Recursive Feasibility of Stochastic Model Predictive Control with Mission-Wide Probabilistic Constraints

Kai Wang and Sébastien Gros

Abstract— **This paper is concerned with solving chance-constrained finite-horizon optimal control problems, with a particular focus on the recursive feasibility issue of stochastic model predictive control (SMPC) in terms of mission-wide probability of safety (MWPS). MWPS assesses the probability that the entire state trajectory lies within the constraint set, and the objective of the SMPC controller is to ensure that it is no less than a threshold value. This differs from classic SMPC where the probability that the state lies in the constraint set is enforced independently at each time instant. Unlike robust MPC, where strict recursive feasibility is satisfied by assuming that the uncertainty is supported by a compact set, the proposed concept of recursive feasibility for MWPS is based on the notion of remaining MWPSs, which is conserved in the expected value sense. We demonstrate the idea of mission-wide SMPC in the linear SMPC case by deploying a scenario-based algorithm.**

I. Introduction

Model predictive control (MPC) has been well established for dealing with complex constrained optimal control problems [1], [2]. In the context of MPC, the system dynamics are required to be known and deterministic. In practice, the system uncertainty, including imprecise model parameters and process noises, is generally unavoidable. Because MPC does not take the uncertainty into account, constraints violations can occur.

For applications where safety is critical, robust MPC (RMPC) strategies have been proposed to explicitly account for uncertainty. However, RMPC schemes can only handle bounded disturbances and the resulting control strategy can be conservative. To overcome these limitations, stochastic MPC (SMPC) methods have been developed to seek a trade-off between control performance and the risk of constraint violations using chance constraints.

There are mainly three forms of chance constraints proposed in the literature: individual chance constraint, stage-wise chance constraint [3] and mission-wide chance constraint [4]. E.g. in path planning for vehicles in the presence of obstacles, individual or stage-wise constraints restrict at every time instant the probability that the vehicle collides with an obstacle. In contrast, a mission-wide chance constraint directly restricts the probability of collision on the overall driving mission. A mission-wide chance constraint is arguably more meaningful than stage-wise constraints. Indeed, the former directly handles the risk of running a mission [5], [4], while the latter does it very indirectly. However, stage-wise chance constraints are easier to handle than mission-wide constraints. Indeed, the latter handles probabilities over entire state trajectories, yielding very large probability spaces. More forms of chance constraints are discussed in, e.g., [3, Section 2.2] and references therein.

The current research on SMPC focuses on developping efficient methods for solving the underlying optimization problem, while recursive feasibility is less explored. Indeed, because of the (possibly unbounded) stochasticity, the recursive feasibility of SMPC typically holds in the probabilistic sense, making its analysis much more involved. Some results exist in specific contexts. When the system uncertainties have finite supports, recursive feasibility can be guaranteed using robust MPC [6], at the cost of yielding very conservative control policies. For linear stochastic systems with infinite support, if the first two moments of the disturbance distribution are known, constraint-tightening methods via the Chebyshev–Cantelli inequality are presented in [7], [8], [9]. Recursive feasibility is guaranteed using backup strategies when an infeasible optimization problem is encountered [7], [8], and using time-varying risk bound [9]. The author in [10], [11] proposed SMPC algorithms that have a certain probability of remaining feasible if the initial condition is feasible. However, none of these methods tackle mission-wide probability of safety (MWPS), nor can provide a meaningful certificate of MWPS. In [12], the problem of maximizing the MWPS is expressed as a stochastic invariance problem and further developed into an optimal control problem, which is solvable via dynamic programming. Unfortunately, problems constraining the MWPS rather than maximizing it cannot necessarily be put in that simple form.

Guaranteeing recursive feasibility of a SMPC problem with MWPS constraints is an open problem, and this paper investigates a tentative solution. The main contributions of this paper is threefold. First, we show that if a policy is designed to achieve a certain MWPS, then the MWPS remaining until the end of the mission remains constant in the expected value sense. Second, we design a recursively feasible control scheme using shrinking horizon policies in the context of SMPC with MWPS guarantee. The proposed scheme treats directly the probability of running a mission successfully and therefore does not introduce artificial conservativeness. Third, an efficient scenario-based algorithm is proposed to deploy the idea in the linear case.

The paper is structured as follows. In Section II we present the problem statement of SMPC with MWPS constraints, and its difference from the classical SMPC with stage-wise probabilistic constraint. Section III details how the MWPS remains constant throughout the mission, and a recursively

All authors are with Department of Engineering Cybernetics, Norwegian University of Science and Technology (NTNU), 7491, Trondheim, Norway. {kai.wang, sebastien.gros}@ntnu.no

feasible policy design is discussed in Section IV. We demonstrate the idea in the linear SMPC case based on an efficient scenario-based algorithm in Section V. Finally, Section VI concludes the paper points to some future work.

Notation. Boldface $\mathbf{a}$ (italic a) is a vector (scalar), and $\mathbf{a}_{0,\ldots,N}$ ($a_{0,\ldots,N}$) denotes a sequence of vectors (scalars). We use $\mathbf{s}_{0,\ldots,N} \in \mathbb{S}$ to denote that a state sequence $\mathbf{s}_{0,\ldots,N}$ lies in a constraint set $\mathbb{S}$ of the state space, i.e., $\mathbf{s}_k \in \mathbb{S}$ for all $k = 0, \ldots, N$. We denote $\mathbb{I}_{[a,b]}$ the set of integers in the interval $[a, b] \subseteq \mathbb{R}$.

