

# Stochastic Model Predictive Control: Controlling the Average Number of Constraint Violations<sup>1</sup>

Milan Korda, Ravi Gondhalekar, Frauke Oldewurtel and Colin N. Jones

**Abstract**—This paper considers linear discrete-time systems with additive bounded disturbances subject to hard control input bounds and a stochastic requirement on the number of state-constraint violations averaged over time. This specification facilitates the exploitation of the information on the number of past constraint violations, and consequently enables a significant reduction in conservatism. For the type of constraint considered we develop a recursively feasible receding horizon scheme, and, as a simple modification of our approach, we show how a bound on the average number of violations can be enforced robustly. The computational complexity (online as well as offline) is comparable to existing model predictive control schemes. The effectiveness of the proposed methodology is demonstrated by means of a numerical example.

## I. INTRODUCTION

There is a significant gap between the theory of model predictive control (MPC) and its practical usage. Indeed, theoretical results on stability and recursive feasibility of MPC are available in nominal as well as robust settings [16, 22]. These results are, however, rarely used in practice: in most applications the problem of recursive feasibility is circumvented by employing soft constraints, and (robust) stability is not enforced by design but evaluated a posteriori. The main reason for this discrepancy is the excessive conservatism of theoretically sound approaches compared to empirical methods, especially in a robust setting [20].

A possible remedy is for the constraint specification to be probabilistic rather than robust, or, more generally, for occasional constraint violations to be allowed in a well-defined, but not necessarily probabilistic, manner. An MPC controller can then exploit these occasional violations to achieve a lower objective cost, thereby reducing the conservatism.

The main theoretical challenge when occasional constraint violations are allowed is recursive feasibility. Previous approaches [5, 6, 7, 8, 14, 15] considered constraints on the marginal distribution of the state, typically point-wise in time probabilistic constraints. In those works the constraints were enforced by controlling the conditional probability of constraint violations between two consecutive time instances without taking into account the past behavior of the state process. Although convenient for a receding horizon implementation, this necessarily limits the achievable benefits of the probabilistic specification. Indeed, even the least-restrictive<sup>2</sup> formulation of [14] is still conservative in certain situations.

Milan Korda and Colin N. Jones are with the Laboratoire d'Automatique, École Polytechnique Fédérale de Lausanne, Switzerland. {milan.korda, colin.jones}@epfl.ch.

Ravi Gondhalekar is with the Department of Chemical Engineering, University of California Santa Barbara, USA. gondhalekar@engineering.ucsb.edu.

Frauke Oldewurtel is with the Power Systems Laboratory, Department of Electrical Engineering, Swiss Federal Institute of Technology in Zurich (ETHZ), Switzerland. oldewurtel@eeh.ee.ethz.ch

<sup>1</sup>This is a corrected version of the one that appeared in the proceedings of CDC 2012.

<sup>2</sup>Least-restrictive in terms of the size of the feasible domain with respect to the one-step conditional constraint (Section II, Eq. (5)).

In contrast, the proposed method exploits the information on the past behavior of the state process – namely the number of past constraint violations – when deciding on the current control input, and thereby is capable of enforcing probabilistic properties on the closed-loop trajectory as a whole. The quantity constrained in this paper is the number of constraint violations averaged over time; we derive an MPC control scheme such that bounds on this quantity are satisfied either in a stochastic manner or robustly. These constraints are equally expressive and practically relevant as the traditional probabilistic constraints (considered in, e.g., [14, 15]), and readily facilitate exploiting the knowledge of the number of past constraint violations. For instance, this type of constraint is natural in building climate control, where comfort specifications for the room temperature are of this form [18].

The presented methodology is based on the first-step stochastic invariance introduced in [14] and originally developed for nominal MPC in [10]. The approach is completely independent of the MPC problem cost function and policy parametrization, and introduces only affine constraints.

This flexibility is facilitated by the use of controlled invariant sets, which can be parametrized either explicitly or implicitly. In either case, the computational requirements are comparable to their respective nominal and robust counterparts – in the case of explicit parametrization to the first-step nominal MPC of [9, 10]; in the case of implicit parametrization to any of the traditional robust / stochastic MPC schemes such as affine disturbance feedback [12, 14], prestabilization [15] or tubes [17] (assuming the presented approach is used with the same parametrization as the traditional one).

The paper is organized as follows. The problem to be solved is formulated in Section II; a general set-based solution is described in Section III for the stochastic constraint specification; Section III-B then presents a modification for the robust one. Section IV shows how the presented methodology can be incorporated into an MPC framework, discusses the explicit and implicit policy parametrizations (IV-A) and some computational aspects (IV-B). A numerical example is presented in Section V.

## A. Notation

Throughout the article  $\mathbb{R}$  denotes the set of reals,  $\mathbb{N}_+$  the set of positive integers,  $\mathbb{N}$  the set of nonnegative integers and  $\mathbb{N}_i^j$  denotes the set of consecutive integers  $\{i, i+1, \dots, j\}$ . Similarly, let also  $x_i^j := (x_i, x_{i+1}, \dots, x_j)$ . An infinite sequence  $(x_i, x_{i+1}, \dots)$ ,  $x_j \in \mathcal{M}$  for all  $j \geq i$ , is denoted by  $\{x_j \in \mathcal{M}\}_{j=i}^\infty$ . Random variables are defined on a common probability space with an associated probability measure  $P(\cdot)$ . The expected value operator is denoted by  $\mathbf{E}\{\cdot\}$ ; the conditional probability and expectation based on information available up to time  $t$  are denoted by  $P_t(\cdot)$  and  $\mathbf{E}_t\{\cdot\}$ , respectively. For a set  $\mathcal{M}$  and  $N \in \mathbb{N}_+$ , let  $\mathcal{M}^N$  be the Cartesian product of the set  $N$ -times with itself. Finally, let  $\lfloor \cdot \rfloor$  denote the floor of a real number.

## II. PROBLEM STATEMENT

We consider the linear time-invariant stochastic dynamic system

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad t \in \mathbb{N} \quad (1)$$

with the state  $x_t \in \mathbb{R}^n$ , the control  $u_t \in \mathbb{R}^m$ , and the i.i.d. disturbance sequence  $w_t \in \mathbb{R}^n$ . We assume that the state  $x_t$  is known at time  $t$  for all  $t \in \mathbb{N}$ , and that the pair  $(A, B)$  is stabilizable. The control inputs are subject to hard bounds of the form

$$u_t \in \mathcal{U} \subseteq \mathbb{R}^m, \quad t \in \mathbb{N}, \quad (2)$$

for some polyhedral input constraint set  $\mathcal{U}$ . We let  $w$  denote any random variable having the common distribution of  $w_t$ ,  $t \in \mathbb{N}$ . The support of  $w$  is assumed to be contained in a compact polyhedron  $\mathcal{W}$ , that is,

$$\text{supp}(w) = \text{supp}(w_t) \subseteq \mathcal{W} \subset \mathbb{R}^n, \quad t \in \mathbb{N}. \quad (3)$$

The essence of the paper is then the handling of probabilistic constraints on the closed-loop state process. These constraints can take various forms, but in loose terms one wants the state to stay within a given polyhedron  $\mathcal{X}$  “most of the time”. The polyhedron  $\mathcal{X}$  is referred to as the *(state) constraint set*.

