$. In this case, it suffices to consider the uncertain dynamical system (21) with the surrogate additive disturbance $\mathbf{\Sigma}\widetilde{\mathbf{w}}$ for solving the robust OCP (4). The synthesized controller $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{\Phi}}_x^{-1} = \mathbf{\Phi}_u \mathbf{\Phi}_x^{-1}$ guarantees robust constraint satisfaction for the original uncertain dynamical system (6). We summarize all the steps of solving the robust OCP in the next section.

5. Formulation of SLS MPC

In this section, we present our solution to the robust OCP (4) and discuss the closed-loop properties of the proposed robust MPC method.

5.1. Constraint tightening of robust OCP

Recall that under constraints (26) and (36), it suffices to consider the uncertain dynamics $\mathbf{x} = Z\hat{\mathbf{A}}\mathbf{x} + Z\hat{\mathbf{B}}\mathbf{u} + \mathbf{\Sigma}\widetilde{\mathbf{w}}$ with $\|\tilde{w}_t\|_\infty \leq 1$ for guaranteeing robust constraint satisfaction of the synthesized controller $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{\Phi}}_x^{-1}$ since the filtered disturbance signal $\mathbf{\Sigma}\widetilde{\mathbf{w}}$ can realize all possible values of the lumped uncertainty $\boldsymbol{\eta}$.

Let us take the state constraint tightening of $x_t \in \mathcal{X}$ as an example. The state constraint is a polyhedral set $\mathcal{X} = \{x \in \mathbb{R}^{n_x} \mid F_x x \leq b_x\}$. Denote the number of linear constraints in defining $\mathcal{X}$ as $n_{\mathcal{X}}$ and $\text{facet}(\mathcal{X}) = \{(F_x(:, i), b_x(i)) \mid i = 1, \ldots, n_{\mathcal{X}}\}$ as the set of all linear constraint parameters of $\mathcal{X}$. For the dynamical system (21) with filtered disturbance $\mathbf{\Sigma}\widetilde{\mathbf{w}}$, the affine constraint (26) parameterizes all achievable closed-loop system responses $(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u)$ such that $\mathbf{x} = \widetilde{\mathbf{\Phi}}_x \widetilde{\mathbf{w}}$, $\mathbf{u} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{w}}$ under an LTV state feedback controller $\mathbf{u} = \mathbf{K}\mathbf{x}$ (see Corollary 1). Using $x_t = \tilde{\Phi}_x^{t,t} x_0 + \sum_{i=1}^{t} \tilde{\Phi}_x^{t,t-i} \tilde{w}_{i-1}$, the following tightened state constraints

$$\begin{aligned} &f^\top \tilde{\Phi}_x^{t,t} x_0 + \sum_{i=1}^{t} \|f^\top \tilde{\Phi}_x^{t,t-i}\|_1 \leq b, \\ &\forall (f, b) \in \text{facet}(\mathcal{X}), \quad t = 0, \ldots, T-1 \end{aligned} \tag{37}$$

guarantee $f^\top x_t \leq b$ for all $(f, b) \in \text{facet}(\mathcal{X})$ robustly under the controller $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{\Phi}}_x^{-1}$. This follows from a direct application of the Hölder's inequality on $f^\top x_t$ and the fact $\|\tilde{w}_t\|_\infty \leq 1$. Similarly, we can tighten the terminal constraint as

$$f^\top \tilde{\Phi}_x^{T,T} x_0 + \sum_{i=1}^{T} \|f^\top \tilde{\Phi}_x^{T,T-i}\|_1 \leq b, \ \forall (f, b) \in \text{facet}(\mathcal{X}_T), \tag{38}$$

and tighten the control input constraints as

$$\begin{aligned} &f^\top \tilde{\Phi}_u^{t,t} x_0 + \sum_{i=1}^{t} \|f^\top \tilde{\Phi}_u^{t,t-i}\|_1 \leq b, \\ &\forall (f, b) \in \text{facet}(\mathcal{U}), \quad t = 0, \ldots, T-1. \end{aligned} \tag{39}$$

Again, constraints (37), (38), (39) can be easily expressed as linear constraints on $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u\}$.

5.2. Convex tightening of the robust OCP

Summarizing all the aforementioned steps, we can synthesize a robust LTV state feedback controller $\mathbf{K}$ as shown in the following theorem.

Theorem 2. *Consider the convex quadratic program*

$$\begin{aligned} \tilde{J}_T^* = \min_{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}} \quad & \tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u) \\ \text{s.t.} \quad & \text{affine constraint (26)} \\ & \text{over-approximation constraint (36)} \\ & \text{tightened constraints (37), (38), (39)} \\ & x_0 = x(k) \end{aligned} \tag{40}$$

where $\widetilde{\mathbf{\Phi}}_x \in \mathcal{L}_{TV}^{T, n_x \times n_x}$, $\widetilde{\mathbf{\Phi}}_u \in \mathcal{L}_{TV}^{T, n_u \times n_x}$, $\mathbf{\Sigma} \in \mathcal{L}_{TV}^{T, n_x \times n_x}$ *is parameterized in Section 4.1, and the cost function is given by*

$$\tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u) = \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x(:, 0) \\ \widetilde{\mathbf{\Phi}}_u(:, 0) \end{bmatrix} x_0 \right\|_2^2 \tag{41}$$

with $\mathbf{Q} = \mathit{blkdiag}(Q, \ldots, Q, Q_T)$, $\mathbf{R} = \mathit{blkdiag}(R, \ldots, R, 0)$. *We have:*

1. *For any feasible solution* $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$ *of Problem (40), the LTV state feedback controller* $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{\Phi}}_x^{-1}$ *is feasible for the robust OCP (4).*
2. *With the additional structural constraint* $\Sigma^{t,t} = 0, t = 1, \ldots, T$ *on the filter* $\mathbf{\Sigma}$, *the optimal cost in (40) is an upper bound on the optimal cost of the robust OCP (4).*

Proof. The proof of the first point follows directly from our derivation of the constraints in (40) from the previous sections. To prove the second point, denote J_T^* the optimal cost of Problem (4) and $\tilde{J}_T^*$ the optimal cost of Problem (40). With the constraint $\Sigma^{t,t} = 0$, $t = 1, \ldots, T$, the nominal dynamics of the surrogate system (21), i.e., when $\tilde{w}_t = 0$ for $t = 0, \ldots, T-1$, equals that of the actual uncertain system (6). Since the set of feasible solutions $\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u$ to Problem (40) only constitutes a subset of all robust LTV state feedback controllers for the robust OCP (4), we have $\tilde{J}_T^* \geq J_T^*$.

Note that since each entry of d_t is lower bounded by $\sigma_w > 0$ from (36), any feasible $\mathbf{\Sigma}$ of Problem (40) is invertible. Theorem 2 shows that by solving the quadratic program (40), we can search for robust LTV state feedback controllers while minimizing an upper bound on J_T^*. Using the surrogate dynamics (21), we can further extend our framework to minimize an upper bound on the worst-case costs of the robust OCP (4).

Corollary 2. *Consider the worst-case cost in the robust OCP (4) :*

$$J_{T,w}^* = \min_{\pi_{0:T-1}} \max_{\substack{(\Delta_A, \Delta_B) \in \mathcal{P} \\ w_t \in \mathcal{W}, 0 \leq t \leq T-1}} \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} \right\|_\infty. \tag{42}$$

Let $\tilde{J}_{T,w}^*$ *be the optimal cost of Problem (40) with*

$$\begin{aligned} \tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u) = & \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x(:, 0) \\ \widetilde{\mathbf{\Phi}}_u(:, 0) \end{bmatrix} x_0 \right\|_\infty + \\ & \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x(:, 1:T) \\ \widetilde{\mathbf{\Phi}}_u(:, 1:T) \end{bmatrix} \right\|_\infty. \end{aligned} \tag{43}$$

Then, we have $\tilde{J}_{T,w}^* \geq J_{T,w}^*$.