II. Problem Statement

We consider a mission spanning a predefined horizon $N \in \mathbb{N}$ to be "safe" if:

$$\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \tag{1}$$

starting from some initial states $\mathbf{s}_0 \in \mathbb{S}$. Here, $\mathbf{s}_k \in \mathbb{R}^n$ is the state at time step k, and $\mathbb{S} \subset \mathbb{R}^n$ is a set in the state space. We assume that the true stochastic system dynamics are given by:

$$\rho[\, \mathbf{s}_+ \,|\, \mathbf{s}, \mathbf{u} \,] \tag{2}$$

providing the probability density underlying transitions from a state-input pair $\mathbf{s}, \mathbf{u}$ to a new state $\mathbf{s}_+$. Throughout the paper, we assume that the states are known and continuous. Notice that the control community typically uses:

$$\mathbf{s}_+ = \mathbf{f}(\mathbf{s}, \mathbf{u}, \mathbf{w}) \tag{3}$$

to describe stochastic dynamics, in which $\mathbf{w}$ denotes the stochastic disturbances and $\mathbf{f}$ is generally a nonlinear function. The input $\mathbf{u}$ is given by a control policy sequence

$$\boldsymbol{\pi} := \{\boldsymbol{\pi}_0, \ldots, \boldsymbol{\pi}_{N-1}\}$$

such that

$$\mathbf{u}_k = \boldsymbol{\pi}_k(\mathbf{s}_k), \;\; \forall k \in \mathbb{I}_{[0,N-1]}.$$

In general, guaranteeing the absolute safety as described in (1) yields very conservative control policies, or is even infeasible if the uncertainty is unbounded. Alternatively, for a given initial condition $\mathbf{s}_0 \in \mathbb{S}$ and a policy sequence $\boldsymbol{\pi}$, we are interested in the Mission-Wide Probability of Safety (MWPS):

$$\mathbb{P}[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}]. \tag{4}$$

The problem we are interested in is then to find a policy sequence solution of

$$\min_{\boldsymbol{\pi}} \quad \mathbb{E}\left[M(\mathbf{s}_N) + \sum_{k=0}^{N-1} L(\mathbf{s}_k, \boldsymbol{\pi}_k(\mathbf{s}_k)) \right] \tag{5a}$$

$$\text{s.t.} \quad \mathbb{P}[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}] \geq S\,, \tag{5b}$$

where $S \in [0, 1]$ is a predefined safety bound, the functions L and M are some given stage and terminal costs, and (5a) is the expectation over the state trajectories resulting from $\mathbf{s}_0, \boldsymbol{\pi}$ and (2).

In practice, calculating an optimal policy sequence for problem (5) exactly is hardly possible, because it involves optimization over an infinite dimensional function space. To tackle this issue, we will be interested in using stochastic MPC formulations to generate policies that enforce the MWPS (4), where a control input $\mathbf{u}_k$ is computed by solving an optimal control problem at every time step. In that context, a key concept will be the remaining MWPS at any time $k \in \mathbb{I}_{[1,N-1]}$ for a given state $\mathbf{s}_k$, defined as:

$$\mathbb{P}[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_k, \boldsymbol{\pi}] \tag{6}$$

Here $\mathbf{s}_k$ is the outcome of a realization $\mathbf{s}_{1,\ldots,k}$ of the Markov Chain and its relationship to (4).

A key observation is that a policy sequence $\boldsymbol{\pi}$ satisfying (5b) does not yield any guarantee on (6). Indeed, an adversarial realization can e.g. bring the system into a state $\mathbf{s}_k$ for which the remaining MWPS is lower than S.

This observation entails that in the proposed context, the notion of recursive feasibility needs to be treated in a different way that is commonly done in robust MPC. We will detail this in Section III. We briefly detail next the motivation for developing methods to treat MWPS constraints outside of the classical chance-constraint framework.

A. Mission-wide constraints versus stage-wise constraints

In this section, we present our motivation for treating MWPS directly rather than via Stage-Wise Probability of Safety (SWPS). In particular, regardless of the desired MWPS level, enforcing it via SWPS becomes very conservative for long missions. SWPS problems seek policies that enforce constraints of the form:

$$\mathbb{P}\left[\mathbf{s}_k \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}\right] \geq s_k \geq s, \quad \forall\, k \in \mathbb{I}_{[1,N]}\,, \tag{7}$$

in which $1 \geq s_k \geq s \geq 0$. One can then easily verify that the Boolean algebra and Booles's inequality entail that:

$$\begin{aligned}
&\mathbb{P}[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}] = 1 - \mathbb{P}\left[\bigcup_{k=1}^{N} \mathbf{s}_k \notin \mathbb{S} \,\middle|\, \mathbf{s}_0, \boldsymbol{\pi} \right] \\
&\geq 1 - \sum_{k=1}^{N} \mathbb{P}\left[\mathbf{s}_k \notin \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}\right] \\
&= 1 - \sum_{k=1}^{N} \left(1 - \mathbb{P}\left[\mathbf{s}_k \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}\right]\right) \\
&\geq 1 - N + \sum_{k=1}^{N} s_k \geq 1 - N + Ns\,.
\end{aligned}$$