For the sake of brevity, we focus exclusively on so-called *joint chance constraints*, where the probabilistic requirements are imposed on the polyhedron  $\mathcal{X}$  as a whole, although an extension to *individual chance constraints*, where the requirements are imposed on each hyperplane defining  $\mathcal{X}$  separately, is straightforward.

### A. Probabilistic constraint formulations

In the remainder of this section we state the specific type of constraints considered in this paper and relate them to other types of probabilistic constraints considered in the past. Our previous work [14] considered constraints of the form<sup>3</sup>

$$P(x_t \in \mathcal{X}) \geq 1 - \alpha, \quad \forall t \in \mathbb{N}_+, \quad (4)$$

with  $\alpha \in [0, 1)$ , which were enforced through the sufficient *one-step (conditional) constraint*

$$P(x_{t+1} \in \mathcal{X} \mid x_t) \geq 1 - \alpha, \quad \forall t \in \mathbb{N}. \quad (5)$$

Being a function of  $x_t$  (and hence of  $w_0^{t-1}$ ), the left-hand side of this constraint is a random variable and the inequality is required to hold for all realizations of  $x_t$  under a given control policy. The use of conditioning may, however, bring about significant conservatism and, consequently, almost negate the benefits of the probabilistic constraint specification. Indeed, the relationship

$$\mathbf{E}\{P(x_{t+1} \in \mathcal{X} \mid x_t)\} = P(x_{t+1} \in \mathcal{X}) \quad (6)$$

clearly manifests the potential conservatism introduced by the conditioning in (5).

In this paper we almost entirely eliminate this conservatism by requiring that the closed-loop state process satisfies

$$\begin{cases} \mathbf{E}_t \left\{ \frac{v_{t+1}}{t+1} \right\} \leq \alpha & \text{if } \frac{v_t}{t} \leq \alpha \\ \frac{v_{t+k}}{t+k} \searrow \alpha \text{ as } k \rightarrow \tau_t & \text{if } \frac{v_t}{t} > \alpha, \end{cases} \quad (7)$$

<sup>3</sup>In [14] we actually considered individual chance constraints only, but the extension to the joint chance constraints (4) and (5) is straightforward.

where  $v_t$  denotes the number of constraint violations<sup>4</sup> up to time  $t$ , that is,

$$v_t := \begin{cases} \sum_{k=1}^t \mathbb{I}_{[x_k \notin \mathcal{X}]} & t \in \mathbb{N}_+, \\ 0 & t = 0, \end{cases} \quad (8)$$

with  $\mathbb{I}_{[x_k \notin \mathcal{X}]}$  denoting the indicator function of the complement of  $\mathcal{X}$  evaluated in  $x_k$ . The integer random variable  $\tau_t$  is then the first return time after  $t$  of the average number of violations below the required level  $\alpha$ , i.e.,

$$\tau_t := \inf\{i > t \mid v_i/i \leq \alpha\}, \quad (9)$$

where, as usual, the infimum of an empty set is infinity. The convergence in the second line of (7) is required to be with probability one (or “almost surely”).

In other words, rather than controlling the probability of constraint violation at each time separately, at each time instant  $t$  we control the average number of constraint violations, where the average is taken over the entire history of the state trajectory up to time  $t$ . The average number of violations is controlled in such a manner that whenever it is below the prescribed level  $\alpha$  (line one of (7)), then it remains below  $\alpha$  in expectation at next time instant; if the average number of violations happens to be above the level of  $\alpha$  (line two of (7)), we require that it converges back towards  $\alpha$  with probability one.

**Remark 1.** *The allowed level of constraint violation  $\alpha \in [0, 1)$  typically comes directly from application requirements, but it can also be viewed as a tuning parameter adjusting the conservativeness of the controller.*

As a simple modification of our approach, we show how to enforce the averaged constraint robustly, that is, how to enforce

$$\frac{1}{t} \sum_{k=1}^t \mathbb{I}_{[x_k \notin \mathcal{X}]} \leq \alpha, \quad \forall t \in \mathbb{N}_+, \quad \forall \{w_t \in \mathcal{W}\}_{t=0}^\infty. \quad (10)$$

Performance under this constraint should be compared to the classical robust MPC rather than to its stochastic counterparts, because in (10) there is no randomness in the *number* of violations allowed to occur up to a time  $t$  – the only uncertainty is about *when* the violations occur.

## III. MAIN RESULTS

In this section we present a recursively feasible receding horizon control policy under which the closed-loop state process satisfies the constraint (7). The main idea is simple: keep track of the number of past constraint violations; if it is “large”, enforce the constraint (5) as is; if it is “small”, loosen the constraint (5) appropriately; and at all times enforce such invariance constraints that the constraint (5) will be feasible at the next time instant. One can then extend this idea and build a family of nested sets in which the state is allowed to climb up if the number of violations is “small” and is forced to climb down if the number gets “large”; the (possibly loosened) probabilistic constraint (5) is then enforced only at the lowermost level.

Now we give a precise mathematical formulation of the above discussion. We will repeatedly employ joint chance constraints of the form

$$P(Ax + Bu + w \in \mathcal{X}) \geq 1 - \delta, \quad (11)$$

<sup>4</sup>The violation at time zero (which cannot be prevented) is omitted just for notational convenience. All our results

for some  $x \in \mathbb{R}^n$ ,  $u \in \mathbb{R}^m$  and some  $\delta \in [0, 1]$ . The left-hand side of these constraints is, however, difficult to evaluate precisely; hence we make use of a  $(1 - \delta)$ -confidence region of  $w$ , which can be any set  $\mathcal{R}_\delta \subseteq \mathcal{W}$  such that

$$P(w \in \mathcal{R}_\delta) \geq 1 - \delta. \quad (12)$$

A sufficient condition for (11) to hold is then

$$Ax + Bu + w \in \mathcal{X} \quad \forall w \in \mathcal{R}_\delta. \quad (13)$$

In the sequel we make use of the following assumption.

**Assumption 1.** A compact confidence region  $\mathcal{R}_\alpha$  for the constraint violation level  $\alpha$  has been determined.

For example, ellipsoidal or polytopic confidence regions are typically straightforward to construct and suitable from a computational point of view. Note that the assumption of  $\mathcal{R}_\alpha$  being compact has been made only for computational tractability. See Section IV-B for computational aspects.

We start the exposition with an instructional, single-layer, approach which is then extended to a more powerful multi-layer approach in Section III-A.

The first key ingredient in our approach is the stochastic feasibility set  $\mathcal{X}_s$ :

**Definition 1** (Stochastic feasibility set [14]). *The stochastic feasibility set of the constraint (5) associated with the confidence region  $\mathcal{R}_\alpha$  is*

$$\mathcal{X}_s := \{x \in \mathbb{R}^n \mid \exists u \in \mathcal{U} \text{ s.t. } Ax + Bu + w \in \mathcal{X} \quad \forall w \in \mathcal{R}_\alpha\}.$$

The subscript  $s$  signifies ‘‘stochastic’’ in contrast to a robust version  $\mathcal{X}_r$  defined in Section III-B.