Proof. Through the triangle inequality and the multiplicativity of the ℓ_∞ norm, we have that $\tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u)$ defined in (43) is an upper bound on

$$\max_{\substack{\|\tilde{w}_t\|_\infty \leq 1, \\ 0 \leq t \leq T-1}} \left\| \begin{bmatrix} \mathbf{Q}^{1/2} & \\ & \mathbf{R}^{1/2} \end{bmatrix} \begin{bmatrix} \widetilde{\mathbf{\Phi}}_x \\ \widetilde{\mathbf{\Phi}}_u \end{bmatrix} \tilde{\mathbf{w}} \right\|_\infty, \tag{44}$$

which is the worst-case cost for the surrogate system (21) under the controller $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u \widetilde{\mathbf{\Phi}}_x^{-1}$. For any feasible solution $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$ to Problem (40), the filtered disturbance $\mathbf{\Sigma}\tilde{\mathbf{w}}$ in the surrogate dynamics (21) over-approximates the lumped uncertainty in the actual uncertain system (6). Therefore, Eq. (44) is an upper bound on $J_{T,w}^*$ and we have $\tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u) \geq J_{T,w}^*$. Taking the minimum over all feasible $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u, \mathbf{\Sigma}\}$ gives $\tilde{J}_{T,w}^* \geq J_{T,w}^*$.

Remark 1. When the ℓ_2 norm is used to define the worst-case cost $J_{T,w}^*$ in (42), we can assume the virtual disturbances $\tilde{w}_t$ are bounded in ℓ_2 norm rather than ℓ_∞ norm to simplify the derivation of an upper bound on $J_{T,w}^*$. With $\|\tilde{w}_t\|_2 \leq 1$, the derivation of constraints (36) to (39) follows similarly as in the ℓ_∞ case.[13] Then, define $\tilde{J}_T(\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u)$ as in (43) but with the ℓ_∞ and ℓ_∞-induced norms replaced by their ℓ_2 counterparts, and let $\tilde{J}_{T,w}^*$ denote the corresponding optimal cost of Problem (40). Following the proof of Corollary 2, we have $\tilde{J}_{T,w}^* \geq J_{T,w}^*$.

5.3. Closed-loop properties

SLS MPC solves the OCP (40) at each time instant and applies the first optimal control input to drive the system (2) to the next state. However, as in the uncertainty over-approximation-based methods (Bujarbaruah et al., 2021, 2022), it is challenging to show that the OCP (40) is recursively feasible with a fixed horizon T. To illustrate this, let $\eta_{t|k}$ denote the predicted lumped uncertainty at time t in the robust OCP with the initial state $x(k)$. We note that the constraints in (40) are time-varying, which means the constraints on bounding $\eta_{t|k}$ and $\eta_{t|k+1}$ are different, defying the use of the standard shifting argument (Borrelli et al., 2017, Chapter 12) in the space of system responses to prove recursive feasibility. On the other hand, the use of the conservative uncertainty over-approximation (36) implies that the search space of (40) does not include all robustly feasible LTV state feedback controllers $\mathbf{u} = \mathbf{Kx}$. Therefore, the shifting argument in the controller space cannot be used either.

To provide closed-loop guarantee of constraint satisfaction and Input-to-State Stability (ISS), we can equip SLS MPC with the shrinking horizon strategy shown in Bujarbaruah et al. (2022). In short, at time $k = 0$, we solve Problem (40) with horizon $T = N$ and a robust forward invariant set $\mathcal{X}_T$ as the terminal set. We can use the feasible solution at $k = 0$ to construct a safe backup policy $\mathbf{u} = \mathbf{Kx}$ which guarantees that we can drive the system state into $\mathcal{X}_T$ in at most N steps. Specifically, at time $0 < k < N$, we can apply the controller $u(k) = \sum_{i=0}^{k} K^{k,k-i} x(i)$ where $K^{k,k-i}$ are the corresponding matrices drawn from backup policy $\mathbf{K}$. Then, at time $k = N, N+1, \ldots$, the robust OCP (4) rather than Problem (40) is solved exactly with horizon $T = 1$, i.e., no conservative uncertainty over-approximation is used, by enumerating the vertices of the model uncertainty set. Due to the robust forward invariance of $\mathcal{X}_T$, the robust OCP (4) with horizon $T = 1$ is always feasible and the closed-loop trajectory will remain inside $\mathcal{X}_T$ while observing all the state and input constraints for $k > N$. We refer the readers to Bujarbaruah et al. (2022) for the details of this strategy and the proof of robust constraint satisfaction and ISS of the closed-loop system.

6. Numerical comparison

In this section, we compare the proposed SLS MPC with existing robust MPC methods in terms of conservatism (measured

[13] With $\|\tilde{w}_t\|_2 \leq 1$, the diagonal blocks of the filter $\mathbf{\Sigma}$ now can only be parameterized as $\Sigma^{t,0} = \sigma_{t-1} I$ rather than $\Sigma^{t,0} = \mathrm{diag}(d_{t-1})$. This allows rewriting constraints like (29) as convex ones even when the ℓ_2 norm is used.

Table 1
Tube-based MPC methods.

Method	Controller	Tube	Vertex enumeration of tube
`Tube-A` (Langson et al., 2004)	Barycentric	Homothetic	Yes
`Tube-B` (Lorenzen et al., 2019)	$Kx_t + v_t$	Homothetic	Yes
`Tube-C` (Köhler et al., 2019)	$Kx_t + v_t$	Homothetic	No
`Tube-D` (Lu & Cannon, 2019)	$Kx_t + v_t$	Hyperplane	No

Table 2
Uncertainty over-approximation-based MPC methods.

Method	Features
`Lumped-Disturbance-MPC` (Bujarbaruah et al., 2021)	Over-approximate lumped uncertainty globally by a norm-bounded additive disturbance signal.
`Offline-Tightening-MPC` (Bujarbaruah et al., 2022)	Compute constraint tightening margins offline. Search controller online.
`SLS-MPC` (our method)	Search lumped uncertainty over-approximation and controller jointly in the space of system responses.

by the size of the feasible domain) and computational complexity (measured by solver time) through extensive simulation. Two classes of baselines, i.e., the tube-based (Köhler et al., 2019; Langson et al., 2004; Lorenzen et al., 2019; Lu & Cannon, 2019) and uncertainty over-approximation-based methods (Bujarbaruah et al., 2021, 2022), are considered. The main features of each method are summarized in Tables 1 and 2. Through the numerical examples, we demonstrate that

(1) **Conservatism** (measured by the size of the feasible domain): SLS MPC consistently outperforms all baseline methods and by a large margin in the face of large uncertainty. Under varying uncertainty parameters, the feasible domain of SLS MPC is always more than 90% of the maximal robust control invariant set (the theoretically largest feasible domain for any robust MPC method) in the tested examples while the baseline methods become overly conservative quickly as the uncertainty becomes large.
(2) **Computational complexity**: SLS MPC achieves comparable solver time as the baselines. Specifically, it has the same level of complexity as the uncertainty over-approximation-based methods.