To ensure the satisfaction of (5b) via imposing (7), requires the choice, $1-N+Ns \geq S$, i.e. a bound for s can be derived as:

$$s \geq \frac{N-1}{N} + \frac{S}{N}\,. \tag{8}$$

Hence enforcing MWPS (4) via SWPS (7) requires selecting s according to (8), which yields a bound s close to one very fast as N increases, see Fig. 1, hence turning the SWPS constraints into hard constraints. While tighter bounds than (8) can be derived[1], treating MWPS via SWPS without introducing conservativeness is difficult. The intuitive reason

[1] e.g. the Bonferroni inequalities allow one to refine the bound in (8) by accounting for some of the correlation between successive states

behind this issue is that SWPS formulations neglect the time-correlation between the constraints violations, and as a result, it offers an incorrect representation of the risks incurred by a system over a mission. Bound (8) corrects that, at the cost of introducing a high conservatism. By using a similar argument to what we developed above, risk-allocation technology proposed in [13] optimizes the risk assigned to each stage-wise constraint instead of using constant risk in (7). This method leads a computationally expensive two-stage optimization problem and is still conservative as depicted in [3, Fig. 1].

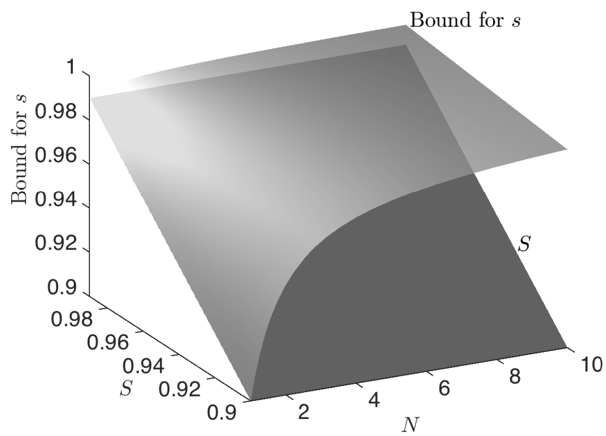

Fig. 1. Illustration of bound (8) for various N and S. The curved manifold displays the bound (8) for the SWPS such that a prescribed MWPS (4) holds.

III. Relation between Remaining MWPS and Initial MWPS

Here, we show that the remaining MWPS is constant in the expected value sense. This offers a novel path for guaranteeing the recursive feasibility of MPC-like control schemes with MWPS constraints.

Lemma 1 *If the policy sequence* $\boldsymbol{\pi}$ *satisfies* (4)*, then*

$$\mathbb{E}_{\{\mathbf{s}_{1,\ldots,k}\in\mathbb{S}\,|\,\mathbf{s}_0,\boldsymbol{\pi}\}}\left[\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\,\mathbf{s}_k,\boldsymbol{\pi}]\right]\geq S \tag{9}$$

for all $k=1,\ldots,N-1$*, i.e. the remaining MWPS at time* k *satisfies the constraint on the prescribed MWPS* (4) *in the expected value sense.*

Proof. We observe that:

$$\begin{aligned}
&\mathbb{P}[\mathbf{s}_{1,\ldots,N}\in\mathbb{S}\,|\,\mathbf{s}_0,\boldsymbol{\pi}]\\
&=\int_{\mathbb{S}^k}\mathbb{P}[\mathbf{s}_{1,\ldots,k}|\mathbf{s}_0,\boldsymbol{\pi}]\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\mathbf{s}_{0,\ldots,k},\boldsymbol{\pi}]\mathrm{d}\mathbf{s}_1\ldots\mathrm{d}\mathbf{s}_k\\
&=\int_{\mathbb{S}^k}\mathbb{P}[\mathbf{s}_{1,\ldots,k}\,|\,\mathbf{s}_0,\boldsymbol{\pi}]\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\,\mathbf{s}_k,\boldsymbol{\pi}]\,\mathrm{d}\mathbf{s}_1\ldots\mathrm{d}\mathbf{s}_k\\
&=\int_{\mathbb{S}}\mathbb{P}[\mathbf{s}_{1,\ldots,k-1}\in\mathbb{S}\wedge\mathbf{s}_k|\mathbf{s}_0,\boldsymbol{\pi}]\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\mathbf{s}_k,\boldsymbol{\pi}]\mathrm{d}\mathbf{s}_k\\
&:=\mathbb{E}_{\{\mathbf{s}_{1,\ldots,k}\in\mathbb{S}\,|\,\mathbf{s}_0,\boldsymbol{\pi}\}}\left[\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\,\mathbf{s}_k,\boldsymbol{\pi}]\right].
\end{aligned}$$

