2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Wasserstein Tube MPC with Exact Uncertainty Propagation

Liviu Aolaritei*, Marta Fochesato*, John Lygeros, and Florian Dörfler

***Abstract*— We study model predictive control (MPC) problems for stochastic LTI systems, where the noise distribution is unknown, compactly supported, and only observable through a limited number of i.i.d. noise samples. Building upon recent results in the literature, which show that distributional uncertainty can be efficiently captured within a Wasserstein ambiguity set, and that such ambiguity sets propagate exactly through the system dynamics, we start by formulating a novel Wasserstein Tube MPC (WT-MPC) problem. We then show that the WT-MPC problem: (1) is a direct generalization of the (deterministic) Robust Tube MPC (RT-MPC) to the stochastic setting; (2) through a scalar parameter, it interpolates between the data-driven formulation based on sample average approximation and the RT-MPC formulation, allowing us to optimally trade between safety and performance; (3) admits a tractable convex reformulation; and (4) is recursively feasible. We conclude with a numerical comparison of WT-MPC and RT-MPC.**

I. Introduction

Tube MPC is an effective control strategy for systems affected by uncertainty that decomposes the system dynamics into two components: (1) a nominal (unperturbed) dynamics, which is utilized for predictions, and (2) an error dynamics, which lies in a tube (e.g., a sequence of polytopes) that contains all possible trajectories of the uncertainty [1].

When the noise has bounded support, and in the absence of additional statistical information on the noise (e.g., samples), robust optimization is employed to formulate a Robust Tube MPC (RT-MPC) problem (see [2], [3] and references therein). However, as the construction results from a worst-case analysis, RT-MPC may often be too conservative. Alternatively, if statistical information about the noise is available, Stochastic Tube MPC (ST-MPC) schemes have been proposed to reduce the conservatism of RT-MPC [4], by relaxing the robust constraints into probabilistic chance constraints with a pre-specified constraint violation probability. While generally intractable, such probabilistic constraints can be dealt with via specific approximations or (often conservative) bounds [5]–[8] if the noise distribution is known, or via scenario-based optimization if the noise distribution is only observable through samples [9]. While alleviating some of the RT-MPC conservatism, the two ST-MPC approaches suffer from a major limitation: by considering a specific noise distribution (e.g., Gaussian), ST-MPC methods fail to guarantee robustness against different (plausible) noise distributions or against *distribution shifts*. To address this shortcoming, more general uncertainty descriptions, which can account for *distributional uncertainty*, i.e., uncertainty about probability distributions, have been recently proposed [10]–[15].

In this paper, we build upon the recent paradigms of *Wasserstein Distributionally Robust Optimization* [16], [17] and *(Optimal Transport-based) Distributional Uncertainty Propagation* [18], [19], and we formulate a novel Tube MPC, which we coin *Wasserstein Tube MPC (WT-MPC)* with Distributionally Robust Conditional Value-at-Risk (DR-CVaR) constraints. Assuming that the noise distribution is compactly supported and only observable through a *limited* number of i.i.d. noise samples, we show that:

- the distributional uncertainty in the state can be expressed as the superposition of a deterministic controlled nominal trajectory and an *autonomous Wasserstein tube* composed of a sequence of Wasserstein ambiguity sets which are *exactly* (in closed-form) propagated through the system dynamics;
- through one scalar parameter, WT-MPC can interpolate between the data-driven formulation based on sample average approximation and RT-MPC, allowing us to optimally trade between safety and performance;
- even in the presence of a small number of samples, WT-MPC can ensure a desired robustness level for the closed-loop system, a smaller closed-loop cost (i.e., increased performance) compared to RT-MPC, and good computational complexity.
- WT-MPC is recursively feasible.

A. Mathematical Preliminaries and Notation

Throughout the paper, $\mathcal{P}(\mathcal{W})$ denotes the space of probability distributions supported on the set $\mathcal{W} \subseteq \mathbb{R}^d$, δ_x denotes the Dirac delta distribution at $x \in \mathbb{R}^d$, and $x \sim \mathbb{P}$ denotes the fact that x is distributed according to $\mathbb{P}$. Moreover, $[i : N]$, with $i \leq N$, denotes the set $\{i, \dots, N\}$, and the symbols $\oplus$ and $\ominus$ denote the Minkowski sum and Pontryagin difference of sets, respectively (see [1, Section 3.1]). Finally, given a matrix A, $A^\dagger$ denotes its Moore-Penrose pseudoinverse.

In this paper, we focus on two classes of transformations of probability distributions: *pushforward* via linear transformations and the *convolution* with a delta distribution.

Definition 1. Let $\mathbb{P} \in \mathcal{P}(\mathcal{W})$ and $A \in \mathbb{R}^{m \times d}$. The pushforward of $\mathbb{P}$ via the linear map $x \mapsto Ax$ is denoted by $A_\# \mathbb{P}$, and is defined by $(A_\# \mathbb{P})(\mathcal{B}) := \mathbb{P}(A^{-1}(\mathcal{B}))$, for all Borel sets $\mathcal{B} \subset A\mathcal{W}$ (where $A\mathcal{W}$ is the image of the set $\mathcal{W}$ through the linear map A).

Intuitively, if $x \sim \mathbb{P}$, then $A_\# \mathbb{P}$ is the probability distribution of the random variable $y = Ax$.

*: Equal contribution.