6.1. Interpretation of the results

SLS MPC achieves significant improvement in conservatism thanks to (a) the LTV state feedback controller parameterization $\mathbf{u} = \mathbf{K}\mathbf{x}$, and (b) the novel uncertainty over-approximation method (see Section 4). We note that there is naturally a tension between controller parameterization and constraint tightening in solving the robust OCP (4), since considering a more complex controller class necessarily makes it harder to guarantee robust satisfaction of constraints. The performance of SLS MPC is attributed to balancing both aspects in a desirable manner.

Conservatism: As shown in Table 1, existing tube-based methods are mainly restricted to using a relatively simple controller parameterization $u_t = Kx_t + v_t$ where the time-invariant pre-stabilizing feedback gain K is chosen offline instead of optimized online. In contrast, SLS MPC utilizes the LTV state feedback controller $\mathbf{u} = \mathbf{K}\mathbf{x}$ which contains $u_t = Kx_t + v_t$ as a subclass and searches for $\mathbf{K}$ online. Compared with the uncertainty over-approximation-based methods (Bujarbaruah et al., 2021, 2022) shown in Table 2, SLS MPC adopts a novel constraint tightening approach that simultaneously over-approximates the lumped uncertainty by a virtual filtered additive disturbance signal and searches for the controller.

Computational complexity: The use of LTV state feedback controllers necessarily leads to the quadratic growth of the number of decision variables in the horizon. This holds true for the uncertainty over-approximation-based methods (Bujarbaruah et al., 2021, 2022) and SLS MPC. In contrast, the number of decision variables is often linear in the horizon in tube-based MPC. Therefore, we expect the solver time of the uncertainty over-approximation-based methods, including SLS MPC, to be larger than that of tube-based MPC. Fig. 7 supports our analysis, but it also suggests that SLS MPC may even be preferable to some tube-based methods when the uncertainty is large.

We provide a detailed explanation of both tube-based and uncertainty over-approximation-based baseline methods and their implementation in Appendix A.3. We note that to the best of our knowledge, such a comprehensive numerical comparison and evaluation of existing robust MPC methods have not been done before. We make our codes for implementing SLS MPC and all listed baselines publicly available at https://github.com/ShaoruChen/Polytopic-SLSMPC. All the experiments in this section were implemented in MATLAB R2019b with YALMIP (Lofberg, 2004) and MOSEK (ApS, 2019) on an Intel i7-6700K CPU.

6.2. Test example

We evaluate the conservatism of SLS MPC and other baseline methods on a 2-dimensional system adapted from Bujarbaruah et al. (2022) in Section 6.3.1. The only difference from the example in Bujarbaruah et al. (2022) is that we use a different model uncertainty assumption, i.e., we only consider uncertainty in one entry of the dynamics matrices A and B, to allow varying the uncertainty levels in a wide range and highlight the differences between the tested methods. In Section 6.3.2, randomly generated systems are used for conservatism comparison. Experiments on the original example in Bujarbaruah et al. (2022) are shown in Appendix A.4.

The system nominal dynamics and problem constraints are given as

$$\begin{aligned} &\hat{A} = \begin{bmatrix} 1 & 0.15 \\ 0.1 & 1 \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} 0.1 \\ 1.1 \end{bmatrix}, \\ &\mathcal{X} = \left\{ x \in \mathbb{R}^2 \middle| \begin{bmatrix} -8 \\ -8 \end{bmatrix} \leq x \leq \begin{bmatrix} 8 \\ 8 \end{bmatrix} \right\}, \\ &\mathcal{U} = \{u \in \mathbb{R} | -4 \leq u \leq 4\}. \end{aligned} \tag{45}$$

We consider the following polytopic model uncertainty:

$$\begin{aligned} &\Delta_A \in \text{Conv}\left\{ \begin{bmatrix} \epsilon_A & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} -\epsilon_A & 0 \\ 0 & 0 \end{bmatrix} \right\}, \\ &\Delta_B \in \text{Conv}\left\{ \begin{bmatrix} 0 \\ \epsilon_B \end{bmatrix}, \begin{bmatrix} 0 \\ -\epsilon_B \end{bmatrix} \right\} \end{aligned} \tag{46}$$

and the norm bounded additive disturbances $\|w_t\|_\infty \leq \sigma_w$ for robust MPC.[14] The uncertainty parameters ϵ_A, ϵ_B and σ_w are to be specified. The cost weights are chosen as $Q = 10I$, $R = 1$, $Q_T = 10I$.

[14] In this case, according to (31), we have $\sigma_{w,i} = \sigma_w$ for $i = 1, \ldots, nx$.

6.3. Conservatism comparison

6.3.1. Feasible domain comparison

We compare the conservatism of our proposed method SLS-MPC using the full filter parameterization and the baseline robust MPC methods for varying values of $(\epsilon_A, \epsilon_B, \sigma_w)$ and the horizon. For each fixed $(\epsilon_A, \epsilon_B, \sigma_w)$, we first apply an iterative algorithm (Grieder et al., 2003) to find the maximal robust control invariant set and use it as the terminal set $\mathcal{X}_T$. By construction, $\mathcal{X}_T$ gives the largest feasible domain for any robust MPC algorithm and can be achieved by the exact yet computationally expensive dynamic programming approach (Borrelli et al., 2017, Chapter 15). We carry out a 10×10 uniform grid search of initial conditions x_0 over $\mathcal{X}_T$ and solve all robust MPC formulations with the sampled x_0 using horizon $T = 3$ and $T = 10$. The coverage of the feasible domain of each MPC method is given by the ratio of sampled feasible initial conditions and a coverage close to 1 indicates minimal conservatism of the method even compared with dynamic programming. In Fig. 2, we plot the coverages of the aforementioned robust MPC methods under varying $(\epsilon_A, \epsilon_B, \sigma_w)$ values for horizon $T = 3$ and $T = 10$, respectively. In Fig. 3, we plot the feasible domain of each robust MPC method under the parameter $\epsilon_A = 0.4, \epsilon_B = 0.1, \sigma_w = 0.1$ and horizon $T = 3$ on a 20×20 grid of initial conditions. The feasible domain is estimated as the convex hull of the feasible initial conditions of each method.

Fig. 2 shows that SLS-MPC consistently outperforms all other methods for a wide range of uncertainty parameters and Fig. 3 provides a visual illustration. In all cases SLS-MPC achieves a coverage greater than 90%. Importantly, from Fig. 2 we observe that the conservatism of SLS-MPC remains almost unaffected as the horizon increases from $T = 3$ to $T = 10$ while there is clear conservatism deterioration for all other methods. For example, the coverage of Offline-Tightening-MPC is similar to those of the tube-based methods with horizon $T = 3$ but drops considerably with horizon $T = 10$ (see Appendix A.3.2.2 for a detailed explanation).

Remark 2. Although we do not provide a recursive feasibility guarantee with a fixed horizon for SLS-MPC, Figs. 2 and 3 suggest that in practice the feasible domain of SLS-MPC with a fixed horizon can be large enough to overcome the effect of the shrinking horizon strategy discussed in Section 5.3.

6.3.2. Randomly generated examples

Following the problem setup in Section 6.2 with the horizon $T = 10$, we fix the uncertainty parameters $\epsilon_A = 0.2, \epsilon_B = 0.1, \sigma_w = 0.2$. Then, we randomly generate nominal dynamics $\hat{A}, \hat{B}$ by sampling their entries uniformly from a bounded interval and scaling the spectral radius of $\hat{A}$ to a random number in $[0.5, 2.5]$. In Fig. 4 we plot the coverages of different robust MPC methods over 130 randomly generated examples, arranged in an ascending order of the coverages of SLS-MPC. We observe that SLS-MPC achieves the best coverage in almost all examples while no method can consistently achieve the second-best coverage. For a large portion of the randomly generated examples, the coverage gap between SLS-MPC and the second-best method is significant.