Here, $\mathbf{s}_{1,\ldots,k}$ is a Markov Chain underlying (2), and therefore a random variable in the high dimensional space $(\mathbb{R}^n)^k$. Given a policy sequence $\boldsymbol{\pi}$, the remaining MWPS at time k, i.e., the term inside $\mathbb{E}_{\{\mathbf{s}_{1,\ldots,k}\in\mathbb{S}\,|\,\mathbf{s}_0,\boldsymbol{\pi}\}}[\,\cdot\,]$, depends only on the random state $\mathbf{s}_k$, which is a particular dimension in the Markov Chain $\mathbf{s}_{1,\ldots,k}$. Hence, the last equation holds because here $\mathbb{E}_{\{\mathbf{s}_{1,\ldots,k}\in\mathbb{S}\,|\,\mathbf{s}_0,\boldsymbol{\pi}\}}[\,\cdot\,]$ is used to denote the expectation value of the remaining MWPS that is taken over all possible realizations of the random Markov Chain $\mathbf{s}_{1,\ldots,k}$ that remains in $\mathbb{S}$. ■

Lemma 1 entails that the MWPS is conserved in the expected value sense throughout the mission if a mission-wide policy sequence has been selected at the beginning of the mission. Result (9) is arguably best interpreted in a frequentist framework. Indeed, even though a specific realization $\mathbf{s}_{0,\ldots,k}$ may be adversarial for the remaining MWPS, we observe that in average the MWPS remains unchanged throughout the mission. As a result, (9) entails that if running missions under policy $\boldsymbol{\pi}$ designed according to (5), the resulting ratio of success will asymptotically be at least S. While this statement may appear tautological, it provides a basic concept of recursive feasibility that can be translated into constraints in a MPC framework to ensure that a prescribed MWPS is achieved. We detail this observation below.

IV. Recursive feasibility of MWPS with shrinking-horizon policies

In this section, we focus on solving the originally proposed mission-wide probability-constrained finite-horizon optimal control problem (5) using shrinking-horizon policies that are updated as the mission progresses. The reason behind this is that the exact optimal policies for (5) is difficult to compute in general.

As a result, in practice, the policy sequence $\boldsymbol{\pi}$ is typically finitely parameterized, and hence restricted to a subset of the set of admissible policies. This introduces sub-optimality, and makes it useful to re-solve problem (5) at every time instant k, according to the latest state realization $\mathbf{s}_k$. We then consider at every time k the control policy sequence:

$$\boldsymbol{\pi}^k=\left\{\boldsymbol{\pi}^k_k,\ldots,\boldsymbol{\pi}^k_{N-1}\right\} \tag{10}$$

lasting to the end of the mission. For the sake of brevity, we will work with a shrinking horizon extending to the end of the mission. The fixed, receding horizon shorter than the mission duration will be the object of our future work.

At every time instant $k\in\mathbb{I}_{[0,N-1]}$, for the corresponding state $\mathbf{s}_k$, we consider solving the following shrinking-horizon, mission-wide and chance-constrained problem:

$$\min_{\boldsymbol{\pi}^k}\quad\mathbb{E}\left[M(\mathbf{s}_N)+\sum_{l=k}^{N-1}L(\mathbf{s}_l,\boldsymbol{\pi}^k_l(\mathbf{s}_l))\right] \tag{11a}$$

$$\text{s.t.}\quad\mathbb{P}[\mathbf{s}_{k+1,\ldots,N}\in\mathbb{S}\,|\,\mathbf{s}_k,\boldsymbol{\pi}^k]\geq S_k \tag{11b}$$

to get a new policy sequence. Here, $S_k\in[0,1]$ is a varying risk-bound that will be specified later. Notice that while (11) computes an entire policy sequence $\boldsymbol{\pi}^k$ for the current state $\mathbf{s}_k$, only the first policy $\boldsymbol{\pi}^k_k$ of that sequence is used to generate the actual control action, as the policy sequence is recalculated at the next time instant $k+1$, in a classic MPC fashion. The inputs eventually applied to the closed-loop system will therefore read as:

$$\mathbf{u}_k = \boldsymbol{\pi}_k^k(\mathbf{s}_k)\ , \quad \forall\, k \in \mathbb{I}_{[0,N-1]} \tag{12}$$

In the context of mission-wide SMPC, we will consider the recursive feasibility issue of employing the policy sequence $\{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{N-1}^{N-1}\}$ resulting from extracting only the first policy $\boldsymbol{\pi}_k^k$ of the policy sequence $\boldsymbol{\pi}^k$ at every time step k, for all $k \in \mathbb{I}_{[0,N-1]}$. We show next that retaining recursive feasibility in the sense of (9) requires only that the new policy sequence produces a remaining MWPS that is not worse than a discounted one achieved by the previous policy for the current state $\mathbf{s}_k$. We formalise this statement in the proposition below.