In plain words,  $\mathcal{X}_s$  is a subset (because of the confidence region possibly strict) of the set of states for which there exists an admissible input such that the conditional constraint (5) is satisfied. Because the state and the input constraint sets  $\mathcal{X}$  and  $\mathcal{U}$  are polyhedra,  $\mathcal{X}_s$  is also a polyhedron for any compact confidence region  $\mathcal{R}_\alpha$  (see Section IV-B).

The second key ingredient is a stochastic robust controlled invariant set (SRCI set)  $\mathcal{S}$ :

**Definition 2** (SRCI set [14]). *A set  $\mathcal{S} \subseteq \mathcal{X}_s$  is a stochastic robust controlled invariant set if it satisfies the following condition:*

$$\forall x \in \mathcal{S} \exists u \in \mathcal{U} \text{ s.t. : } Ax + Bu + w \in \mathcal{S} \quad \forall w \in \mathcal{W}, \\ Ax + Bu + w \in \mathcal{X} \quad \forall w \in \mathcal{R}_\alpha.$$

We need the following assumption on  $\mathcal{S}$ :

**Assumption 2.** A nonempty polyhedral SRCI set  $\mathcal{S}$  exists and has been characterized.

The requirement of  $\mathcal{S}$  being polyhedral has been imposed for the sake of computational tractability (see Section IV); all theoretical results presented in this paper hold for any nonempty SRCI set  $\mathcal{S}$ .

Now we proceed to define a quantity called *probability leeway* that will control the loosening of the one-step conditional constraint. It will be the only parameter adjusting constraint tightness in the single-layer approach, and it will act at the lowermost level of the multi-layer approach. The first line constraint (7) is then equivalent to

$$\mathbf{E}_t\{v_{t+1}\} \leq \alpha(t+1), \quad (14)$$

and therefore (7) will be satisfied at time  $t$  if we ensure that

$$(t+1)\alpha - \mathbf{E}_t\{v_{t+1}\} = \alpha + \alpha t - v_t - P_t(x_{t+1} \notin \mathcal{X})$$

is nonnegative. Consequently, the quantity to be controlled at time  $t$  is  $P_t(x_{t+1} \notin \mathcal{X})$ , the conditional probability of violating the state constraint at time  $t+1$  given information at time  $t$ . To that effect define the probability leeway  $\beta_t \in [\alpha, 1]$  at time  $t \in \mathbb{N}$  as

$$\beta_t := \max\{\min\{\alpha + \alpha t - v_t, 1\}, \alpha\}. \quad (15)$$

In other words,  $\beta_t$  is the highest probability level such that whenever  $\beta_t > \alpha$ , enforcing

$$P_t(x_{t+1} \notin \mathcal{X}) \leq \beta_t \quad (16)$$

guarantees the satisfaction of the first line of (7) at time  $t$ . If, on the other hand,  $\beta_t = \alpha$  (or more precisely  $\alpha t < v_t$ ), then enforcing (16) does not guarantee the satisfaction of  $\mathbf{E}_t\{v_{t+1}/(t+1)\} \leq \alpha$ . However, the probability leeway  $\beta_t$  can be used to define a control law that guarantees the satisfaction of (7) as a whole.

To this end define for all  $t \in \mathbb{N}$

$$\tilde{\mathcal{U}}_t(x_t, v_t) := \{u \in \mathcal{U} \text{ s.t.} \\ Ax_t + Bu_t + w \in \mathcal{S} \quad \forall w \in \mathcal{W} \quad (17a)$$

$$\beta_t < 1 \Rightarrow Ax_t + Bu_t + w \in \mathcal{X} \quad \forall w \in \mathcal{R}_{\beta_t}\}, \quad (17b)$$

where  $\mathcal{R}_{\beta_t}$  is a  $(1 - \beta_t)$ -confidence region of  $w$  and  $\beta_t$  is defined through  $v_t$  and  $t$  by (15). Note in particular that the constraint (17b) is disabled for  $\beta_t = 1$ . We require the following (natural) assumption on the confidence regions:

**Assumption 3.**  $\mathcal{R}_{\beta_t} \subseteq \mathcal{R}_\alpha$  for all  $\beta_t \in [\alpha, 1)$ .

A basic single-layer set-valued control policy under which the satisfaction of (7) is guaranteed can now be defined by

$$\tilde{\kappa}_t(x_t, v_t) \in \tilde{\mathcal{U}}_t(x_t, v_t), \quad t \in \mathbb{N}. \quad (18)$$

In other words, at time  $t$  we enforce the invariance constraint (17a), which ensures that the one-step constraint (17b) is feasible at all times whenever it is feasible at time zero (since we assume  $\mathcal{R}_{\beta_t} \subseteq \mathcal{R}_\alpha$ ); if the number of past violations  $v_t$  is at least  $1 - \alpha$  less than the specified level  $\alpha t$  (i.e.,  $\beta_t = 1$ ), we do not enforce any additional constraints; if, on the other hand, the number of past violations exceeds the specified level  $\alpha t$  (i.e.,  $\beta_t = \alpha$ ), we also enforce the one-step conditional constraint (5) through (17b) with  $\beta_t = \alpha$ ; finally, if the difference between the specified level and the current number of violations is between zero and  $1 - \alpha$  (i.e.,  $\beta_t \in (\alpha, 1)$ ), we enforce the one-step constraint loosened as much as possible such that (7) is still guaranteed to be (conditionally) satisfied at the next time instant.

**Remark 2.** *The control policy (18) is time-varying and non-Markovian since  $\beta_t$  depends on time  $t$  and on the number of past constraint violations  $v_t$ .*

**Remark 3.** *It is not necessary to construct a confidence region for each value of  $\beta_t \in [\alpha, 1)$ . First, in view of (15), for any rational value of  $\alpha$ , there are only finitely many possible values of  $\beta_t$  between  $\alpha$  and one. Second, for arbitrary value of  $\alpha$ , one can always choose fixed values  $\alpha =: \hat{\beta}_1 < \hat{\beta}_2 < \dots < \hat{\beta}_{n_\beta} < 1$  for which confidence regions  $\mathcal{R}_{\hat{\beta}_i}$  are precomputed and then round  $\beta_t < 1$  to the nearest lower value of  $\hat{\beta}_i$ . In addition, the confidence regions can, and*

typically are, chosen such that only the scaling varies while the shape remains the same for each value of  $\beta_i$ . The storage of these regions then amounts to the storage of the shape of a single region and  $n_{\beta}$  real numbers representing the scalings.

A formal proof of the intuitively obvious fact that, given  $\tilde{\mathcal{U}}_t(x_0, 0) \neq \emptyset$ , the closed loop state process under the control law  $u_t = \tilde{\kappa}_t(x_t, v_t)$  is defined at all times and satisfies the constraint (7) is given in a more general, multi-layer, setting in Theorem 1.