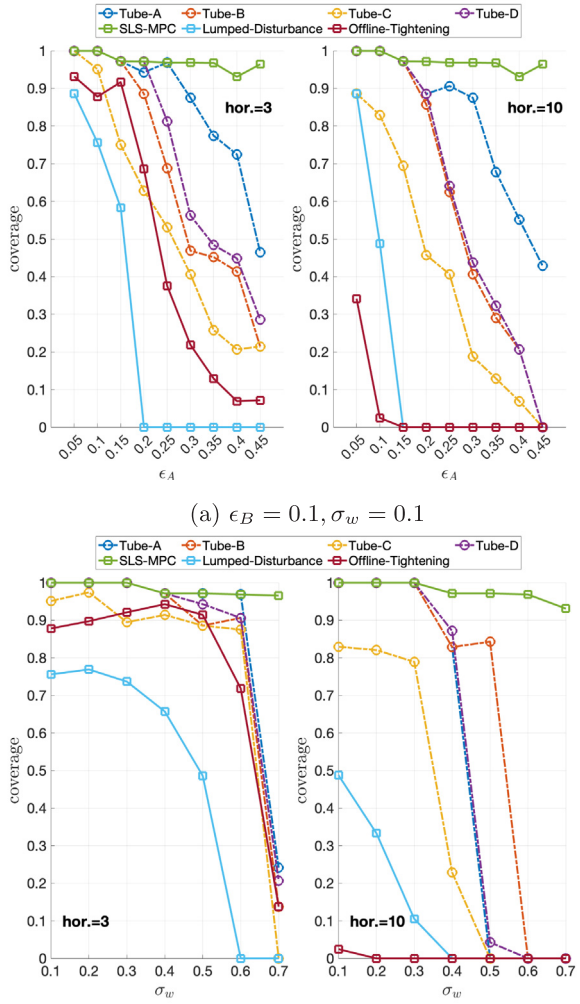

(a) $\epsilon_B = 0.1, \sigma_w = 0.1$

(b) $\epsilon_A = 0.1, \epsilon_B = 0.1$.

Fig. 2. Coverage comparison of different robust MPC methods with different uncertainty parameters and horizons (Left: horizon = 3. Right: horizon = 10). SLS-MPC consistently outperforms all other baselines and achieves coverages over 90%.

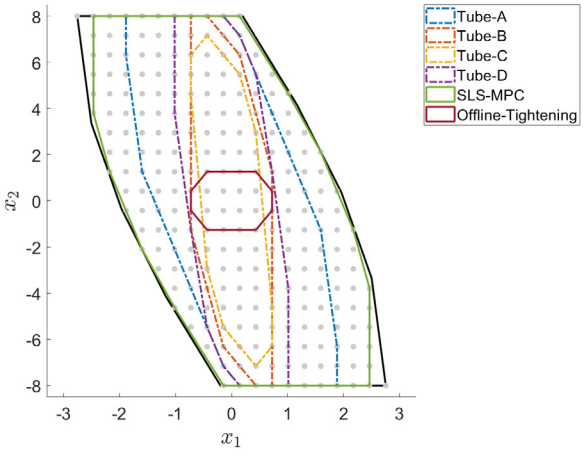

Fig. 3. Feasible domain comparison of robust MPC methods with uncertainty parameters $\epsilon_A = 0.4, \epsilon_B = 0.1, \sigma_w = 0.1$, and horizon $T = 3$. The gray dots denote the initial states sampled from a 20×20 grid.[8]

[8] The feasible domain of Lumped-Disturbance-MPC is empty in this example and therefore is not plotted.

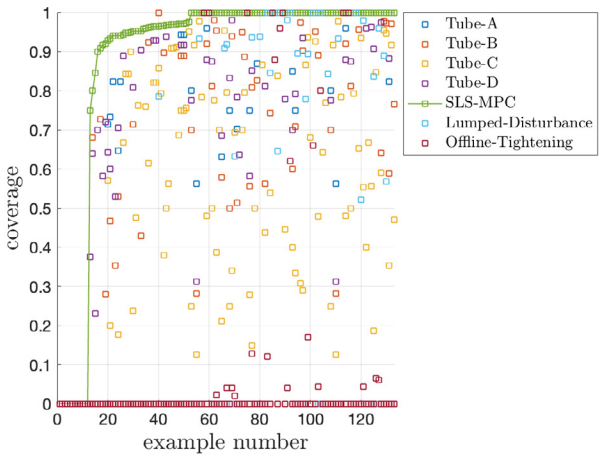

Fig. 4. Coverage comparison of robust MPC methods with horizon $T = 10$ on randomly generated examples.

6.3.3. Effects of tube cross-sections

In our experiments, we note that the design of the tube cross-section can significantly affect the performance of tube-based MPC. Therefore, for each problem setup, we solved tube-based MPC with three common choices of tube cross-sections in the literature: the minimal robust forward invariant set (Kouramas et al., 2005), the maximal robust forward invariant set (Pluymers et al., 2005), the λ-contractive set (Raković, 2016, Chapter 5.6). These methods require first computing a robustly stabilizing controller $u_t = Kx_t$ for the uncertain system $x_{t+1} = (\hat{A} + \Delta_A)x_t + (\hat{B} + \Delta_B)u_t$ with $(\Delta_A, \Delta_B) \in \mathcal{P}$, which can be achieved by solving an SDP (Boyd et al., 1994). We choose such K as the pre-stabilizing feedback gain for tube-based MPC and use it to find the tube cross-section.

Fig. 5 demonstrates the dependence of the conservatism of tube-based robust MPC methods on the choice of the tube cross-section. We observe that (i) the choice of the cross-section can greatly affect the conservatism, and (ii) there is no particular choice that consistently performs better than the others for different levels of uncertainties. Importantly, improperly chosen tube cross-sections may cause tube-based MPC to be overly conservative even when the uncertainty level is small, as demonstrated by the cases of $\epsilon_A = 0.1$ in Fig. 5(d) and $\epsilon_A = 0.25$ in Fig. 5(b). Therefore, in practice it requires manual tuning to figure out which cross-section to use, and the computation of the aforementioned tube cross-sections can be challenging when the system dimension increases. In the feasible domain comparison shown in Figs. 2 and 4, for each tube-based method, we report the *best coverage* out of all three tube cross-sections.

6.3.4. Evaluation of filter parameterization in SLS-MPC

We now investigate how the parameterization of the filter $\mathbf{\Sigma}$ affects the conservatism of `SLS-MPC`. With the full parameterization, the filter $\mathbf{\Sigma}$ is a block-lower triangular matrix while with the diagonal parameterization, all off-diagonal blocks are set zero. Fig. 6 shows that the full parameterization of the filter reduces the conservatism of `SLS-MPC`. When either the model uncertainty or the additive disturbance becomes large, the diagonal parameterization of the filter is not sufficient to mitigate the effects of uncertainties. The average solver time of `SLS-MPC` with the full (diagonal) parameterization is 0.0703(0.0601) seconds in Fig. 6(a) and 0.0764(0.0685) seconds in Fig. 6(b).