The authors are with the Automatic Control Laboratory, Department of Electrical Engineering and Information Technology at ETH Zürich, Switzerland {aliviu,mfochesato,jlygeros,dorfler}@ethz.ch

979-8-3503-0123-6/23/$31.00 ©2023 IEEE

Example 1. Let $\widehat{\mathbb{P}} = \frac{1}{n}\sum_{i=1}^{n} \delta_{\widehat{x}^{(i)}}$ be an empirical distribution supported on the samples $\{\widehat{x}^{(i)}\}_{i=1}^{n}$. Then, $A_{\#}\widehat{\mathbb{P}} = \frac{1}{n}\sum_{i=1}^{n} \delta_{A\widehat{x}^{(i)}}$ is empirical as well, supported on the propagated samples $\{A\widehat{x}^{(i)}\}_{i=1}^{n}$.

Moreover, given $x \sim \mathbb{P}$ on $\mathbb{R}^d$ and $y \in \mathbb{R}^d$, $x + y$ is distributed according to the convolution $\delta_y * \mathbb{P}$ defined below.

Definition 2. Let $\mathbb{P} \in \mathcal{P}(\mathbb{R}^d)$ and $y \in \mathbb{R}^d$. Then, the convolution of $\mathbb{P}$ and δ_y is denoted by $\delta_y * \mathbb{P}$, and is defined by $(\delta_y * \mathbb{P})(\mathcal{A}) = \mathbb{P}(\mathcal{A} \ominus y)$, for all Borel sets $\mathcal{A} \subset \mathbb{R}^d$.

II. Wasserstein Tube MPC

We consider the discrete-time linear time-invariant system

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + w_t \\ u_t &= Kx_t + c_t, \end{aligned} \tag{1}$$

where $A \in \mathbb{R}^{d\times d}$, $B \in \mathbb{R}^{d\times m}$ are known, the initial condition $x_0 \in \mathbb{R}^d$ is known and deterministic, and the stochastic noise sequence $\{w_t\}_{t\in\mathbb{N}} \subset \mathbb{R}^d$ is i.i.d. according to an *unknown* distribution $\mathbb{P}$. Moreover, we consider a *fixed stabilizing* feedback gain matrix K, i.e., $A + BK$ is Schur stable.

Assumption 1.

(i) $\mathbb{P}$ has compact support $\mathcal{W} = \{\xi \in \mathbb{R}^d : F\xi \leq g\}$.

(ii) The origin belongs to $\mathcal{W}$, i.e., $0 \in \mathcal{W}$.

(iii) We have access to n_0 i.i.d. samples $\{\widehat{w}^{(i)}\}_{i=1}^{n_0}$ from $\mathbb{P}$.

Since we only have access to a finite number of samples from the unknown noise distribution $\mathbb{P}$, we are faced with *distributional uncertainty*, i.e., uncertainty about probability distributions. In what follows, we employ the methods developed in [18], [19], which lay the foundation to capture and propagate distributional uncertainty in dynamical systems.

A. Capture and Propagate Distributional Uncertainty

We start by defining, for any $t \in \mathbb{N}$, the vector $\mathbf{w}_{[t-1]} = \begin{bmatrix} w_{t-1}^\top & \dots & w_0^\top \end{bmatrix}^\top$. Then, the distributional uncertainty in the state x_t is naturally inherited from the distributional uncertainty in the noise trajectory $\mathbf{w}_{[t-1]}$. Therefore, in order to capture the distributional uncertainty in x_t, we first need to capture the distributional uncertainty in $\mathbf{w}_{[t-1]}$, and then to propagate it, through the system dynamics (1), to x_t.

We start by constructing $n \in \mathbb{N}$ noise sample *trajectories* $\widehat{\mathbf{w}}_{[t-1]}^{(i)} := \begin{bmatrix} (\widehat{w}_{t-1}^{(i)})^\top & \dots & (\widehat{w}_0^{(i)})^\top \end{bmatrix}^\top$, for $i \in [1:n]$, using the n_0 available noise samples. Since the noise is i.i.d., we can easily construct such sample trajectories by letting each entry $\widehat{w}_j^{(i)}$, with $j \in \{0,\dots,t-1\}$ and $i \in \{1,\dots,n\}$, be an arbitrary sample from $\{\widehat{w}^{(i)}\}_{i=1}^{n_0}$. We then define the empirical probability distribution on the product set $\mathcal{W}^t = \mathcal{W} \otimes \dots \otimes \mathcal{W}$, with t terms,

$$\widehat{\mathbb{P}}_{[t-1]} := \frac{1}{n}\sum_{i=1}^{n} \delta_{\widehat{\mathbf{w}}_{[t-1]}^{(i)}}.$$

We capture the distributional uncertainty in $\mathbf{w}_{[t-1]}$ via *Wasserstein ambiguity sets*, i.e., balls of probability distributions, defined using the Wasserstein distance, and centered at the empirical distribution $\widehat{\mathbb{P}}_{[t-1]}$. For $\mathbb{Q} \in \mathcal{P}(\mathcal{W}^t)$, the (type-1) *Wasserstein distance* between $\mathbb{Q}$ and $\widehat{\mathbb{P}}_{[t-1]}$ is defined by

$$W^{\|\cdot\|_2}(\mathbb{Q}, \widehat{\mathbb{P}}_{[t-1]}) := \inf_{\pi\in\Pi} \int_{\mathcal{W}^t\times\mathcal{W}^t} \|x_1 - x_2\|_2 \,\mathrm{d}\pi(x_1, x_2),$$