#### A. Multi-layer version

Under the control policy (18), the invariance constraint (17a) is independent of the number of past constraint violations, no matter how small it is. As a consequence, the invariance constraint (17a) can be a major, if not sole, source of conservatism for small  $v_t$  (i.e., for  $\beta_t$  close or equal to one).

The multi-layer approach presented in this section alleviates this by loosening the invariance constraints with a decreasing number of past violations. The idea is to construct a family of nested one-step reachability sets around the SRCI set  $\mathcal{S}$ . To this end, define the robust reachability operator of a set  $\mathcal{M} \subseteq \mathbb{R}^n$  as

$$\text{Reach}(\mathcal{M}) \quad (19)$$

$$:= \{x \in \mathbb{R}^n \mid \exists u \in \mathcal{U} \text{ s.t. } Ax + Bu + w \in \mathcal{M} \forall w \in \mathcal{W}\}.$$

The nested family of length  $n_s$  is then given by

$$\mathcal{S}_1 := \mathcal{S}, \quad (20a)$$

$$\mathcal{S}_{k+1} := \text{Reach}(\mathcal{S}_k), \quad k = 1, \dots, n_s - 1. \quad (20b)$$

**Remark 4.** The nested property  $\mathcal{S}_k \subseteq \mathcal{S}_{k+1}$  follows from the fact that  $\mathcal{M} \subseteq \text{Reach}(\mathcal{M})$  whenever the set  $\mathcal{M}$  is robust controlled invariant and the fact that the  $\text{Reach}(\cdot)$  operator preserves the invariance.

**Remark 5.** When the sets  $\mathcal{M}$  and  $\mathcal{U}$  are polyhedral, the evaluation of  $\text{Reach}(\mathcal{M})$  amounts to a single polyhedral projection; consequently,  $\text{Reach}(\mathcal{M})$  is also polyhedral (unless empty).

Now we proceed to the definition of an integer random variable that will play the role of a layer index that controls to which layer (i.e., to which set  $\mathcal{S}_k$ ) the state is permitted to move. Observe that

$$\beta_t = 1 \Leftrightarrow \alpha t - v_t \geq 1 - \alpha,$$

and, for any  $k \in \mathbb{N}$ ,


$$\begin{aligned} \beta_{t+k} = 1 &\Leftrightarrow \alpha(t+k) - v_{t+k} \geq 1 - \alpha \\ &\Leftrightarrow \alpha t - v_t \geq (k+1)(1 - \alpha). \end{aligned}$$

In other words, if we know at time  $t$  that

$$\alpha t - v_t \geq (k+1)(1 - \alpha),$$

then we are guaranteed that the first time the probabilistic constraint (17b) needs to be enforced is  $t+k+1$ ; therefore we can permit the state  $x_{t+1}$  to move to the set  $\mathcal{S}_{k+1}$ , from which we are guaranteed to get back to  $\mathcal{S}_1 = \mathcal{S}$  at the time  $t+k+1$  if necessary. This suggests that the sets  $\mathcal{S}_k$  should be indexed by the number of integer multiples of  $1 - \alpha$  in  $\alpha t - v_t$ . Thus, define the layer index  $r_t \in \mathbb{N}_1^{n_s}$  as

$$r_t := \max \left\{ \min \left\{ \left\lfloor \frac{\alpha t - v_t}{1 - \alpha} \right\rfloor, n_s \right\}, 1 \right\}, \quad t \in \mathbb{N}. \quad (21)$$



**Fig. 1:** Structure of the sets involved and some of the possible evolutions of the state over several time steps. Note the nested property  $\mathcal{S} = \mathcal{S}_1 \subseteq \mathcal{S}_2 \subseteq \mathcal{S}_3$  (the colored sets), the inclusion  $\mathcal{S} = \mathcal{S}_1 \subseteq \mathcal{X}_s$ , and that  $\mathcal{X}_s \not\subseteq \mathcal{X}$  nor  $\mathcal{X} \not\subseteq \mathcal{X}_s$ . The dashed lines show some of the feasible evolutions for the given values of  $\beta_t$  and  $r_t$ . We start with  $x_t \in \mathcal{S}_1$  and  $r_t = 2$ ; hence the state is permitted to move to  $\mathcal{S}_2$ . From there we assume two scenarios: if there is no violation at  $t+1$ , then  $r_t$  may increase to  $r_{t+1} = 3$  in which case the state is allowed to move to  $\mathcal{S}_3$ ; if there is a violation at  $t+1$ , then  $r_{t+1} = 1$  (and  $\beta_{t+1} = 1$ ) and the state has to move back to  $\mathcal{S}_1$ . The latter scenario is further expanded: if there is no violation then we have  $r_{t+2} = 1$  (and  $\beta_{t+2} = 1$ ) in which case only the inclusion to  $\mathcal{S}_1$  is enforced; if there is a violation at  $t+2$ , then we need to enforce the (loosened) probabilistic constraint, here with  $\beta_{t+2} = 0.7$ . The former scenario is expanded still further assuming that a violation occurred at time  $t+3$ , which requires enforcing a tighter probabilistic constraint than before, here with  $\beta_{t+3} = 0.3$ . It is important to note that the evolution of  $r_t$  and  $\beta_t$  depends on the specific values of  $\alpha t - v_t$  and  $\alpha$ , which means that the scenarios shown are not the only ones possible. Consistent values for these scenarios are  $\alpha = 0.3$  and  $\alpha t - v_t = 1.8$ .

Now we can define a set-valued multi-layer control policy. To this end define for all  $t \in \mathbb{N}$  the sets

$$\mathcal{U}_t(x_t, v_t) := \{u \in \mathcal{U} \text{ s.t. } Ax_t + Bu_t + w \in \mathcal{S}_{r_t} \forall w \in \mathcal{W}, \quad (22a)$$

$$\beta_t < 1 \Rightarrow Ax_t + Bu_t + w \in \mathcal{X} \forall w \in \mathcal{R}_{\beta_t}\}, \quad (22b)$$

and

$$\Pi_t := \{(x_t, v_t) \mid \mathcal{U}_t(x_t, v_t) \neq \emptyset\}. \quad (23)$$

The multi-layer control policy is then defined by

$$\kappa_t(x_t, v_t) \in \mathcal{U}_t(x_t, v_t), \quad t \in \mathbb{N}. \quad (24)$$

The Remarks 2 and 3 about the single-layer policy (18) apply also to this control policy. Note also that, since  $\mathcal{S}_1 = \mathcal{S}$ , we recover the single-layer control policy for  $n_s = 1$ .

A typical evolution of the state within the nested structure is depicted in Figure 1. The sets are drawn as ellipses although in fact they are polyhedra under Assumption 2.

Now we can state and prove our main result:

**Theorem 1.** Under the control law  $u_t = \kappa_t(x_t, v_t)$  the following holds:

- I. If  $x_0 \in \mathcal{S}$ , then  $(x_0, 0) \in \Pi_0$  (initial feasibility).
- II. If  $(x_t, v_t) \in \Pi_t$ , then  $(x_{t+1}, v_{t+1}) \in \Pi_{t+1}$  (recursive feasibility).
- III. If  $(x_0, 0) \in \Pi_0$ , then  $x_t$  satisfies the constraint (7) (closed-loop constraint satisfaction).