Proposition 1 *Assume that the initial policy sequence $\boldsymbol{\pi}^0$ satisfies the MWPS cosntraint:*

$$\mathbb{P}[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}^0] \geq S_0 \geq S \tag{13}$$

and that each policy sequence $\boldsymbol{\pi}^k$ is built under the constraint:

$$\mathbb{P}[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_k, \boldsymbol{\pi}^k] \geq S_k \tag{14}$$

where

$$S_k = \gamma_k \mathbb{P}[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_k, \boldsymbol{\pi}^{k-1}]$$

holds and with $\gamma_k \in (0,1]$, for all $k \in \mathbb{I}_{[1,N-1]}$. Then the MWPS under $\mathbf{u}_k = \boldsymbol{\pi}_k^k(\mathbf{s}_k)$ and $k \in \mathbb{I}_{[0,N-1]}$ reads as:

$$\mathbb{P}\left[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,\middle|\, \mathbf{s}_0, \{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{N-1}^{N-1}\}\right] \geq \prod_{k=1}^{N-1} \gamma_k S_0\,. \tag{15}$$

Proof. We will prove this by induction. Consider

$$\begin{aligned}
&\mathbb{P}\left[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,\middle|\, \mathbf{s}_0, \{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_k^k, \ldots, \boldsymbol{\pi}_{N-1}^k\}\right] \\
&= \int_{\mathbb{S}} \mathbb{P}\left[\mathbf{s}_{1,\ldots,k-1} \in \mathbb{S} \wedge \mathbf{s}_k \,|\, \mathbf{s}_0, \{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{k-1}^{k-1}\}\right] \\
&\qquad\qquad \cdot \mathbb{P}\left[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,\middle|\, \mathbf{s}_k, \boldsymbol{\pi}^k\right] \,\mathrm{d}\mathbf{s}_k \\
&\geq \int_{\mathbb{S}} \mathbb{P}\left[\mathbf{s}_{1,\ldots,k-1} \in \mathbb{S} \wedge \mathbf{s}_k \,|\, \mathbf{s}_0, \{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{k-1}^{k-1}\}\right] \\
&\qquad\qquad \cdot \gamma_k \mathbb{P}\left[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,\middle|\, \mathbf{s}_k, \boldsymbol{\pi}^{k-1}\right] \mathrm{d}\mathbf{s}_k \\
&= \gamma_k \mathbb{P}\left[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,\middle|\, \mathbf{s}_0, \{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{k-1}^{k-1}, \ldots, \boldsymbol{\pi}_{N-1}^{k-1}\}\right]\,,
\end{aligned}$$

where the last equality holds because the last integral describes the MWPS associated to applying the policy sequence $\{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{k-1}^{k-1}, \ldots, \boldsymbol{\pi}_{N-1}^{k-1}\}$. Hence an induction from

$$\mathbb{P}[\mathbf{s}_{1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_0, \boldsymbol{\pi}^0] \geq S_0$$

yields (15). ■

Let us introduce the following corollaries, showing the practical implications of Proposition 1:

Corollary 1 ***Guarantee of MWPS:*** *The choice:*

$$\prod_{k=1}^{N-1} \gamma_k S_0 = S \tag{16}$$

together with the policy update constraint (14) *yields a sequence of policies $\{\boldsymbol{\pi}_0^0, \ldots, \boldsymbol{\pi}_{N-1}^{N-1}\}$ that satisfies the prescribed MWPS constraint* (13).

Proof. The update constraint (14) ensures that (15) holds. Condition (16) imposed on the factors $\gamma_{1,\ldots,N-1}$ then ensures that (5b) is satisfied. ■

Corollary 2 ***Recursive Feasibility:*** *Constraint* (14) *is always feasible for any $\gamma \leq 1$*

Proof. We observe that (14) is feasible for $\boldsymbol{\pi}^k = \boldsymbol{\pi}^{k-1}$. ■

V. A Scenario-based Linear SMPC approach with Mission-Wide Guarantees

In this section we deploy the mission-wide SMPC idea developed so far in the linear case. Let us consider that the stochastic dynamics (3) are explicitly given by:

$$\mathbf{s}_{k+1} = A\mathbf{s}_k + B\mathbf{u}_k + \mathbf{w}_k\,, \tag{17}$$

and that the safe set $\mathbb{S}$ is polytopic, i.e.

$$\mathbb{S} = \{\, \mathbf{s} \,|\, C\mathbf{s} + \mathbf{c} \leq 0\}\,. \tag{18}$$

Here we assume that the disturbances $\mathbf{w}_k$, $k \in \mathbb{I}_{[0,N-1]}$ are i.i.d., and zero-mean for the sake of notation convenience.

At each time instants $k \in \mathbb{I}_{[0,N-1]}$, the predicted state $\mathbf{s}_t$ for all $t = k, k+1, \ldots, N$, can be split into a nominal part and an stochastic error part, i.e., $\mathbf{s}_t = \bar{\mathbf{s}}_t + \mathbf{e}_t$. we consider the policy sequence $\boldsymbol{\pi}_t^k$ parameterized via $\bar{\mathbf{u}}_t$, K, given by:

$$\mathbf{u}_t = \boldsymbol{\pi}_t^k(\mathbf{s}_t) := \bar{\mathbf{u}}_t + K\mathbf{e}_t, \quad \forall t \in \mathbb{I}_{[k,N-1]}$$

where K is a stabilizing feedback matrix for the nominal dynamics:

$$\bar{\mathbf{s}}_{t+1} = A\bar{\mathbf{s}}_t + B\bar{\mathbf{u}}_t \quad \bar{\mathbf{s}}_k = \mathbf{s}_k. \tag{19}$$

The stochastic error dynamics are then given by:

$$\mathbf{e}_{t+1} = (A + BK)\,\mathbf{e}_t + \mathbf{w}_t, \quad \mathbf{e}_k = 0\,. \tag{20}$$