where $\Pi := \Pi(\mathbb{Q}, \widehat{\mathbb{P}}_{[t-1]})$ is the set of all probability distributions over $\mathcal{W}^t \times \mathcal{W}^t$ with marginals $\mathbb{Q}$ and $\widehat{\mathbb{P}}_{[t-1]}$. The semantics are as follows: we seek the minimum cost to transport the probability distribution $\mathbb{Q}$ onto the probability distribution $\widehat{\mathbb{P}}_{[t-1]}$, when transporting a unit of mass from x_1 to x_2 costs $\|x_1 - x_2\|_2$. Intuitively, $W^{\|\cdot\|_2}(\mathbb{Q}, \widehat{\mathbb{P}}_{[t-1]})$ quantifies the discrepancy between $\mathbb{Q}$ and $\widehat{\mathbb{P}}_{[t-1]}$ and it naturally provides us with a definition of ambiguity in $\mathcal{P}(\mathcal{W}^t)$. Specifically, the *Wasserstein ambiguity set* (henceforth simply referred to as *ambiguity set*) of radius ε, centered at $\widehat{\mathbb{P}}_{[t-1]}$, and with support $\mathcal{W}^t$ is defined by

$$\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[t-1]}) := \{\mathbb{Q} \in \mathcal{P}(\mathcal{W}^t) : W^{\|\cdot\|_2}(\mathbb{Q}, \widehat{\mathbb{P}}_{[t-1]}) \leq \varepsilon\}.$$

In words, $\mathbb{B}_\varepsilon^c(\widehat{\mathbb{P}}_{[t-1]})$ includes all probability distributions on $\mathcal{W}^t$ onto which $\widehat{\mathbb{P}}_{[t-1]}$ can be transported with a budget of at most ε. Such ambiguity sets are shown in [18], [19] to be a very natural and principled tool to capture distributional uncertainty, enjoying powerful geometrical, statistical, and computational features and guarantees. Moreover, they are easily propagated through linear maps, and the result of the propagation is itself an ambiguity set.

Remark 1. The ambiguity radius ε is a tunable parameter that encapsulates the robustness (or risk aversion) level. Higher ε translates to more distributions being captured within $\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[t-1]})$, and consequently more robustness against unforeseen noise realizations being introduced.

We are now ready to study the propagation of the distributional uncertainty from the noise $\mathbf{w}_{[t-1]}$ to the state x_t. With the aim of formulating a Tube MPC (see [1, Chapter 3.2]), we start by rewriting the state as $x_t = z_t + e_t$, i.e., the sum of a deterministic *nominal state* z_t, and a stochastic *error state* e_t. This gives rise to the equivalent system dynamics

$$z_{t+1} = Az_t + Bv_t \tag{2a}$$

$$v_t = Kz_t + c_t \tag{2b}$$

$$e_{t+1} = A_K e_t + w_t, \tag{2c}$$

with $A_K := A + BK$, and with initial conditions $z_0 = x_0$ and $e_0 = 0$. Such state separation will allow us to represent the distributional uncertainty in the state trajectory as the superposition of the deterministic controlled nominal state trajectory z_t and an *autonomous Wasserstein tube*. To do so, we start by rewriting the error dynamics (2c) in the form

$$\begin{aligned} e_t &= \mathbf{D}_{t-1}\mathbf{w}_{[t-1]} \\ \mathbf{D}_{t-1} &:= \begin{bmatrix} I & A_K & \dots & A_K^{t-1} \end{bmatrix}. \end{aligned} \tag{3}$$

Moreover, we denote by $\{\widehat{e}_t^{(i)}\}_{i=1}^{n}$ the n error state samples, obtained by feeding into (3) the n noise sample trajectories $\{\widehat{\mathbf{w}}_{[t-1]}^{(i)}\}_{i=1}^{n}$, i.e.,

$$\widehat{e}_t^{(i)} := \mathbf{D}_{t-1}\widehat{\mathbf{w}}_{[t-1]}^{(i)}, \quad \forall i \in [1:n].$$

The following proposition shows that the distributional uncertainty in x_t can be exactly captured.

Proposition 1. *Let Assumption 1 hold, and consider the linear control system* (1)*, with i.i.d. noise* $\{w_t\}_{t\in\mathbb{N}}$*. Moreover, let* $\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[t-1]})$ *capture the distributional uncertainty in the noise trajectory* $\mathbf{w}_{[t-1]}$*. Then the distributional uncertainty in* x_t *is exactly captured by the ambiguity set*

$$\mathbb{S}_t := \mathbb{B}_\varepsilon^{\|\cdot\|_2\circ\mathbf{D}_{t-1}^\dagger}(\widehat{\mathbb{P}}_{x_t}), \tag{4}$$

centered at $\widehat{\mathbb{P}}_{x_t} := \frac{1}{n}\sum_{i=1}^n \delta_{z_t+\widehat{e}_t^{(i)}}$*, and with support* $z_t \oplus \mathcal{E}_t$*, for* $\mathcal{E}_t := \mathbf{D}_{t-1}\mathcal{W}^t$*.*

The proof of this and all subsequent results are deferred to the online extended version [20].

Remark 2. Expression (4) reveals that the distributional uncertainty in the state trajectory can be represented in the probability space $\mathcal{P}(\mathbb{R}^d)$ as the superposition of the nominal state trajectory δ_{z_t} (the probabilistic representation of z_t) and the *Wasserstein tube* $\mathbb{B}_\varepsilon^{\|\cdot\|_2\circ\mathbf{D}_{t-1}^\dagger}(\widehat{\mathbb{P}}_{e_t})$, with center $\widehat{\mathbb{P}}_{e_t} := \frac{1}{n}\sum_{i=1}^n \delta_{\widehat{e}_t^{(i)}}$ and support $\mathcal{E}_t$. This follows immediately from [19, Corollary 8], which ensures that $\mathbb{S}_t = \delta_{z_t} * \mathbb{B}_\varepsilon^{\|\cdot\|_2\circ\mathbf{D}_{t-1}^\dagger}(\widehat{\mathbb{P}}_{e_t})$. Wasserstein tubes are a natural generalization of the standard robust tubes (used in RT-MPC), as explained next. Let diam$(\mathcal{E}_t)$ denote the diameter of $\mathcal{E}_t$, measured using the distance $\|\cdot\|_2\circ\mathbf{D}_{t-1}^\dagger$. Then, through the choice of ε, ranging from 0 to diam$(\mathcal{E}_t)$, the Wasserstein tube $\mathbb{B}_\varepsilon^{\|\cdot\|_2\circ\mathbf{D}_{t-1}^\dagger}(\widehat{\mathbb{P}}_{e_t})$ interpolates between the empirical distribution $\frac{1}{n}\sum_{i=1}^n \delta_{\widehat{e}_t^{(i)}}$ and the probabilistic representation of the robust tube $\mathcal{E}_t$, i.e., the set of all distributions δ_ξ, $\forall \xi \in \mathcal{E}_t$.