*Proof.* I. At time zero we have  $r_t = 1$ , and feasibility of (22a) and (22b) then follows from the definition of  $\mathcal{S} = \mathcal{S}_1$ .

II. Consider first  $r_t = 1$ . Given feasibility at time  $t$ , we know that the state at time  $t+1$  will be in  $\mathcal{S}_1 = \mathcal{S}$ . Therefore, by definition of  $\mathcal{S}$ , the invariance constraint (22a) as well as the probabilistic constraint (22b) will be feasible at time  $t+1$ . Note further that, from (15),  $\beta_{t+1} < 1$  implies  $r_t = 1$ ; hence,  $r_t = 1$  is the only case when the constraint (22b) can be active at time  $t+1$ .

Next, for  $r_t > 1$  we know that the state at time  $t+1$  will be in  $\mathcal{S}_{r_t} = \text{Reach}(\mathcal{S}_{r_t-1})$ . Feasibility of the invariance constraint (22a) at time  $t+1$  follows since  $r_t$  can decrease by at most one over one time step ( $r_{t+1} \geq r_t - 1$ ).

III. Consider first  $t \in \mathbb{N}_+$  such that  $v_t/t \leq \alpha$ . In that case we have, by definition of  $\beta_t$ , that

$$\begin{aligned} \mathbf{E}_t\{v_{t+1}\} &= v_t + P_t(x_{t+1} \notin \mathcal{X}) \leq v_t + \beta_t \\ &= v_t + \alpha + t\alpha - v_t = (t+1)\alpha \end{aligned}$$

as desired by the first line of (7).

Consider now a time instant  $t \in \mathbb{N}_+$  such that  $v_t/t > \alpha$  and let  $\tau_t$  defined in (9) be the first return time of  $v_t/t$  below  $\alpha$ . Clearly, for those events on which  $\tau_t < \infty$ , the second line of (7) is satisfied. Assume therefore  $\tau_t = \infty$ . Define further the process  $\gamma_k := v_{t+k} - (t+k)\alpha$  for  $k \in \mathbb{N}$ . Then

$$\begin{aligned} \gamma_{k+1} &= v_{t+k} - (t+k)\alpha + \mathbb{I}_{[x_{t+k+1} \notin \mathcal{X}]} - \alpha \\ &= \gamma_k + \mathbb{I}_{[x_{t+k+1} \notin \mathcal{X}]} - \alpha. \end{aligned}$$

Now, since  $v_{t+k}/(t+k) > \alpha$ , we have that that  $\beta_{t+k} = \alpha$  for all  $k \in \mathbb{N}$  and therefore

$$\begin{aligned} \mathbf{E}_{t+k}\{\gamma_{k+1}\} &= \gamma_k + \mathbf{E}_{t+k}\{\mathbb{I}_{[x_{t+k+1} \notin \mathcal{X}]}\} - \alpha \\ &= \gamma_k + P_{t+k}(x_{t+k+1} \notin \mathcal{X}) - \alpha \leq \gamma_k. \end{aligned}$$

As a result, the process  $\gamma_k$ ,  $k \in \mathbb{N}$ , is a nonnegative supermartingale. Now by Doob's martingale convergence theorem,  $\gamma_k$  converges to some random variable  $l \geq 0$ . Therefore,

$$\lim_{k \rightarrow \infty} \frac{v_{t+k}}{t+k} - \xi = \frac{\gamma_k}{t+k} = \lim_{k \rightarrow \infty} \frac{l}{t+k} = 0$$

with probability one as desired.  $\square$

### B. Robust average constraint

In this section we outline a modification of the presented approach for the robust averaged constraint (10). Unlike the stochastic requirement (7), satisfaction of (10) guarantees a hard bound on the average number of constraint violations, and hence may be preferable in situations where mere satisfaction in expectation is not sufficient (i.e., it is "too random").

The modification consists of replacing the confidence region  $\mathcal{R}_\alpha$  by the whole support of the disturbance  $\mathcal{W}$  in the construction of the stochastic feasibility set and the SRCI set. Equivalently, the stochastic feasibility set is now replaced by the one-step robust reachability set  $\mathcal{X}_r$  of the constraint set  $\mathcal{X}$ , that is,

$$\mathcal{X}_r := \text{Reach}(\mathcal{X}).$$

The SRCI set  $\mathcal{S}$  is then replaced by a robust controlled invariant (RCI) subset of  $\mathcal{X}_r$ , that is, by a set  $\hat{\mathcal{S}} \subseteq \mathcal{X}_r$  that satisfies

$$\begin{aligned} \forall x \in \hat{\mathcal{S}} \exists u \in \mathcal{U} \text{ s.t. : } Ax + Bu + w \in \hat{\mathcal{S}} \quad \forall w \in \mathcal{W}, \\ Ax + Bu + w \in \mathcal{X} \quad \forall w \in \mathcal{W}. \end{aligned}$$

The construction of the nested family is the same as before:

$$\hat{\mathcal{S}}_1 := \hat{\mathcal{S}}, \quad (25a)$$

$$\hat{\mathcal{S}}_{k+1} := \text{Reach}(\hat{\mathcal{S}}_k), \quad k = 1, \dots, n_s - 1. \quad (25b)$$

Similarly, as before, define now for all  $t \in \mathbb{N}$

$$\hat{\mathcal{U}}_t(x_t, v_t) := \{u \in \mathcal{U} \text{ s.t.}$$

$$Ax_t + Bu_t + w \in \hat{\mathcal{S}}_{r_t} \quad \forall w \in \mathcal{W}, \quad (26a)$$

$$\beta_t < 1 \Rightarrow Ax_t + Bu_t + w \in \mathcal{X} \quad \forall w \in \mathcal{W}\}, \quad (26b)$$

and

$$\hat{\Pi}_t := \{(x_t, v_t) \mid \hat{\mathcal{U}}_t(x_t, v_t) \neq \emptyset\}, \quad (27)$$

where  $\beta_t$  and  $r_t$  are defined in (15) and (21). The robust, multi-layer control policy is then given by

$$\hat{\kappa}_t(x_t, v_t) \in \hat{\mathcal{U}}_t(x_t, v_t), \quad t \in \mathbb{N}. \quad (28)$$

The structure of the control policy is the same as that of (24) – only the sets to which a robust inclusion is enforced differ. In particular the probabilistic constraint (22b), a robust inclusion to a confidence region, is replaced by (26b), a robust inclusion to the constraint set  $\mathcal{X}$ . The activation of the two constraints is, however, governed by the same random variable  $\beta_t$  defined in (15).

The following theorem is analogous to Theorem 1.