Our goal is to solve the following mission-wide probability constrained optimal control problem at every time instant k:

$$\begin{aligned}
\min_{\bar{\mathbf{u}}_{k,\ldots,N-1}} \quad & \mathbb{E}\left[\mathbf{s}_N^\top Q_N \mathbf{s}_N + \sum_{t=k}^{N-1}\left(\mathbf{s}_t^\top Q \mathbf{s}_t + \mathbf{u}_t^\top R \mathbf{u}_t\right)\right] && \text{(21a)}\\
\text{s.t.} \quad & \bar{\mathbf{s}}_k = \mathbf{s}_k && \text{(21b)}\\
& \bar{\mathbf{s}}_{t+1} = A\bar{\mathbf{s}}_t + B\bar{\mathbf{u}}_t, \qquad \forall t \in \mathbb{I}_{[k,N-1]} && \text{(21c)}\\
& \mathbf{e}_{t+1} = (A+BK)\,\mathbf{e}_t + \mathbf{w}_t, \ \forall t \in \mathbb{I}_{[k,N-1]} && \text{(21d)}\\
& \mathbf{s}_{t+1} = \bar{\mathbf{s}}_{t+1} + \mathbf{e}_{t+1}, \qquad \forall t \in \mathbb{I}_{[k,N-1]} && \text{(21e)}\\
& \mathbb{P}[C\mathbf{s}_{t+1} + \mathbf{c} \leq 0,\ \forall t \in \mathbb{I}_{[k,N-1]}] \geq S_k\,. && \text{(21f)}
\end{aligned}$$

Here, Q, Q_N are semi-positive definite, R is positive definite, and the value

$$S_k = \gamma_k \mathbb{P}[\mathbf{s}_{k+1,\ldots,N} \in \mathbb{S} \,|\, \mathbf{s}_k, \boldsymbol{\pi}^{k-1}] \tag{22}$$

will be estimated at every time instant k using Monte Carlo simulation based on the real, closed-loop state $\mathbf{s}_k$ and the previous policy sequence $\boldsymbol{\pi}^{k-1}$.

A. An efficient scenario-based SMPC algorithm

Cost Function. Since $\mathbb{E}[\mathbf{s}_t] = \bar{\mathbf{s}}_t$ and $\mathbf{e}_t$ is zero mean (since $\mathbf{w}_t$ is zero mean by assumption), the cost function (21a) can be written explicitly as

$$\bar{\mathbf{s}}_N^\top Q_N \bar{\mathbf{s}}_N + \sum_{t=k}^{N-1} \left(\bar{\mathbf{s}}_t^\top Q \bar{\mathbf{s}}_t + \bar{\mathbf{u}}_t^\top R \bar{\mathbf{u}}_t \right) + \sigma,$$

where σ is a constant term that can be excluded from the cost function.

Chance Constraint. Substituting (21d) and (21e) into (21f), the constraints can be rewritten as

$$\mathbb{P}[C(\bar{\mathbf{s}}_{t+1} + (A+BK)\,\mathbf{e}_t + \mathbf{w}_t) + \mathbf{c} \le 0, \forall t \in \mathbb{I}_{[k,N-1]}] \ge S_k$$

and further be written as

$$\mathbb{P}[\underbrace{\mathcal{C}\mathcal{A}\mathbf{w}_{k,\ldots,N-1}^\top + [\mathbf{c},\ldots,\mathbf{c}]^\top}_{:=H} + \mathcal{C}\bar{\mathbf{s}}_{k+1,\ldots,N}^\top \le \mathbf{0}] \ge S_k \quad (23)$$

with matrix $\mathcal{A}$ obtained by condensing the dynamic (21d) and matrix $\mathcal{C}$ being block-diagonal with C as blocks.

Scenario Approximation. In general, providing a closed form for (23) is difficult. Fortunately, this problem can be handled efficiently with a scenario-based approach. Constraints (23) is replaced by a finite, sufficiently large number N_k of deterministic constraints resulting from sampling the disturbance sequence $\mathbf{w}_{k,\ldots,N-1}$. For a given time instant k, we define the i^{th} sample for all $i \in \mathbb{I}_{[1,N_k]}$ as

$$\mathbf{w}_{k,\ldots,N-1}^{(i)} := \{\mathbf{w}_k^{(i)}, \ldots, \mathbf{w}_{N-1}^{(i)}\},$$

Hence, the chance constraint (23) can be converted to

$$H^{(i)} + \mathcal{C}\bar{\mathbf{s}}_{k+1,\ldots,N}^\top \le \mathbf{0}, \quad \forall i \in \mathbb{I}_{[1,N_k]}, \quad (24)$$

where

$$H^{(i)} = \mathcal{C}\mathcal{A}(\mathbf{w}_{k,\ldots,N-1}^{(i)})^\top + [\mathbf{c},\ldots,\mathbf{c}]^\top.$$

In order to guarantee that (24) approximates (23) with a high probability $1-\beta$, where β is typically set to be very small (e.g., $\beta = 10^{-6}$), N_k must satisfy the following inequality [14]:

$$\sum_{n=1}^{d_k} \binom{N_k}{n} (1-S_k)^n S_k^{N_k - n} \le \beta,$$

where d_k is the number of optimization variables. The explicit lower bound of N_k can be further derived as [15]:

$$N_k \ge \frac{2}{1-S_k}\left(\ln\frac{1}{\beta} + d_k\right). \quad (25)$$

To further reduce the conservatism of the scenario-based approach, a sample removal approach is proposed in [16] and several variants are proposed. Their use here is beyond the scope of this paper.