In what follows, we inspect the four components of (4).

- *Ambiguity radius* ε. This quantity is naturally inherited from the ambiguity set $\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[t-1]})$ that models the distributional uncertainty in the noise trajectory.
- *Center* $\widehat{\mathbb{P}}_{x_t}$. This is an empirical distribution over the n points $\{z_t + \widehat{e}_t^{(i)}\}_{i=1}^n$. Notice that the position of these points in $\mathbb{R}^d$ is controlled by the feedforward input c_t.
- *Transportation cost* $\|\cdot\|_2 \circ \mathbf{D}_{t-1}^\dagger$. This is defined as $\left(\|\cdot\|_2 \circ \mathbf{D}_{t-1}^\dagger\right)(\xi) := \|\mathbf{D}_{t-1}^\dagger \xi\|_2$, and it influences the shape of the ambiguity set, as explained next. Using the SVD decomposition $\mathbf{D}_{t-1} = U\Sigma V^\top$, with $\{\sigma_i\}_{i=1}^d$ the singular values of $\mathbf{D}_{t-1}$ and $\{u_i\}_{i=1}^d$ the orthonormal columns of U, the transportation cost boils down to

$$\|\mathbf{D}_{t-1}^\dagger(x_1 - x_2)\|_2 = \sqrt{\sum_{i=1}^d \frac{1}{\sigma_i^2}\left|u_i^\top(x_1 - x_2)\right|^2}. \tag{5}$$

  This shows that the cost of moving probability mass from the center distribution in the direction u_i costs $\|x_1 - x_2\|/\sigma_i$. The feedback gain matrix K has an indirect influence on the amount of mass moved in this direction through the singular value σ_i of the matrix $\mathbf{D}_{t-1}$ (e.g., the higher the value of σ_i, the more probability mass is moved in the direction u_i).
- *Support set* $z_t \oplus \mathcal{E}_t$. Since $\mathcal{W}$ is compact and polyhedral, we have that the set $z_t \oplus \mathcal{E}_t$ is compact and polyhedral, and can be written as

$$z_t \oplus \mathcal{E}_t = \{\xi \in \mathbb{R}^d : F_t\xi \le g_t + F_t z_t\}. \tag{6}$$

  for some $q_t \in \mathbb{N}$, $F_t \in \mathbb{R}^{q_t\times d}$, and $g_t \in \mathbb{R}^{q_t}$. Moreover, F_t and g_t can be obtained from the following iteration: $\mathcal{E}_t = A_K\mathcal{E}_{t-1} \oplus \mathcal{W}$, $\mathcal{E}_0 = \{0\}$.

Remark 3. The inspection of the ambiguity set (4) reveals that the feedforward term c_t can control (through z_t) the *position* in $\mathbb{R}^d$ of the center distribution $\widehat{\mathbb{P}}_{x_t}$. However, c_t has no influence over the *shape and size* of the ambiguity set (i.e., the transportation cost and radius). In particular, these are exclusively influenced by the feedback gain matrix of K (through $\mathbf{D}_{t-1}^\dagger$). For more details on the decomposition of the roles of c_t and K, we refer to [18, Section IV].

B. Distributionally Robust CVaR Constraints

Armed with the closed-form expression $\mathbb{S}_t$ for the ambiguity set that captures the distributional uncertainty in the state x_t, we can now study how to impose constraints. Specifically, we define the polyhedral constraint set $\mathcal{X} := \left\{x \in \mathbb{R}^d : \max_{j\in[1:J]} a_j^\top x + b_j \le 0,\ J \in \mathbb{N}\right\}$, and we want to guarantee that, for some $\gamma \in (0,1)$, the Distributionally Robust Conditional Value-at-Risk (DR-CVaR) constraint

$$\sup_{\mathbb{Q}\in\mathbb{S}_t} \mathbf{CVaR}_{1-\gamma}^{\mathbb{Q}}\left(\max_{j\in[1:J]} a_j^\top x_t + b_j\right) \le 0 \tag{7}$$

is satisfied (see [18, Equation (1)] for the definition of CVaR). Such constraints are very natural for control tasks in the face of distributional uncertainty, where safety is of interest. We motivate this in what follows.

- *Safety considerations.* CVaR constraints are by now standard in risk averse optimization. From [21] we know that (7) implies the following nonconvex distributionally robust chance constraint (DR-CC):

$$\inf_{\mathbb{Q}\in\mathbb{S}_t} \mathbb{Q}(x_t \in \mathcal{X}) \ge 1 - \gamma.$$

  Using Proposition 1, this guarantees that $x_t \in \mathcal{X}$, with probability $1-\gamma$, for all the noise trajectory distributions in $\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[t-1]})$. Additionally, (7) guarantees that $x_t \in \mathcal{X}$ in expectation for the most averse noise realizations of probability γ. This naturally controls the distance of x_t from $\mathcal{X}$ for the remaining probability γ. Differently, notice that if we only impose the DR-CC, then x_t could be arbitrarily far from $\mathcal{X}$ with probability γ.
- *Computational tractability.* Using results from distributionally robust optimization (see [16] and [17]), in Proposition 2 we show that the DR-CVaR constraint (7) can be exactly reformulated as a finite set of deterministic convex constraints.