**Theorem 2.** Under the control law  $u_t = \hat{\kappa}_t(x_t, v_t)$  the following holds:

- I. If  $x_0 \in \hat{\mathcal{S}}$ , then  $(x_0, 0) \in \hat{\Pi}_0$ .
- II. If  $(x_t, v_t) \in \hat{\Pi}_t$ , then  $(x_{t+1}, v_{t+1}) \in \hat{\Pi}_{t+1}$ .
- III. If  $(x_0, 0) \in \hat{\Pi}_0$ , then  $x_t$  satisfies the constraint (10).

*Proof.* I., II. The first two parts of the proof are almost verbatim copies of the corresponding parts of Theorem 1 and therefore omitted.

III. Let  $\tau$  be the last time before  $t$  such that  $\beta_\tau = 1$  and zero if there is no such time. For  $\tau = t$ , the result follows trivially from the definition of  $\beta_t$  (15). Next, for  $\tau = 0$  the result follows since then there are no constraint violations up to time  $t$ . Finally, for  $\tau \in \mathbb{N}_1^{t-1}$ , we have from (15) that

$$v_\tau \leq \alpha\tau + \alpha - 1,$$

and, since there are no violations from  $\tau+2$  to  $t$ , we get

$$v_t = v_{\tau+1} \leq v_\tau + 1 \leq \alpha\tau + \alpha = \alpha(\tau+1) \leq \alpha t,$$

which is equivalent to (10).  $\square$

## IV. IMPLEMENTATION

In this section we discuss how the general theory developed in previous sections can be employed within an MPC framework. For the most part, the procedure is identical for both constraints considered, (7) and (10). Therefore, we focus on stochastic constraint (7), which is our primary concern, and outline modifications for the robust constraint (10) whenever necessary.

The set valued control policy (24) (or (28)) gives rise to a generic MPC problem

**Problem 1.**

$$\text{minimize } \{J \mid u_t \in \mathcal{U}_t(x_t, v_t)\},$$

where the cost function  $J$  is completely free to choose, as well as is the prediction horizon and the policy parametrization with respect to which the cost function  $J$  is minimized.

For each value of  $t$ , the constraint set of Problem 1 is a polyhedron (by Assumption 2), and hence if  $J$  is convex in the decision variables of the problem (i.e., the policy parametrization), then Problem 1 is convex. Computational aspects of this problem are discussed in the remainder of this section.

*A. Parametrization of SRCI sets*

A crucial step for the application of Problem 1 is the parametrization of a family of SRCI sets  $\mathcal{S}_k$  (20). One way of doing so is an explicit construction of an SRCI set  $\mathcal{S}$  (Definition 2) and a subsequent application of the  $\text{Reach}(\cdot)$  operator (Eq. (19) and Remark 5). It is desirable that the set  $\mathcal{S}$  (and hence all  $\mathcal{S}_k$ ) be large, preferably *maximal*. However, the computation of maximum (stochastic) (robust) controlled invariant sets is known to be difficult in larger dimensions, as polyhedral projections are required, and the maximal set may not be polyhedral [1, 2]. Nevertheless, there are effective algorithms for the computation of low-complexity *polyhedral* controlled invariant under-approximations of these sets; see, e.g., [3, 4, 21].

The computation of maximum controlled invariant sets (or large under-approximations thereof) can be avoided if a family of SRCI sets  $\mathcal{S}_k$  is parametrized implicitly. The implicit inclusion to a family of SRCI sets can be achieved via the traditional MPC with a terminal invariant set [16, 22]; the procedure is now briefly described.

For the traditional MPC scheme, at a time  $t$  and on a prediction horizon  $N$ , the control input predictions  $u_{t+k}$  are given by an explicit policy parametrization (the decision variable) for  $k \in \mathbb{N}_0^{N-1}$ , and by a fixed terminal controller for  $k \geq N$ . Let the explicit policy (in general a causal state-sequence feedback or, equivalently, causal disturbance feedback) be  $\pi := (\pi_0, \dots, \pi_{N-1})$ , that is,

$$u_{t+k} = \pi_k(x_t^{t+k}), \quad k \in \mathbb{N}_0^{N-1},$$

and let the terminal state-feedback controller be  $\kappa_f$ , that is,

$$u_{t+k} = \kappa_f(x_{t+k}), \quad k \geq N.$$

The constraint satisfaction is enforced explicitly along the prediction horizon through constraints on the policy parametrization  $\pi$  and implicitly beyond the prediction horizon by constraining the terminal state  $x_{t+N}$  to a positively invariant set associated with the terminal controller  $\kappa_f$ . Specifically, the terminal set  $\mathcal{X}_{\kappa_f}$  employed is a subset of the stochastic feasibility set  $\mathcal{X}_s$  such that

$$Ax + B\kappa_f(x) + w \in \mathcal{X}_{\kappa_f} \quad \forall w \in \mathcal{W}, \quad (29a)$$

$$Ax + B\kappa_f(x) + w \in \mathcal{X} \quad \forall w \in \mathcal{R}_\alpha, \quad (29b)$$

$$\kappa_f(x) \in \mathcal{U} \quad (29c)$$

is satisfied for all  $x \in \mathcal{X}_{\kappa_f}$ .

At a time  $t$ , given  $x_t$  and  $v_t$  (and hence  $r_t$  and  $\beta_t$ ), we wish to ensure that  $u_t \in \mathcal{U}_t(x_t, v_t)$ . The probabilistic constraint (22b) remains unchanged. The invariance constraint (22a) is enforced implicitly as follows. First, by

definition of  $\mathcal{S}_k$  and  $\mathcal{S}$ , the inclusion  $x_{t+1} \in \mathcal{S}_{r_t}$  for all  $w \in \mathcal{W}$  is ensured if

$$P(x_{t+r_t+i} \in \mathcal{X} \mid x_{t+r_t+i-1}, x_t, v_t) \geq 1 - \alpha \quad (30)$$

is satisfied for a given  $(x_t, v_t)$ , all  $i \in \mathbb{N}_+$  and all possible  $x_{t+r_t+i-1}$  generated by all possible  $w_t^{t+r_t+i-2} \in \mathcal{W}^{r_t+i-1}$  under the given policy  $\pi$  and the terminal controller  $\kappa_f$ . The constraint (30) is then enforced explicitly for  $i \in \mathbb{N}_1^{N-r_t}$  by constraints on  $\pi$  and implicitly for  $i > N - r_t$  by requiring that  $x_{t+N} \in \mathcal{X}_{\kappa_f}$  for all  $w_t^{t+N-1}$  under  $\pi$ .

In principle there are no restrictions on the policy parametrization  $\pi$  and the terminal controller  $\kappa_f$  as long as they are “compatible” in the sense that the shifted solution is feasible at the next time instant. In fact, most of the widely used robust and stochastic MPC parametrizations such as affine disturbance feedback [12, 14], pre-stabilization [15] or tubes [17] can be used.

For the robust constraint (10), the constraint (30) is replaced by a robust inclusion of  $x_{t+r_t+i}$  to  $\mathcal{X}$ , and  $\mathcal{R}_\alpha$  is replaced by  $\mathcal{W}$  in (29b) when constructing a terminal region.