6.4. Solver time comparison

In Fig. 7, we plot the average solver time of each robust MPC method in the conservatism comparison experiment shown in Section 6.3.1 with varying uncertainty parameters[9] and horizon $T = 10$. Data is left blank if a robust MPC method was not feasible on any sampled initial conditions. We observe that the solver time of `Tube-C` and `Tube-D` is normally an order of magnitude smaller than that of the other methods. The average solver times of the uncertainty over-approximation-based methods `SLS-MPC`, `Lumped-Disturbance-MPC`, `Offline-Tightening-MPC` are close since these methods all search for an LTV state feedback controller. We observe that the solver time of robust MPC methods can significantly increase when a challenging problem instance (e.g., those with large uncertainty parameters) is encountered.

7. Conclusion

We proposed a novel robust MPC method for uncertain linear systems subject to both polytopic model uncertainty and additive disturbances. Using System Level Synthesis, our method searches for a robust LTV state feedback controller in the space of closed-loop system responses which allows efficient over-approximation of the effects of uncertainty and optimization over upper bounds on the worst-case costs with respect to the model uncertainty. Numerical examples demonstrate that our method can significantly reduce the conservatism compared with a wide range of tube-based and uncertainty over-approximation-based robust MPC methods.

Appendix

A.1. Proof of Lemma 1

As commented in Section 3.2, for any realization of $\{\mathbf{\Delta}_A, \mathbf{\Delta}_B, \mathbf{w}\}$, the values of the lumped uncertainty $\boldsymbol{\eta}$ are given by (16) and uniquely defined. We first assume Eq. (22), i.e., $\mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\}) \subseteq \mathcal{R}(\mathbf{\Sigma}\widetilde{\mathbf{w}})$ holds. This means for any realization of uncertainty and the induced lumped uncertainty values $\boldsymbol{\eta}^\star$, we can always find $\widetilde{\mathbf{w}}^\star \in \mathcal{W}_{\widetilde{\mathbf{w}}}$ such that $\boldsymbol{\eta}^\star = \mathbf{\Sigma}\widetilde{\mathbf{w}}^\star$. Plugging $\widetilde{\mathbf{w}}^\star$ into (16) readily gives that $\widetilde{\mathbf{w}}^\star$ is a solution to (23). Therefore, condition (22) indicates that Eq. (23) holds robustly.

Next, we assume Eq. (23) holds robustly. For any realization of $\{\mathbf{\Delta}_A, \mathbf{\Delta}_B, \mathbf{w}\}$, let $\widetilde{\mathbf{w}}^\star$ denote the solution to (23) and define $\boldsymbol{\xi}^\star = \mathbf{\Sigma}\widetilde{\mathbf{w}}^\star$. Let $\boldsymbol{\eta}^\star$ denote the solution to (16) with the same uncertainty parameters. Due to the uniqueness of the solution, we have $\boldsymbol{\eta}^\star = \boldsymbol{\xi}^\star = \mathbf{\Sigma}\widetilde{\mathbf{w}}^\star$. Therefore, for any $\boldsymbol{\eta}^\star \in \mathcal{R}(\boldsymbol{\eta}; \{\mathbf{\Phi}_x, \mathbf{\Phi}_u\})$, we can find $\widetilde{\mathbf{w}}^\star \in \mathcal{W}_{\widetilde{\mathbf{w}}}$ such that $\boldsymbol{\eta}^\star = \mathbf{\Sigma}\widetilde{\mathbf{w}}^\star$, and hence (22) holds.

A.2. Proof of Corollary 1

1. For a given controller $\mathbf{K}$, similar to (11), the system responses from $\widetilde{\mathbf{w}}$ to $(\mathbf{x}, \mathbf{u})$ are given by $\widetilde{\mathbf{\Phi}}_x = (I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\mathbf{\Sigma}$, $\widetilde{\mathbf{\Phi}}_u = \mathbf{K}(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\mathbf{\Sigma}$ and satisfy the affine constraint (26).

2. First note that the block diagonal of $\widetilde{\mathbf{\Phi}}_x$ and $\mathbf{\Sigma}$ are equal according to (26). Then, $\mathbf{\Sigma}$ being invertible indicates that $\widetilde{\mathbf{\Phi}}_x^{-1}$ exists. For any $\{\widetilde{\mathbf{\Phi}}_x, \widetilde{\mathbf{\Phi}}_u\}$ satisfying (26), we synthesize the state feedback controller as $\mathbf{K} = \widetilde{\mathbf{\Phi}}_u\widetilde{\mathbf{\Phi}}_x^{-1}$ which satisfies $I - Z\hat{\mathbf{A}} - Z\hat{\mathbf{B}}\mathbf{K} = \mathbf{\Sigma}\widetilde{\mathbf{\Phi}}_x^{-1}$ by multiplying Eq. (26) with $\widetilde{\mathbf{\Phi}}_x^{-1}$ from right. Then we have $(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\mathbf{\Sigma} = \widetilde{\mathbf{\Phi}}_x\mathbf{\Sigma}^{-1}\mathbf{\Sigma} = \widetilde{\mathbf{\Phi}}_x$ and $\mathbf{K}(I - Z(\hat{\mathbf{A}} + \hat{\mathbf{B}}\mathbf{K}))^{-1}\mathbf{\Sigma} = \widetilde{\mathbf{\Phi}}_u\widetilde{\mathbf{\Phi}}_x^{-1}\widetilde{\mathbf{\Phi}}_x = \widetilde{\mathbf{\Phi}}_u$.

[9] For tube-based MPC, the cross-section that gives the best coverage is used to generate this figure.

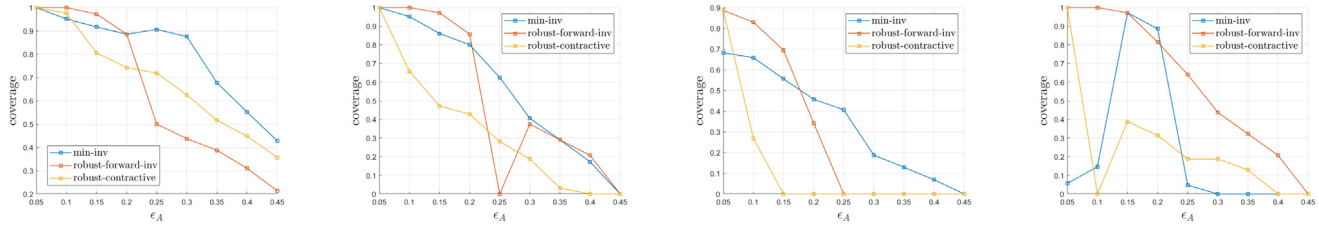

(a) Tube-A (b) Tube-B (c) Tube-C (d) Tube-D

Fig. 5. The coverage of tube-based robust MPC methods with different choices of the tube cross-section. The tube cross-section can be chosen as the minimal nominal forward invariant set (min-inv), maximal robust forward invariant set (robust-foward-inv), or the maximal robust contractive set (robust-contractive). See Section 6.3.3 for details. The coverage is reported for $\epsilon_B = 0.1, \sigma_w = 0.1$ and varying sizes of ϵ_A.

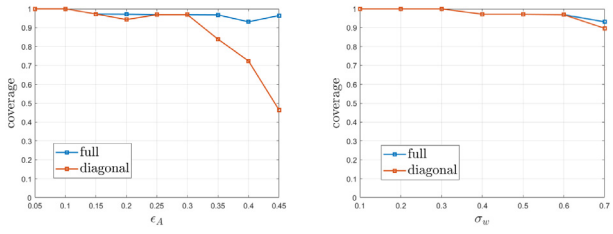

(a) $\epsilon_B = 0.1, \sigma_w = 0.1$ (b) $\epsilon_A = 0.1, \epsilon_B = 0.1$.