Proposition 2. *Let Assumption 1 hold. Then, the constraint* (7) *is equivalent to the following set of convex* *constraints, whose feasible region we denote by Γ_t:*

$$\forall i \in [1:n],\ \forall j \in [1:J+1]:$$
$$\begin{cases} \tau \in \mathbb{R}, \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \zeta_{ij} \in \mathbb{R}_+^{q_t} \\ \lambda \varepsilon n + \sum_i^n s_i \leq 0 \\ \alpha_j^\top (z_t + \widehat{e}_t^{(i)}) + \beta_j(\tau) + \zeta_{ij}^\top \left(g_t - F_t \widehat{e}_t^{(i)} \right) \leq s_i \\ \left\| \mathbf{D}_{t-1}^\dagger \left((\mathbf{D}_{t-1}^\dagger)^\top \mathbf{D}_{t-1}^\dagger \right)^{-1} \left(F_t^\top \zeta_{ij} - \alpha_j \right) \right\|_2 \leq \lambda, \end{cases}$$

with $\alpha_j := a_j/\gamma$ and $\beta_j(\tau) := (b_j + \gamma\tau - \tau)/\gamma$, for $j \in [1:J]$, as well as $\alpha_{J+1} := 0$ and $\beta_{J+1}(\tau) := \tau$.

Proposition 2 guarantees that the following equivalence

$$\sup_{\mathbb{Q} \in \mathbb{S}_t} \text{CVaR}_{1-\gamma}^{\mathbb{Q}} \left(\max_{j \in [1:J]} a_j^\top x_t + b_j \right) \leq 0 \iff z_t \in \Gamma_t \quad (8)$$

holds. Interestingly, (8) reveals that the DR-CVaR constraint on the distributionally uncertain state x_t can be equivalently reformulated as a set of deterministic constraints on the nominal state z_t. In the rest of the paper, we will write $z \in \Gamma_t$ to denote the fact that there exist $\tau \in \mathbb{R}, \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \zeta_{ij} \in \mathbb{R}_+^{q_t}$, for $i \in [1:n]$ and $j \in [1:J+1]$, which satisfy the constraints in Proposition 2 for $z_t = z$.

C. MPC with DR-CVaR Constraints

We are now ready to formulate our Wasserstein Tube MPC (WT-MPC). We let $N \in \mathbb{N}$ denote the MPC horizon of interest, and we use the subscript $k|t$, for $k \in [0:N]$, to denote the (open-loop) predicted dynamics at time k given the (closed-loop) time step t. Moreover, as required by engineering applications, we consider the *robust* constraint $u_t \in \mathcal{U}$ on the input, for all $t \in \mathbb{N}$. Then, the WT-MPC reads

$$\begin{aligned} \min\ & \sum_{k=0}^{N-1} \left(\|z_{k|t}\|_Q^2 + \|v_{k|t}\|_R^2 \right) \\ \text{s.t.}\ & c_{k|t}, v_{k|t} \in \mathbb{R}^m, z_{k|t} \in \mathbb{R}^d && \forall k \in [0:N] \\ & z_{k+1|t} = A z_{k|t} + B v_{k|t} && \forall k \in [0:N-1] \\ & v_{k|t} = K z_{k|t} + c_{k|t} && \forall k \in [0:N-1] \\ & v_{k|t} \in \mathcal{U} \ominus K\mathcal{E}_k && \forall k \in [0:N-1] \\ & z_{k|t} \in \mathcal{Z}_k && \forall k \in [1:N-1] \\ & z_{N|t} \in \mathcal{Z}_f \\ & z_{0|t} = x_t, \end{aligned}$$

with

- *Initial condition $z_{0|t} = x_t$.* The open-loop nominal state $z_{0|t}$ is initialized at the (measured) closed-loop state value x_t (leaving $e_{0|t} = 0$). This guarantees that the open-loop dynamics are exactly as in (2).
- *Nominal constraint sets $\mathcal{Z}_k$.* The choice $\mathcal{Z}_k := \Gamma_k$ guarantees that the DR-CVaR constraint (7) is satisfied by $x_{k|t}$ (recall the equivalence (8)). In this case, the constraint $z_{k|t} \in \Gamma_k$ will be enforced through the set of constraints provided in Proposition 2. Consequently, the WT-MPC will inherit the decision variables $\tau \in \mathbb{R}, \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \zeta_{ij} \in \mathbb{R}_+^{q_t}$, $\forall i \in [1:n], \forall j \in [1:J+1]$. However, this choice does not ensure recursive feasibility. In Section III we show that this issue can be resolved through an appropriate constraint tightening, which results in a choice $\mathcal{Z}_k \subsetneq \Gamma_k$.
- *Terminal nominal set $\mathcal{Z}_f$.* In Section III we explain how this set should be chosen to ensure recursive feasiblity.
- *Input constraint sets $\mathcal{U} \ominus K\mathcal{E}_k$.* Since $u_k = K(z_k + e_k) + c_k = v_k + K e_k$, and since e_k is supported on $\mathcal{E}_k$, this choice guarantees that $u_k \in \mathcal{U}$ is satisfied.