*B. Enforcing the constraints*

In this section, we briefly outline how to enforce the constraint  $u_t \in \mathcal{U}_t(x_t, v_t)$  of Problem 1. In the case of implicit policy parametrization, the specific form of the constraint (30) depends on the policy parametrization and is not discussed here; see [14] for a detailed treatment with affine disturbance feedback in a similar setting. Hence, in the remainder of this section we assume that a family of SRCI sets  $\mathcal{S}_k$  has been characterized explicitly.

From a computational viewpoint we are faced with a single type of constraint – the robust inclusion of the successor state to a given polyhedron. Invariance constraints (e.g., (22a)) and one-step probabilistic and robust constraints (e.g., (22b) and (26b)) translate to this type of constraints.

The robust inclusions translate immediately to affine constraints. Indeed, given a polyhedron  $\mathcal{P} = \{x \mid a_j^T x \leq b_j\}$  and a set  $\mathcal{M}$ , we have

$$\begin{aligned} Ax + Bu + w \in \mathcal{P} \quad \forall w \in \mathcal{M} &\Leftrightarrow \\ &\Leftrightarrow a_j^T (Ax + Bu) \leq b_j - \max_{w \in \mathcal{M}} a_j^T w, \end{aligned}$$

where  $\max_{w \in \mathcal{M}} a_j^T w$  are fixed finite numbers for any compact set  $\mathcal{M}$ . In particular, these numbers can be found by solving a linear program for a polyhedral  $\mathcal{M}$  or found analytically for an ellipsoidal  $\mathcal{M}$  (details are omitted for brevity).

*C. Computational complexity discussion*

This section briefly compares the computational complexity of the presented approaches with existing MPC formulations.

If the SRCI sets are parametrized explicitly, then the *offline* complexity is governed by computation of maximum controlled invariant sets (or under-approximations thereof) and as such is comparable to the nominal first-step MPC of [9, 10]. The *online* computational requirements are then governed by the type of the cost function  $J$ , policy parametrization and, to a smaller extent, by the complexity of the SRCI sets employed.

If the SRCI sets are parametrized implicitly, then the *offline* computational complexity is determined by the policy parametrization  $\pi$  and the terminal controller  $\kappa_f$ , which are also the main factors in *online* complexity along with the type of the cost function  $J$ . Most importantly, online as well as offline computational requirements are completely analogous

to those of the traditional robust and stochastic MPC schemes with the same parametrization, be it affine disturbance feedback [12, 14], prestabilization [15] or tubes [17].

It should be stressed that with explicitly parametrized SRCI sets as well as with implicit parametrization using any of the above-mentioned policies, the constraint set of Problem 1 is *polyhedral* and hence the class of the problem (e.g., quadratic / linear program) is not altered by introducing either of the constraints (7) and (10).

## V. NUMERICAL EXAMPLE

We consider the system (1) given by the matrices  $A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$ ,  $B = \begin{bmatrix} 1 \\ 0.7 \end{bmatrix}$ . The i.i.d. disturbance sequence has a componentwise-uncorrelated Gaussian distribution with each component truncated at 3, i.e.,  $\mathcal{W} = \{w \in \mathbb{R}^n \mid \|w\|_\infty \leq 3\}$ . The input constraints are  $|u| \leq 12$ . The state is constrained to lie (most of the time) in a rectangular region given by  $|x_1| \leq 7$ ,  $x_2 \geq 0$  and  $x_2 \leq 12$ . The allowed violation level  $\alpha$  is 0.2. The cost function  $J$  is the standard expectation of a quadratic function given by the weighting matrices  $Q = \text{diag}(0, 1)$  and  $R = 0.1$ . The prediction horizon is  $N = 8$ . As a policy parametrization we chose the affine disturbance feedback (see, e.g., [14]), which allows for the exact evaluation of  $J$ . A family of SRCI sets  $\mathcal{S}_k$  with six layers ( $n_s = 6$ ) was parametrized explicitly, starting with the maximum SRCI set  $\mathcal{S}_1 = \mathcal{S}$ . The confidence regions  $\mathcal{R}_\alpha$  and  $\mathcal{R}_{\beta_t}$  were chosen as scaled symmetric boxes around the origin.

We compared the performance of the proposed approaches using both the stochastic constraint (7) (Av-Exp) and the robust constraint (10) (Av-Rob) against the first-step approach to the one-step conditional constraint (5) (One-Step) proposed in [14], and against the standard affine disturbance feedback robust MPC (Robust) of, e.g., [12]. As a benchmark we used the unconstrained infinite-horizon LQ optimal controller; this controller outperforms any other controller and was included only to get insight into the absolute performance of the constrained approaches. Rather than performing a Monte Carlo analysis, we compared the policies over one long realization ( $T = 1000$ ). The performance in terms of the cost function and the number of constraint violations at the final time is summarized in Table I. The stochastic version, Av-Exp, of the presented approach outperforms the other policies by a significant margin, and, as expected, the average number of violations is almost on the 20% boundary. Also the robust version, Av-Rob, outperforms the standard robust MPC by a large margin. Note in particular that the number of violations with Av-Rob is also almost tight; however, the benefit gained by violating the constraint is partly negated by the need to return well into the interior of the constraint set in order to satisfy the constraint robustly in the future. This is why the one-step stochastic MPC performs better than Av-Rob despite fewer constraint violations. These observations are only confirmed by inspecting the trajectories of  $x_2$  shown in Figure 2 along with the trajectories of  $\beta_t$ ,  $v_t/t$  and  $r_t$ . Indeed, the trajectory of  $x_2$  for Av-Rob has significantly higher variance than for the other policies.

## VI. CONCLUSION

This paper presented a framework to handle the constraints on the number of state-constraint violations averaged over time, where the average number of violations can be bounded either in a stochastic manner or robustly. The key ingredient of our approach is the explicit incorporation of the information

on the past number of state-constraint violations into the decision on the current control input. The approach significantly reduces the conservatism of previous stochastic MPC formulations, which is confirmed by a numerical example. The computational requirements (both online and offline) are comparable to those of the recursively feasible nominal MPC with a first-step constraint of [9, 10] in the case of explicit parametrization of controlled invariant sets, or comparable to those of robust MPC in the case of implicit parametrization (assuming the same parametrization for the presented and for the robust approach).