Fig. 6. Feasible domain coverage of `SLS-MPC` using the full (blue) and diagonal (orange) parameterization of the filter for different uncertainty parameters.

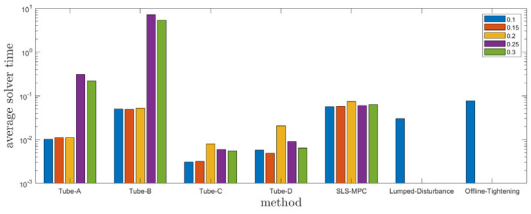

(a) Average solver time of each robust MPC method with horizon $T = 10$, $\epsilon_B = 0.1, \sigma_w = 0.1$ and varying values of ϵ_A as labeled in the legend.

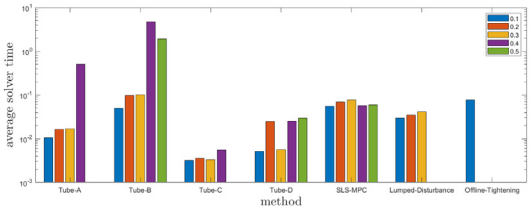

(b) Average solver time of each robust MPC method with horizon $T = 10$, $\epsilon_A = 0.1, \epsilon_B = 0.1$ and varying values of σ_w as labeled in the legend.

Fig. 7. Average solver time of different robust MPC methods with horizon $T = 10$ in the coverage evaluation (see Fig. 2) under different uncertainty parameters.

A.3. Robust MPC baselines

Here we explain the features and implementation of the robust MPC baselines used in Section 6. These methods mainly vary in (i) how to parameterize the control policy π_t and (ii) how to handle the polytopic model uncertainty such that the state and input constraints are robustly satisfied under the synthesized policy π_t.

A.3.1. Tube-based robust MPC

Tube-based methods, which aim to synthesize a robust controller and an associated tube to contain all possible realization of system trajectories, are popular in robust MPC. Tube-based MPC typically follows the following scheme:

(1) Parameterize the control policy π_t.
(2) Parameterize the tube cross-sections $\mathbb{X}_t$.
(3) Enforce the tube containment constraints inductively:

$$x_t \in \mathbb{X}_t \Rightarrow x_{t+1} \in \mathbb{X}_{t+1}, \quad \forall (\Delta_A, \Delta_B) \in \mathcal{P}, \forall w_t \in \mathcal{W}, t = 0, \dots, T-1. \tag{A.1}$$

and guarantee robust satisfaction of all state and input constraints:

$$\mathbb{X}_t \subseteq \mathcal{X}, u_t = \pi_t(x_t) \in \mathcal{U}, \ \forall x_t \in \mathbb{X}_t, t = 0, \dots, T-1, \quad x_0 \in \mathbb{X}_0, \mathbb{X}_T \in \mathcal{X}_T. \tag{A.2}$$

(4) Specify the cost function and solve a robust OCP with constraints (A.1) and (A.2).

Steps 1 to 3 in the above scheme jointly determine the conservatism and computational complexity of a tube-based robust MPC method. Next, we review all four tube-based robust MPC baselines shown in Table 1.

A.3.1.1. Tube-A. The first method which we denote `Tube-A` is from Langson et al. (2004). In `Tube-A`, a homothetic tube (Raković et al., 2012b) is applied:

$$\mathbb{X}_t = z_t + \alpha_t \mathbb{X}, t = 1, \dots, T, \tag{A.3}$$

where $z_t \in \mathbb{R}^{n_x}$ and $\alpha_t \geq 0$ is a scalar. The tube cross-sections $\mathbb{X}_t$ at different time instants are restricted to be transitions and dilations of a given polytope $\mathbb{X}$.

Let n_J denote the number of vertices of $\mathbb{X}$ and $\{x^j\}_{j=1}^{n_J}$ denote the vertices of $\mathbb{X}$. Since $\mathbb{X}_t$ has the same number of vertices as $\mathbb{X}$, we can define $\{x_t^j\}_{j=1}^{n_J}$ as the vertices of $\mathbb{X}_t$ similarly. For controller design, `Tube-A` associates a control input u_t^j with each vertex x_t^j of $\mathbb{X}_t$, and parameterize the feedback policy π_t in a barycentric manner:

$$u_t = \pi_t(x_t) = \sum_{j=1}^{n_J} \lambda_j u_t^j, \quad \text{where } \lambda_j \text{ satisfies } x_t = \sum_{j=1}^{n_J} \lambda_j x_t^j, \lambda_j \geq 0, \sum_{j=1}^{n_J} \lambda_j = 1. \tag{A.4}$$

With the homothetic tube and barycentric controller parameterization, constraints (A.1) and (A.2) can be enforced by only considering the vertex states x_t^j and control inputs u_t^j. In this process, both the vertices of the tube cross-section $\mathbb{X}_t$ and the polytopic model uncertainty set $\mathcal{P}$ are enumerated.

A.3.1.2. Tube-B. In Lorenzen et al. (2019), the homothetic tube parameterization (A.3) is applied with a pre-stabilizing controller

Table A.1
Coverage comparison of robust MPC methods with horizon $T = 3$ and $T = 10$ for the example from Bujarbaruah et al. (2022). The 3-tuple denotes the coverages of tube-based methods using different cross-sections.

Method	Tube-A	Tube-B	Tube-C	Tube-D	SLS-MPC	Lumped-Disturbance	Offline-Tightening
Coverage (hor. = 3)	(0.66, **1**, 1)	(0, **1**, 1)	(0.12, 0.78, **0.88**)	(0.95, **1**, 1)	0.98	0.80	0.85
Coverage (hor .= 10)	(0.34, **1**, 1)	(0, **1**, 0.88)	(0, **0.39**, 0.05)	(0.93, **1**, 0.89)	0.98	0.49	0.02

$\pi_t(x_t) = Kx_t + v_t$. We denote this method as Tube-B where the time-invariant feedback gain K is found offline by solving an SDP such that $u_t = Kx_t$ is robustly stabilizing for the uncertain system (2) with polytopic model uncertainty. The bias terms v_t of the control inputs are optimized online together with the center z_t and scaling parameter α_k of the cross-sections. Constraints (A.1) and (A.2) are enforced using a dual formulation and enumeration of the vertices of each cross-section $\mathbb{X}_t$.

A.3.1.3. Tube-C. Köhler et al. (2019) propose a tube-based robust MPC method that applies a homothetic tube (A.3) and a pre-stabilizing feedback controller $u_t = Kx_t + v_t$. Compared with Tube-B, this method bounds the effects of uncertainty on the dilation of the tube cross-sections offline instead of using vertex enumeration of the tube cross-section. Furthermore, in Köhler et al. (2019) the centers of the cross-sections $\mathbb{X}_t$ are chosen as the nominal trajectory $\hat{x}_t$ evolved according to $\hat{x}_{t+1} = \hat{A}\hat{x}_t + \hat{B}(K\hat{x}_t + v_t)$; in contrast, in Tube-A and Tube-B, the centers z_t of $\mathbb{X}_t$ are treated as free variables. We label the robust MPC method proposed in Köhler et al. (2019) as Tube-C.