Remark 4. The only difference between WT-MPC and RT-MPC stands in the choice of the nominal constraint sets. Recall that in RT-MPC these sets are chosen as $\mathcal{X} \ominus \mathcal{E}_k$ [2]. Since the DR-CVaR constraint (7) relaxes the set $\mathcal{X}$, and since (7) is equivalent to $z_k \in \Gamma_k$, we have that $\mathcal{X} \ominus \mathcal{E}_k \subset \Gamma_k$. Therefore, WT-MPC reduces the conservatism of RT-MPC (at the expense of a pre-defined constraint violation probability). Finally, WT-MPC is a direct generalization of the RT-MPC to the stochastic setting. Indeed, through the choice of ε, the WT-MPC interpolates between the data-driven formulation based on sample average approximation, for $\varepsilon \to 0$, and the RT-MPC formulation, for $\varepsilon \to \text{diam}(\mathcal{E}_N)$. This follows immediately from Remark 2.

III. Recursive Feasibility

Since the WT-MPC problem is solved in a receding horizon fashion, it is fundamental to ensure that it is *recursively feasible*. As pointed out in the previous section, the choice $\mathcal{Z}_k = \Gamma_k$, with Γ_k defined in Proposition 2, does not guarantee recursive feasibility. To see this, recall from (8) that $z_{k|t} \in \Gamma_k$ is equivalent to $x_{k|t}$ satisfying the DR-CVaR constraint (7). The latter, in turn, guarantees that $x_{k|t} \in \mathcal{X}$ only with high probability, i.e., for almost (*but not*) all possible noise realizations. Since WT-MPC is initialized at the closed-loop state x_t, there is a low, but *strictly positive* probability of infeasibility. Notice that this issue does not appear in RT-MPC, which robustifies against all possible noise realizations by imposing the constraint $z_{k|t} \in \mathcal{X} \ominus \mathcal{E}_k$. Such conservative choice, which considers the support of the noise, automatically guarantees recursive feasibility.

In what follows, we will show that an appropriate constraint tightening resolves this issue. We start by defining the nominal constraint sets as

$$\mathcal{Z}_k := \bigcap_{p=1}^{k} \left(\Gamma_p \ominus \left(\bigoplus_{r=p}^{k-1} A_K^r \mathcal{W} \right) \right), \quad \forall k \in [1:N]. \quad (9)$$

By construction, we have that $\mathcal{Z}_k \subset \Gamma_k$. Such constraint tightening is consistent with the stochastic MPC literature. Specifically, (9) is closely related to what is done in [5] to guarantee recursive feasibility.

Remark 5. Using the fact that $\mathcal{X} \ominus \mathcal{E}_k \subset \Gamma_k$, it can be easily derived that $\mathcal{X} \ominus \mathcal{E}_k \subset \mathcal{Z}_k$. This shows that the constraint tightening (9), which guarantees recursive feasibility (see Theorem 3), results in a WT-MPC problem which is generally less conservative than RT-MPC.

The nominal constraints $z_{k|t} \in \mathcal{Z}_k$ can be easily enforced in the WT-MPC problem (see Remark 6 in the online

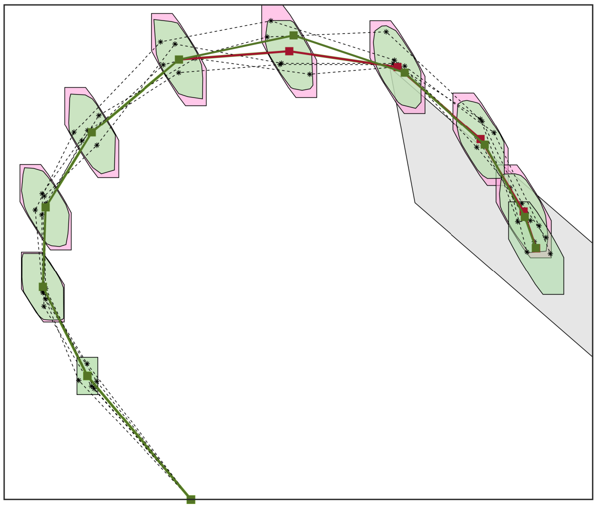

Fig. 1: Tubes evolution along the prediction horizon. The red and green solid lines denote the nominal state trajectory under the RT-MPC and WT-MPC policies, respectively. Black dashed lines represent random realizations of the error trajectory. The pink areas denote the robust tube $\{\mathcal{E}_k\}_{k=0}^N$ and the light green areas denote the tube $\{\mathcal{X}-\Gamma_k\}_{k=0}^N$ constructed numerically using Proposition 2. The grey area represents the terminal invariant set.

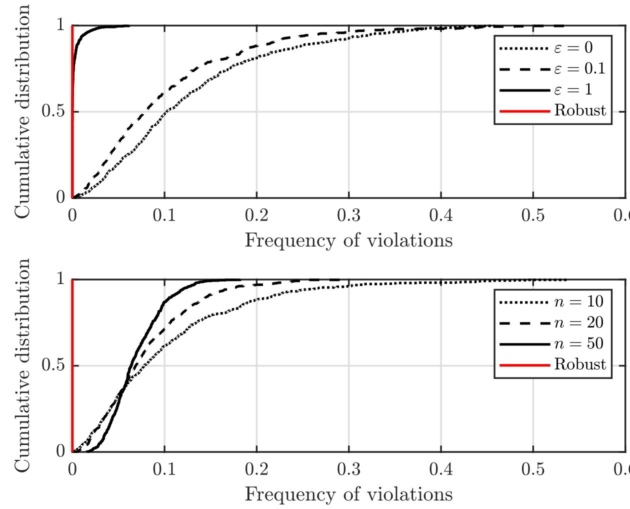

Fig. 2: Effect of Wasserstein radius (top panel) and number of samples (bottom panel) on the empirical frequency of violation during in open-loop. The case $\varepsilon = 0.01$ is almost-identical to $\varepsilon = 0$, hence it is not reported to ease the plot readability.

extended version [20]). For recursive feasibility, we need the following standard assumption on the terminal nominal set.