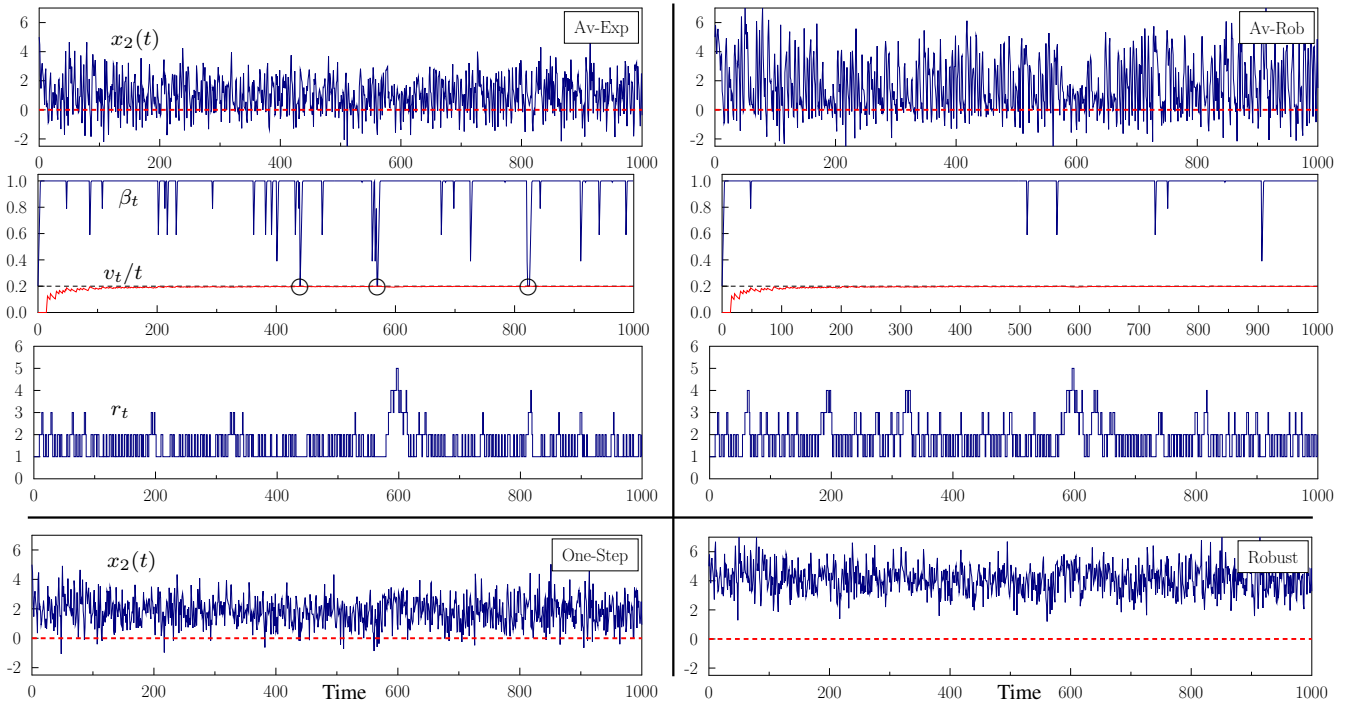
Future work will focus on extending the presented methodology to output feedback, time-varying systems and to different types of constraints, e.g., integrated chance constraints [19], spectrum constraints [11, 13], time-varying constraints and weighted-average constraints. Another direction is optimization of closed-loop performance under the constraints considered here; for instance, online optimization of the confidence regions and prediction strategies that take into account the nested structure are topics worth investigating.

## REFERENCES

- [1] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.
- [2] F. Blanchini. *Set-Theoretic Methods in Control*. Birkhäuser Boston, first edition, 2007.
- [3] F. Blanchini, S. Miani, and C. Savorgnan. Dynamic augmentation and complexity reduction of set-based constrained control. In *17th IFAC World Congress, Seoul, Korea*, pages 14324–14329, 2008.
- [4] T. B. Blanco, M. Cannon, and B. D. Moor. On efficient computation of low-complexity controlled invariant sets for uncertain linear systems. *Int. Journal of Control*, 83(7):1339–1346, 2010.
- [5] M. Cannon, B. Kouvaritakis, and D. Ng. Probabilistic tubes in linear stochastic model predictive control. *Systems & Control Letters*, 58(10-11):747–753, 2009.
- [6] M. Cannon, B. Kouvaritakis, S. V. Raković, and Q. Cheng. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Transactions on Automatic Control*, 56(1):194–200, 2011.
- [7] M. Cannon, B. Kouvaritakis, and X. Wu. Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, 45(1):167–172, 2009.
- [8] M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic Constrained MPC for Multiplicative and Additive Stochastic Uncertainty. *IEEE Transactions on Automatic Control*, 54(7):1626–1632, 2009.
- [9] R. Gondhalekar and J. Imura. Least-restrictive move-blocking model predictive control. *Automatica*, 46(7):1234–1240, 2010.
- [10] R. Gondhalekar, J. Imura, and K. Kashima. Controlled invariant feasibility – A general approach to enforcing strong feasibility in MPC applied to move-blocking. *Automatica*, 45(12):2869–2875, 2009.
- [11] R. Gondhalekar, C. N. Jones, T. Besselmann, J. H. Hours, and M. Mercangöz. Constrained spectrum control using MPC. In *50th IEEE Conference on Decision and Control and European Control Conference*, pages 1219–1226, 2011.
- [12] P. J. Goulart, E. C. Kerrigan, and J. M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523–533, 2006.
- [13] J. H. Hours, M. N. Zeilinger, R. Gondhalekar, and C. N. Jones. Spectrogram MPC: Enforcing hard constraints on system’s output spectra. In *American Control Conference*, pages 2010–2017, 2012.
- [14] M. Korda, R. Gondhalekar, J. Cigler, and F. Oldewurtel. Strongly feasible stochastic model predictive control. In *50th IEEE Conference on Decision and Control and European Control Conference*, pages 7381–7386, 2011.
- [15] B. Kouvaritakis, M. Cannon, S. V. Raković, and Q. Cheng. Explicit use of probabilistic distributions in linear predictive control. *Automatica*, 46(10):1719–1724, 2010.

**TABLE I:** Comparison of the control policies over the simulation time  $T = 1000$ . First row: cost increase over the unconstrained LQ optimal controller. Second row: number of constraint violations averaged over time at the end of the simulation. The allowed level of violation is  $\alpha = 0.2$ .

policy	LQ	Av-Exp	One-Step	Av-Rob	Robust
$J/J_{LQ}$	1	2.45	3.64	5.64	13.77
$v_T/T$	0.494	0.199	0.027	0.192	0



**Fig. 2:** Sample paths of  $x_2$ ,  $\beta_t$ ,  $v_t/t$  and  $r_t$  for Av-Exp and Av-Rob, and sample paths of  $x_2$  for One-Step and Robust policies. Note that with the stochastic constraint (top left) there are three time instants (circled) with  $v_t/t > \alpha$ , whereas  $v_t/t \leq \alpha$  for all  $t$  with the robust averaged constraint (top right). This reflects the difference between the two specifications: the robust constraint requires  $v_t/t \leq \alpha$  for each realization of the disturbance process, whereas the stochastic constraint allows for occasional violations of  $v_t/t \leq \alpha$ , thereby improving the cost performance.

- [16] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- [17] D. Q. Mayne, M. M. Seron, and S. V. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.
- [18] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:15–27, 2012.
- [19] A. Prékopa. *Stochastic Programming*. Springer, first edition, 1995.
- [20] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11:733–764, 2003.
- [21] S. V. Raković, E. C. Kerrigan, D. Q. Mayne, and K. I. Kouramas. Optimized robust control invariance for linear discrete-time systems: Theoretical foundations. *Automatica*, 43(5):831–841, 2007.
- [22] J. B. Rawlings and D. Q. Mayne. *Model Predictive Control: Theory and Design*. Nob Hill Publishing, first edition, 2009.