For each scenario i, n_c linear constraints are generated in (24). It is clear that n_c is equal to the number of rows of matrix $H^{(i)}$. We additionally observe that for the constraint of index $j \in \mathbb{I}_{[1,n_\mathrm{c}]}$ in (24), the following inequality holds:

$$[H^{(i)}]_j + [\mathcal{C}]_j \bar{\mathbf{s}}_{k+1,\ldots,N} \le \max_{q \in \mathbb{I}_{[1,N_k]}} [H^{(q)}]_j + [\mathcal{C}]_j \bar{\mathbf{s}}_{k+1,\ldots,N}$$

for all $i \in \mathbb{I}_{[1,N_k]}$, where $[\star]_j$ denotes the j^{th} row of the matrix $\star$. Note that this inequality is tight, i.e., for all constraint of index j there always exists at least one sample of index i that ensures the above inequality tight. Hence j^{th} constraint is satisfied for all realizations i if they are satisfied for the one having the largest $[H^{(i)}]_j$.

Let us label:

$$\mathcal{I}_j = \max_{i \in \mathbb{I}_{[1,N_k]}} [H^{(i)}]_j, \quad \forall j \in \mathbb{I}_{[1,n_\mathrm{c}]}.$$

then we have that constraint (24) is equivalent to the following constraints:

$$\mathcal{I}_j + [\mathcal{C}]_j \bar{\mathbf{s}}_{k+1,\ldots,N} \le 0, \quad \forall j \in \mathbb{I}_{[1,n_\mathrm{c}]}.$$

Note that calculating $\mathcal{I}_j$, for all $j \in \mathbb{I}_{[1,n_\mathrm{c}]}$, requires only n_c (vector) maximum operations that are easy to implement and computationally efficient.

Now, (21) is equivalent to the following QP:

$$\begin{aligned} \min_{\bar{\mathbf{u}}_{k,\ldots,N-1}} \quad & \bar{\mathbf{s}}_N^\mathrm{T} Q_N \bar{\mathbf{s}}_N + \sum_{t=k}^{N-1} \left(\bar{\mathbf{s}}_t^\mathrm{T} Q \bar{\mathbf{s}}_t + \bar{\mathbf{u}}_t^\mathrm{T} R \bar{\mathbf{u}}_t \right) && (26a) \\ \text{s.t.} \quad & \bar{\mathbf{s}}_k = \mathbf{s}_k && (26b) \\ & \bar{\mathbf{s}}_{t+1} = A\bar{\mathbf{s}}_t + B\bar{\mathbf{u}}_t + \bar{\mathbf{w}}_t, \ \forall t \in \mathbb{I}_{[k,N-1]} && (26c) \\ & \mathcal{I}_j + [\mathcal{C}]_j \bar{\mathbf{s}}_{k+1,\ldots,N} \le 0, \quad \forall j \in \mathbb{I}_{[1,n_\mathrm{c}]} && (26d) \end{aligned}$$

yielding a regular QP of the same complexity as a normal linear MPC.

A systematic overview of the proposed scenario-based mission-wide linear SMPC scheme is summarized in Algorithm 1.

Algorithm 1: linear SMPC with MWPS constraints

Initialization: S_0, $\gamma_{1,\ldots,N-1}$, initial state $\mathbf{s}_0$
for $k = 0 : N-1$, **do**
  1) **if** $k \ge 1$ **then**
    Evaluate S_k in (22) through Monte Carlo simulation
  2) Generate N_k scenarios according to (25)
  3) Get the solution $\bar{\mathbf{u}}_{k,\ldots,N-1}^*$ by solving (26)
  4) Send $\bar{\mathbf{u}}_k^*$ to the actual system and update state: $\mathbf{s}_{k+1} = A\mathbf{s}_k + B\bar{\mathbf{u}}_k^* + \mathbf{w}_k$

B. Numerical Case Study

We consider the linear system (17) with

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.5 \\ 1 \end{bmatrix}$$

and the uncertainty is assumed to have a Gauss distribution

$$\mathbf{w}_k \sim \mathcal{N}(0, 0.04 \cdot I).$$

The safe (constraint) set $\mathbb{S}$ (18) is given by matrices

$$C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & 0 \\ 0 & -1 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} -2 \\ -2 \\ -10 \\ -2 \end{bmatrix}.$$

The matrices $Q = I$, $R = 0.1$, and

$$K = [-0.6167, -1.2703], \quad Q_N = \begin{bmatrix} 2.0599 & 0.5916 \\ 0.5916 & 1.4228 \end{bmatrix}$$

are computed from the corresponding LQR solution.