Assumption 2. There exists a terminal set $\mathcal{Z}_f$ satisfying:
(i) $K\mathcal{Z}_f \subseteq \mathcal{U} \ominus K\mathcal{E}_N$.
(ii) $A_K\mathcal{Z}_f \oplus A_K^N \mathcal{W} \subseteq \mathcal{Z}_f$.
(iii) $\mathcal{Z}_f \subseteq \mathcal{Z}_N$.

Conditions (i)-(iii) in Assumption 2 are naturally inherited from robust MPC [2], and they guarantee that $\mathcal{Z}_f$ is forward invariant at time N, robustly with respect to the initial condition, for zero control input.

Theorem 3. *Let Assumption 1 hold, let $\mathcal{Z}_k$ be defined as in* (9)*, and let $\mathcal{Z}_f$ satisfy Assumption 2. Then, WT-MPC is recursively feasible.*

In addition to recursive feasibility, whenever the ambiguity set $\mathbb{B}_\varepsilon^{\|\cdot\|_2}(\widehat{\mathbb{P}}_{[0]})$ contains the true distribution $\mathbb{P}$ of the noise, Theorem 3 guarantees that the closed-loop system satisfies $x_t \in \mathcal{X}$, with probability $1-\gamma$, $\forall t \in \mathbb{N}$.

IV. Numerical Experiments

To numerically validate the proposed WT-MPC, we consider the following system

$$x_{k+1} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0.5 \\ 1 \end{bmatrix} u_k + w_k,$$

borrowed from [2], with initial condition $x_0 = [-5, -2]^\top$. We let the support of the uncertainty $\mathcal{W}$ be the box $[-0.15, 0.15] \otimes [-0.15, 0.15]$, and we consider a noise which is uniformly distributed over $\mathcal{W}$, i.e., $w_t \sim \mathcal{U}(\mathcal{W})$, $\forall t \in \mathbb{N}$. Moreover, we assume to have access to a dataset consisting of n disturbance trajectories of length $N = 10$, for $n \in \{10, 20, 50\}$. For the WT-MPC problem, we consider the stage cost $\|z_{k|t}\|_Q^2 + \|v_{k|t}\|_R^2$ with $Q = \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right]$, and $R = 0.1$. Moreover, we consider the state constraint set

$$\begin{aligned} \mathcal{X} = \{x \in \mathbb{R}^2 : \max\{&\begin{bmatrix} 1 & 0 \end{bmatrix} x - 2, \begin{bmatrix} -1 & 0 \end{bmatrix} x - 10, \\ &\begin{bmatrix} 0 & 1 \end{bmatrix} x - 2, \begin{bmatrix} 0 & -1 \end{bmatrix} x - 2\} \leq 0\}, \end{aligned}$$

and the input constraint set $\mathcal{U} = \{u \in \mathbb{R} : -1 \leq u \leq 1\}$. Finally, we set $\gamma = 0.2$ in the DR-CVaR (7), and we choose an ambiguity radius $\varepsilon \in \{0, 0.01, 0.1, 1\}$. We compare:

- RT-MPC, which robustifies against all possible noise realizations within the support set $\mathcal{W}$.
- WT-MPC, which exploits the availability of data (the n noise trajectories), and robustifies against distributional uncertainty by tuning the ambiguity radius ε.

A. Open-Loop Analysis

We first compare the two methods in open-loop. Fig. 1 compares the open-loop behaviour of the system under the RT-MPC policy and the WT-MPC policy. Through the DR-CVaR constraint (7), WT-MPC relaxes the robust constraints into probabilistic ones, allowing a fraction of the state error trajectories $\{e_k\}_{k=0}^N$ to result in a user-defined probability of violating the constraint $x_k \in \mathcal{X}$. This is equivalent to considering the tube $\{\mathcal{X} - \Gamma_k\}_{k=0}^N$, which is visibly smaller than to the robust tube $\{\mathcal{E}_k\}_{k=0}^N$.

In the second open-loop experiment, we compute the empirical probability of violating the constraint $x_t \in \mathcal{X}$, $\forall t \in [1 : N]$, for 10000 noise trajectory realizations. For each parameter configuration (n, ε) we repeat the procedure 500 times, each time considering a different realization of the center distribution $\widehat{\mathbb{P}}_{[t-1]}$. Fig. 2 shows a parametric study of the WT-MPC policy by tuning the ambiguity radius ε (for fixed $n = 20$) and the number of noise sample trajectories n (for fixed $\varepsilon = 0.01$). We see that the constraint satisfaction increases with larger radii (i.e., more distributional robustness) and larger sample sizes (i.e., more knowledge about the true noise distribution). However, we observe that ε has a higher influence on the frequency of violations compared to n. This is primarily due to the fact that we use a very limited amount of noise sample trajectories: only $n \in \{10, 20, 50\}$ from the distribution of $\mathbf{w}_{[N-1]}$, which lives in dimension $Nd = 20$. Using few samples can be highly desirable in real-time applications, since a higher n translates to more constraints in the WT-MPC (see Proposition 2), and therefore a higher computational complexity. In that case, Fig. 2 shows that this can be done without sacrificing robustness, by simply picking a higher radius, which comes at no added computational complexity of the WT-MPC problem.

B. Closed-Loop Analysis

Next, we compare the performance of the two methods in closed-loop for a control task of length $T = 15$. Again, we consider $n \in \{10, 20, 50\}$ (for fixed $\varepsilon = 0.01$) and $\varepsilon \in \{0, 0.01, 0.1, 1\}$ (for fixed $n = 20$), and for each parameter configuration we repeat the procedure 100 times, each time considering differents realization of the noise. The out-of-sample performance in terms of both the empirical probability of violating the constraint $x_t \in \mathcal{X}$, $\forall t \in [1 : T]$, and the closed-loop cost are reported in Fig. 3.