A.3.1.4. Tube-D. We denote the robust MPC method proposed in Lu and Cannon (2019) as Tube-D. It uses the pre-stabilizing feedback controller $u_t = Kx_t + v_t$ but parameterizes the cross-sections as $\mathbb{X}_t = \{x \in \mathbb{R}^{n_x} \mid Vx \leq \beta_t\}$, where $V \in \mathbb{R}^{r \times n_x}$ denotes a fixed set of hyperplanes that define a polytope and the hyperplane offsets $\beta_t \in \mathbb{R}^r$ are optimized online. This formulation of $\mathbb{X}_t$ is more flexible than the homothetic tube (A.3), but treating β_t as optimization variables give rise to bilinear terms of the form $\Lambda\beta_t$ in enforcing constraints (A.1) and (A.2), where Λ denotes the dual variables in a Lagrangian function. To avoid numerical intractability, the dual variables Λ are fixed offline by solving a set of linear programs.

A.3.2. Uncertainty over-approximation-based MPC

For the uncertain system $x_{t+1} = \hat{A}x_t + \hat{B}u_t + w_t$ with only additive disturbances, Goulart et al. (2006) prove that the LTV state feedback controller $\mathbf{u} = \mathbf{L}\mathbf{x} + \mathbf{g}$ is equivalent to the disturbance feedback controller $\mathbf{u} = \mathbf{M}\mathbf{w} + \mathbf{v}$.[10] To extend the state/disturbance feedback approach to handle both the polytopic model uncertainty and the additive disturbances, the effects of uncertainty have to be properly over-approximated for the robust OCP to be numerically tractable.

A.3.2.1. Lumped-disturbance-MPC. A simple extension of state/disturbance feedback controllers to handle model uncertainties is to uniformly over-approximate the perturbation $\Delta_A x_t + \Delta_B u_t + w_t$ by a norm-bounded disturbance $\bar{w}_t$, and apply standard state/disturbance feedback controller design on the system $x_{t+1} = \hat{A}x_t + \hat{B}u_t + \bar{w}_t$. Since the state constraint $\mathcal{X}$ and input constraint $\mathcal{U}$ are compact sets, an upper bound on $\|\Delta_A x_t + \Delta_B u_t + w_t\|_\infty$ can be easily obtained using the triangle inequality and sub-multiplicativity of norms as shown in Bujarbaruah et al. (2021). We denote this method proposed in Bujarbaruah et al. (2021) as Lumped-Disturbance-MPC. Despite being simple, the uniform norm-bounded uncertainty over-approximation can be conservative since it is agnostic to the current state of the system. Our proposed method, SLS MPC, resolves this issue by maintaining the dependence of the uncertainty over-approximation on the system states.

A.3.2.2. Offline-tightening-MPC. Bujarbaruah et al. (2022) propose a constraint tightening method for the robust OCP (4) such that the disturbance feedback controller parameters $\mathbf{M}$ and $\mathbf{v}$ explicitly control the tightening margin. However, the constraint tightening parameters also heavily rely on offline computed relaxation bounds that hold for all possible open-loop control inputs, which could already become conservative before the disturbance feedback controller is plugged in. We denote this method from Bujarbaruah et al. (2022) as Offline-Tightening-MPC. In the implementation of this method, the constraint tightening margins are computed offline using a hybrid approach of uncertainty set vertex enumeration and norm bounds over-approximation (see Bujarbaruah et al. (2020, Appendix A.4)). Due to the combinatorial complexity, we only enumerate the vertices of the uncertainty set (46) for a truncation horizon of $N = 3$ and use norm bounds over-approximation for the rest $T - N$ predictive steps. Therefore, the conservatism of such offline computed margins increases as the MPC horizon T becomes large.

A.4. Additional conservatism comparison results

In this section, we compare the coverages of all robust MPC methods for the example from Bujarbaruah et al. (2022) where the model uncertainty is given by

$$\Delta_A \in \mathrm{Conv}\left\{\begin{bmatrix} 0 & \pm\epsilon_A \\ \pm\epsilon_A & 0 \end{bmatrix}\right\}, \quad \Delta_B \in \mathrm{Conv}\left\{\begin{bmatrix} 0 \\ \pm\epsilon_B \end{bmatrix}, \begin{bmatrix} \pm\epsilon_B \\ 0 \end{bmatrix}\right\}, \tag{A.5}$$

with $\epsilon_A = 0.1, \epsilon_B = 0.1$, and the additive disturbances are bounded by $\|w_t\|_\infty \leq 0.1$. This model uncertainty set has 16 vertices in the joint space of (Δ_A, Δ_B), for which the maximal robust control invariant set becomes empty at $\epsilon_A = 0.14$. This prevents us from comparing different methods over a range of uncertainty parameters.

In Table A.1, we report the coverages of the robust MPC methods with horizon $T = 3$ and $T = 10$. For each tube-based MPC method, we report the coverage by (c_1, c_2, c_3) where c_1, c_2, c_3 denote the coverage obtained with the minimal robust forward invariant set, the maximal robust forward invariant set, and the λ-contractive set as the tube cross-sections, respectively (see Section 6.3.3 for details). The best coverage is highlighted in boldface. We observe that the coverage of SLS-MPC is among the best and the conservatism of the tube-based methods depends on the cross-section used. These observations are consistent with what was shown in Section 6.3.

[10] We refer the readers to Goulart et al. (2006) for the exact parameterization of $\mathbf{L}, \mathbf{g}, \mathbf{M}, \mathbf{v}$.

References

Anderson, J., Doyle, J. C., Low, S. H., & Matni, N. (2019). System level synthesis. *Annual Reviews in Control*, *47*, 364–393.

ApS, M. (2019). The MOSEK optimization toolbox for MATLAB manual. Version 9.0. URL http://docs.mosek.com/9.0/toolbox/index.html.

Borrelli, F., Bemporad, A., & Morari, M. (2017). *Predictive control for linear and hybrid systems*. Cambridge University Press.

Boyd, S., Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Boyd, S., El Ghaoui, L., Feron, E., & Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. SIAM.

Bujarbaruah, M., Rosolia, U., Stürz, Y. R., & Borrelli, F. (2021). A simple robust MPC for linear systems with parametric and additive uncertainty. In *2021 American control conference (ACC)* (pp. 2108–2113). IEEE.

Bujarbaruah, M., Rosolia, U., Stürz, Y. R., Zhang, X., & Borrelli, F. (2020). Robust MPC for linear systems with parametric and additive uncertainty: A novel constraint tightening approach. arXiv preprint arXiv:2007.00930.

Bujarbaruah, M., Rosolia, U., Stürz, Y. R., Zhang, X., & Borrelli, F. (2022). Robust MPC for LPV systems via a novel optimization-based constraint tightening. *Automatica*, *143*, Article 110459.

Chen, S., Li, N.-Y., Preciado, V. M., & Matni, N. (2022). Robust model predictive control of time-delay systems through system level synthesis. In *2022 IEEE 61st conference on decision and control (CDC)* (pp. 6902–6909). IEEE.

Fleming, J., Kouvaritakis, B., & Cannon, M. (2014). Robust tube MPC for linear systems with multiplicative uncertainty. *IEEE Transactions on Automatic Control*, *60*(4), 1087–1092.

Gesser, R. S., Lima, D. M., & Normey-Rico, J. E. (2018). Robust model predictive control: Implementation issues with comparative analysis. *IFAC-PapersOnLine*, *51*(25), 478–483.

Goulart, P. J., Kerrigan, E. C., & Maciejowski, J. M. (2006). Optimization over state feedback policies for robust control with constraints. *Automatica*, *42*(4), 523–533.