We select $N = 11$, $S_0 = 0.98$ and $\gamma_{1,\dots,10} = 0.99$, resulting in $S = \prod_{k=1}^{10} \gamma_k S_0 = 0.8863$. The number N_k of disturbance sample is selected from (25). The bound S_k given by (22) is evaluated from Monte Carlo simulation and $\beta = 10^{-6}$. In the simulations, we observed that $S_k \approx 0.99$ for all $k \in \mathbb{I}_{[1,10]}$. This is due to N_k calculated from (25) is conservative, such that the remaining MWPS at time k achieved by the previous policy sequence $\{\boldsymbol{\pi}_k^{k-1}, \dots, \boldsymbol{\pi}_{N-1}^{k-1}\}$, is much higher than that is actually required.

A Monte Carlo simulation that simulates 10^5 missions shows that the resulting ratio of mission success is 99.88%. This result is larger than $S = 88.63\%$. The reason for this discrepancy is that the scenario-based method adopted is conservative. Fig. 2 shows the state trajectories of 10^3 missions.

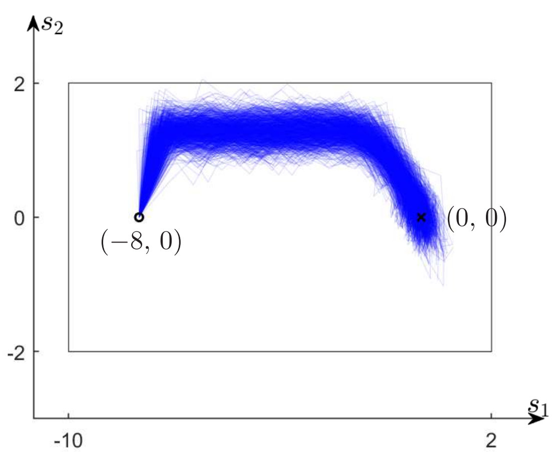

Fig. 2. State trajectories plot obtained by running 10^3 number of missions starting from the initial sate $\mathbf{s}_0 = [-8,\ 0]^\top$. The reference point is $[0,\ 0]^\top$. The rectangular area depicts the safe set $\mathbb{S}$.

VI. Conclusions and Future Work

We investigated optimal policies satisfying Mission-Wide Probability of Safety constraints, i.e. constraints imposing the safety of a system over an entire mission. This is in contrast with classical stochastic MPC, where safety constraints are imposed independently at every time stage. We show that recursive feasibility holds in the expected value sense for the concept of Mission-Wide Probability of Safety, opening a simple and practically meaningful concept of recursive feasibility for stochastic MPC. Optimal control with mission-wide probabilistic constraints is challenging. However, a computationally efficient scenario-based approach is proposed to solve this issue for linear stochastic problems. For the sake of brevity, a shrinking-horizon approach was presented in this paper. The scenario-based approach proposed here relies on classical Monte-Carlo sampling. More advanced methods will be developed in the future for the proposed method.

References

[1] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.

[2] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[3] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints – a review," *Journal of Process Control*, vol. 44, pp. 53 – 67, 2016.

[4] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram, "Chance-constrained dynamic programming with application to risk-aware robotic space exploration," *Auton. Robots*, vol. 39, no. 4, pp. 555–571, 2015.

[5] M. Ono and B. C. Williams, "An efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure," in *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pp. 1376–1382, AAAI Press, 2008.

[6] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Constraint-tightening and stability in stochastic model predictive control," *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3165–3177, 2017.

[7] M. Farina, L. Giulioni, L. Magni, and R. Scattolini, "An approach to output-feedback mpc of stochastic linear discrete-time systems," *Automatica*, vol. 55, pp. 140–149, 2015.

[8] J. A. Paulson, E. A. Buehler, R. D. Braatz, and A. Mesbah, "Stochastic model predictive control with joint chance constraints," *International Journal of Control*, vol. 93, no. 1, pp. 126–139, 2020.

[9] S. Yan, P. Goulart, and M. Cannon, "Stochastic model predictive control with discounted probabilistic constraints," in *2018 European Control Conference (ECC)*, pp. 1003–1008, 2018.

[10] M. Ono, "Joint chance-constrained model predictive control with probabilistic resolvability," in *American Control Conference*, pp. 435–441, IEEE, 2012.

[11] M. Ono, "Closed-loop chance-constrained MPC with probabilistic resolvability," in *Proceedings of the 51th IEEE Conference on Decision and Control*, pp. 2611–2618, 2012.

[12] G. Pola, J. Lygeros, and M. Benedetto, "Invariance in stochastic dynamical control systems," in *International Symposium on Mathematical Theory of Networks and Systems*, 2006.

[13] M. Ono and B. C. Williams, "Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint," in *Proceedings of the 47th IEEE Conference on Decision and Control, CDC Cancún, Mexico*, pp. 3427–3432, IEEE, 2008.

[14] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Transactions on Automatic Control*, vol. 51, no. 5, pp. 742–753, 2006.

[15] G. C. Calafiore, "A note on the expected probability of constraint violation in sampled convex programs," in *Proceedings of the IEEE International Conference on Control Applications, CCA and of the International Symposium on Intelligent Control, ISIC*, pp. 1788–1791, IEEE, 2009.

[16] M. C. Campi and S. A. Garatti, "A sampling-and-discarding approach to chance-constrained optimization: Feasibility and optimality," *Journal of Optimization Theory and Applications*, vol. 148, no. 2, p. 257–280, 2011.