Fig. 3 strengthens the observation made in the open-loop analysis: even in the presence of a small number of samples n, we can ensure (1) a desired robustness level for the closed-loop system, (2) smaller closed-loop cost (i.e., increased performance) compared to RT-MPC, and (3) good computational complexity, by simply adjusting the value of ε. In the bottom plot of Fig. 3, we observe that WT-MPC returns a policy which *optimally* trades between safety (i.e., constraint violation) and performance (i.e., closed-loop cost).

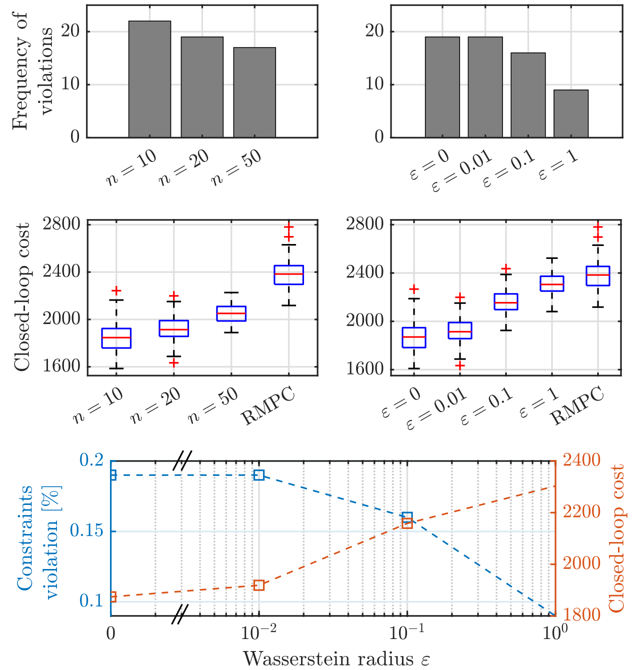

Fig. 3: Closed-loop analysis. Effect of sample size (top-left panel) and Wasserstein radius (top-right panel) on the frequency of violation in closed-loop. Closed-loop cost sensitivity to the two tuning knobs is reported in middle-left panel and middle-right panel, respectively. The bottom plot reports the trade-off between safety and performance as a function of Wasserstein radius.

REFERENCES

[1] Basil Kouvaritakis and Mark Cannon. Model predictive control. *Switzerland: Springer International Publishing*, 38, 2016.

[2] David Q Mayne, María M Seron, and SV Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.

[3] Wilbur Langson, Ioannis Chryssochoos, SV Raković, and David Q Mayne. Robust model predictive control using tubes. *Automatica*, 40(1):125–133, 2004.

[4] Ali Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.

[5] Basil Kouvaritakis, Mark Cannon, Saša V Raković, and Qifeng Cheng. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46(10):1719–1724, 2010.

[6] Mark Cannon, Basil Kouvaritakis, and Desmond Ng. Probabilistic tubes in linear stochastic model predictive control. *Systems & Control Letters*, 58(10-11):747–753, 2009.

[7] Mark Cannon, Basil Kouvaritakis, Saša V Raković, and Qifeng Cheng. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Transactions on Automatic Control*, 56(1):194–200, 2010.

[8] Matthias Lorenzen, Fabrizio Dabbene, Roberto Tempo, and Frank Allgöwer. Constraint-tightening and stability in stochastic model predictive control. *IEEE Transactions on Automatic Control*, 62(7):3165–3177, 2016.

[9] Lukas Hewing and Melanie N Zeilinger. Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control. *IEEE Control Systems Letters*, 4(2):450–455, 2019.

[10] Insoon Yang. Wasserstein distributionally robust stochastic control: A data-driven approach. *IEEE Transactions on Automatic Control*, 66(8):3863–3870, 2020.

[11] Christoph Mark and Steven Liu. Data-driven distributionally robust mpc: An indirect feedback approach. *arXiv preprint arXiv:2109.09558*, 2021.

[12] Marta Fochesato and John Lygeros. Data-driven distributionally robust bounds for stochastic model predictive control. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 3611–3616. IEEE, 2022.

[13] Francesco Micheli, Tyler Summers, and John Lygeros. Data-driven distributionally robust mpc for systems with uncertain dynamics. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4788–4793. IEEE, 2022.

[14] Atharva Navsalkar and Ashish R Hota. Data-driven risk-sensitive model predictive control for safe navigation in multi-robot systems. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1442–1448. IEEE, 2023.

[15] Alireza Zolanvari and Ashish Cherukuri. Data-driven distributionally robust iterative risk-constrained model predictive control. In *2022 European Control Conference (ECC)*, pages 1578–1583. IEEE, 2022.

[16] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.

[17] Soroosh Shafieezadeh-Abadeh, Liviu Aolaritei, Florian Dörfler, and Daniel Kuhn. New perspectives on regularization and computation in optimal transport-based distributionally robust optimization. *arXiv preprint arXiv:2303.03900*, 2023.

[18] Liviu Aolaritei, Nicolas Lanzetti, and Florian Dörfler. Capture, propagate, and control distributional uncertainty. *arXiv preprint arXiv:2304.02235*, 2023.

[19] Liviu Aolaritei, Nicolas Lanzetti, Hongruyu Chen, and Florian Dörfler. Distributional uncertainty propagation via optimal transport. *arXiv preprint arXiv:2205.00343*, 2023.

[20] Liviu Aolaritei, Marta Fochesato, John Lygeros, and Florian Dörfler. Wasserstein tube mpc with exact uncertainty propagation. *arXiv preprint arXiv:2304.12093*, 2023.

[21] Arkadi Nemirovski and Alexander Shapiro. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17(4):969–996, 2007.