Grieder, P., Parrilo, P. A., & Morari, M. (2003). Robust receding horizon control-analysis & synthesis. In *2003 42nd IEEE international conference on decision and control, Vol. 1* (pp. 941–946). IEEE.

Köhler, J., Andina, E., Soloperto, R., Müller, M. A., & Allgöwer, F. (2019). Linear robust adaptive model predictive control: Computational complexity and conservatism. In *2019 IEEE 58th conference on decision and control (CDC)* (pp. 1383–1388). IEEE.

Kothare, M. V., Balakrishnan, V., & Morari, M. (1996). Robust constrained model predictive control using linear matrix inequalities. *Automatica*, *32*(10), 1361–1379.

Kouramas, K. I., Rakovic, S. V., Kerrigan, E. C., Allwright, J. C., & Mayne, D. Q. (2005). On the minimal robust positively invariant set for linear difference inclusions. In *Proceedings of the 44th IEEE conference on decision and control* (pp. 2296–2301). IEEE.

Kouvaritakis, B., Rossiter, J. A., & Schuurmans, J. (2000). Efficient robust predictive control. *IEEE Transactions on Automatic Control*, *45*(8), 1545–1549.

Langson, W., Chryssochoos, I., Raković, S., & Mayne, D. Q. (2004). Robust model predictive control using tubes. *Automatica*, *40*(1), 125–133.

Lofberg, J. (2003). Approximations of closed-loop minimax MPC. In *42nd IEEE international conference on decision and control (IEEE cat. no. 03CH37475), Vol. 2* (pp. 1438–1442). IEEE.

Lofberg, J. (2004). YALMIP: A toolbox for modeling and optimization in MATLAB. In *2004 IEEE international conference on robotics and automation* (pp. 284–289). IEEE.

Lorenzen, M., Cannon, M., & Allgöwer, F. (2019). Robust MPC with recursive model update. *Automatica*, *103*, 461–471.

Lu, X., & Cannon, M. (2019). Robust adaptive tube model predictive control. In *2019 American control conference (ACC)* (pp. 3695–3701). IEEE.

Mayne, D. Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, *50*(12), 2967–2986.

Mayne, D. Q., Seron, M. M., & Raković, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, *41*(2), 219–224.

Muñoz-Carpintero, D., Cannon, M., & Kouvaritakis, B. (2015). Robust MPC strategy with optimized polytopic dynamics for linear systems with additive and multiplicative uncertainty. *Systems & Control Letters*, *81*, 34–41.

Pluymers, B., Rossiter, J. A., Suykens, J. A., & De Moor, B. (2005). The efficient computation of polyhedral invariant sets for linear systems with polytopic uncertainty. In *Proceedings of the 2005, American control conference, 2005* (pp. 804–809). IEEE.

Raković, S. V. (2016). Model predictive control: classical, robust, and stochastic [bookshelf]. *IEEE Control Systems Magazine*, *36*(6), 102–105.

Raković, S. V., Kouvaritakis, B., Cannon, M., Panos, C., & Findeisen, R. (2012a). Parameterized tube model predictive control. *IEEE Transactions on Automatic Control*, *57*(11), 2746–2761.

Raković, S. V., Kouvaritakis, B., Findeisen, R., & Cannon, M. (2012b). Homothetic tube model predictive control. *Automatica*, *48*(8), 1631–1638.

Raković, S. V., Levine, W. S., & Açikmese, B. (2016). Elastic tube model predictive control. In *2016 American control conference (ACC)* (pp. 3594–3599). IEEE.

Schuurmans, J., & Rossiter, J. (2000). Robust predictive control using tight sets of predicted states. *IEE Proceedings-Control Theory and Applications*, *147*(1), 13–18.

Sieber, J., Bennani, S., & Zeilinger, M. N. (2021). A system level approach to tube-based model predictive control. *IEEE Control Systems Letters*, *6*, 776–781.

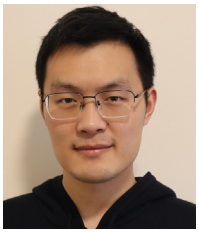

Shaoru Chen received the Ph.D. degree in electrical and systems engineering from the University of Pennsylvania, PA, USA, in 2022, and the B.E. degree in Automation from Zhejiang University, Hangzhou, China, in 2017. He is currently a postdoctoral researcher in Microsoft Research, NYC. His research interests include safe learning-based control, robust model predictive control, and verification of neural networks.

Victor M. Preciado received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA. He is currently an Associate Professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA, USA. He was a recipient of the 2017 National Science Foundation CAREER Award, the 2018 Best Paper Award by the IEEE Control Systems Magazine for his work on control of epidemics, and a runner-up of the 2019 Best Paper Award by the IEEE Transactions on Network Science and Engineering for his work on analysis of epidemics in temporal networks. His main research interests lie at the intersection of Networks, Dynamics, and Data Sciences; in particular, in using innovative mathematical and computational approaches to model and control complex, high-dimensional dynamical systems.

Manfred Morari received the Diploma degree in chemical engineering from ETH Zürich, Zürich, Switzerland, and the Ph.D. degree in chemical engineering from the University of Minnesota, Minneapolis, MN, USA. He was a Professor and the Head of the Department of Information Technology and Electrical Engineering, ETH Zürich. He was the McCollumCorcoran Professor of chemical engineering and the Executive Officer of control and dynamical systems with the California Institute of Technology (Caltech), Pasadena, CA, USA. He was a Professor at the University of Wisconsin, Madison, WI, USA. He is currently with the University of Pennsylvania, Philadelphia, PA, USA. He supervised more than 80 Ph.D. students. Dr. Morari is a fellow of AIChE, IFAC, and the U.K. Royal Academy of Engineering. He is a member of the U.S. National Academy of Engineering. He was a recipient of numerous awards, including Eckman, Ragazzini, and Bellman Awards from the American Automatic Control Council (AACC); Colburn, Professional Progress, and CAST Division Awards from the American Institute of Chemical Engineers (AIChE); Control Systems Technical Field Award and Bode Lecture Prize from IEEE; Nyquist Lectureship and Oldenburger Medal from the American Society of Mechanical Engineers (ASME); and the IFAC High Impact Paper Award. He was the President of the European Control Association. He served on the technical advisory boards of several major corporations.

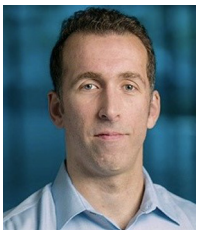

Nikolai Matni is an Assistant Professor in the Department of Electrical and Systems Engineering at the University of Pennsylvania, where he is also a member of the GRASP Lab and PRECISE Center. Nikolai is also a Visiting Research Faculty with Google Robotics, NYC. Prior to joining Penn, Nikolai was a postdoctoral scholar in EECS at UC Berkeley. He has also held a position as a postdoctoral scholar in the Computing and Mathematical Sciences at Caltech. He received his Ph.D. in Control and Dynamical Systems from Caltech in June 2016. He also holds B.A.Sc. and M.A.Sc. in Electrical Engineering from the University of British Columbia, Vancouver, Canada. His research interests broadly encompass the use of learning, optimization, and control in the design and analysis of safety-critical data-driven autonomous systems. Nikolai is a recipient of the NSF CAREER Award (2021), a Google Research Scholar Award (2021), the IEEE ACC 2017 Best Student Paper Award (as co-advisor), and the IEEE CDC 2013 Best Student Paper Award (first ever sole author winner), and the 2021 IEEE George S. Axelby Award.