

Provably-Stable Stochastic MPC for a Class of Nonlinear Contractive Systems

Cordiano, Francesco; Fochesato, Marta; Huang, Linbin; De Schutter, Bart

DOI

[10.1016/j.ifacol.2023.10.313](https://doi.org/10.1016/j.ifacol.2023.10.313)

Publication date

2023

Document Version

Final published version

Published in

IFAC-PapersOnLine

Citation (APA)

Cordiano, F., Fochesato, M., Huang, L., & De Schutter, B. (2023). Provably-Stable Stochastic MPC for a Class of Nonlinear Contractive Systems. *IFAC-PapersOnLine*, 56(2), 11242-11248. <https://doi.org/10.1016/j.ifacol.2023.10.313>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Provably-Stable Stochastic MPC for a Class of Nonlinear Contractive Systems

Francesco Cordiano* Marta Fochesato** Linbin Huang**
Bart De Schutter*

* *Delft Center for Systems and Control, TU Delft, 2628CD, Delft, Netherlands (e-mail: {f.cordiano, b.deschutter}@tudelft.nl)*

** *Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland (e-mail: {mfochesato, linhuang}@ethz.ch)*

Abstract: We present a model predictive control framework for a class of nonlinear systems affected by additive stochastic disturbances with (possibly) unbounded support. We consider hard input constraints and chance state constraints and we employ the unscented transform method to propagate the disturbances over the nonlinear dynamics in a computationally efficient manner. The main contribution of our work is the establishment of sufficient conditions for stability and recursive feasibility of the closed-loop system, based on the design of a terminal cost and a terminal set. We focus here on a special class of nonlinear systems that exhibit contractive properties in the dynamics. By assuming this property, we propose a novel approach to efficiently compute the terminal conditions without the need of performing any linearization of the dynamics. Finally, we provide an illustrative example to corroborate our theoretical findings.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Nonlinear model predictive control, stochastic systems, chance-constrained optimal control

1. INTRODUCTION

Model Predictive Control (MPC, (Rawlings et al., 2017)) is a well-established control paradigm for dynamical systems, owing to its strong theoretical properties and its ability of systematically handling state and input constraints. However, many real-world applications are subject to stochastic uncertainties that can affect the performance of the system in terms of cost, safety, and reliability. Two main approaches have been proposed in the literature to deal with uncertainties. On one side, robust MPC (Bemporad and Morari, 1999) addresses the worst-case disturbances in a bounded uncertainty set. However, the resulting policy is often tagged as too conservative (Köhler et al., 2019). On the other side, stochastic MPC (SMPC, (Mesbah, 2016)) provides constraints satisfaction with a desired level of probability in favor of better closed-loop performances, allowing to take into account possibly unbounded disturbances.

While SMPC has been successfully applied in the context of linear systems (see for example (Farina et al., 2013; Hewing and Zeilinger, 2018), and (Farina et al., 2016) for a complete review), stochastic model predictive control of nonlinear systems (SNMPC) has received relatively little attention in the literature. This is mainly due to (i) the computational complexity associated with the propagation of the uncertainty through the nonlinear dynamics; and (ii) the difficulty in encoding tractable sufficient conditions to ensure stability and recursive feasibility guarantees. While the first problem has been extensively studied in the literature with Gaussian mixture approximations (Weissel et al., 2009), unscented transformation (Völz and Graichen, 2015), and polynomial moments-based methods

(Paulson et al., 2015), the second problem has not found yet a satisfactory solution in the literature.

The analysis of stability guarantees in SNMPC is typically addressed by assuming the existence of a Lyapunov function that exhibits a decrease in expectation, usually enforced by means of a terminal cost in the objective function of the MPC program. For example, (McAllister and Rawlings, 2021) assume the existence of a Lyapunov function, without providing a way of computing it explicitly. In other works, such as (Buehler et al., 2016; Paulson et al., 2015), the computation of a suitable Lyapunov function is tackled via linearization of the dynamics, allowing to use the controller only where the linear approximation holds, which typically translates in a restricted region of attraction. Concerning recursive feasibility, in (McAllister and Rawlings, 2021) the problem is addressed by assuming the existence of a robust invariant terminal set, which, however, requires the uncertainty to be bounded. Furthermore, computing a terminal set with an invariance property can be challenging for a system with nonlinear dynamics (Lazar and Tetteroo, 2018; Yu et al., 2013). In summary, deriving sufficient conditions to enforce stability and recursive feasibility of SNMPC schemes is an open problem, due to the complexity of explicitly designing a terminal cost and a terminal set, and due to potentially large disturbances that might steer the state of the system arbitrarily far from a desired constraint set.

In this work, we consider a class of nonlinear systems that exhibit contractive properties in the dynamics. Contractive systems appear in many applications, such as system biology (Russo et al., 2011) and control of biochemical reactors (Aminzare and Sontag, 2014). Furthermore, con-

traction theory has important applications in stability of nonlinear systems (Köhler et al., 2020). We assume that the system is affected by additive stochastic disturbances with a possibly unbounded support, and that it is subject to hard input constraints and probabilistic state constraints. We provide an MPC framework equipped with guarantees, both for stability, intended as a bound on the expected value of the infinite-horizon closed-loop cost, and recursive feasibility. For this class of nonlinear systems, we propose a novel approach to efficiently compute the terminal cost and the terminal set leading to the sought closed-loop properties. Compared with the previously mentioned works, the proposed algorithm ensures stability and recursive feasibility globally without resorting to linearization. We finally propose an illustrative example to show our theoretical findings.

Notation. Let $\mathbb{R}_{\geq 0}$ denote the set of non-negative real numbers. Given a square matrix A , $\|x\|_A^2$ denotes the quadratic form defined as $x^\top Ax$, while $\rho(A)$ is its spectral radius and $\text{tr}(A)$ its trace. We write $A \geq 0$ to say that the matrix A is positive semi-definite. Let x be a n -dimensional random vector. We denote $\mu := \mathbb{E}[x] \in \mathbb{R}^n$ and $\Sigma := \mathbb{V}[x] \in \mathbb{R}^{n \times n}$, where $\mathbb{E}[\cdot]$, $\mathbb{V}[\cdot]$ denote the operators associated to the expected value and the covariance. The diagonal elements of Σ , namely the variance of each entry of x , are denoted by σ_i^2 . The n -dimensional identity matrix is denoted by I_n .

2. PRELIMINARIES

2.1 Problem Formulation

We consider a nonlinear discrete-time system affected by additive stochastic disturbances:

$$x(k+1) = f(x(k), u(k)) + w(k), \quad (1)$$

where k is the discrete time index, $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the system dynamics, $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $w \in \mathbb{R}^n$ are respectively the state, the control input, and the disturbances. The disturbances are assumed to be distributed according to a zero-mean distribution with possible unbounded support. We denote as Σ_w its covariance matrix, which is assumed to be element-wise bounded.

We assume that the system has an equilibrium in $(\bar{x}, \bar{u}) = (0, 0)$, which is a non-restrictive assumption since there always exists a linear transformation that maps a generic equilibrium (\bar{x}, \bar{u}) to $(0, 0)$. The goal of the controller is to stabilize the system around the origin by minimizing the expected value of a quadratic cost function, while fulfilling chance state constraints and hard input constraints of the form, respectively:

$$\mathbb{P}(x_i(k) \leq x_i^{\text{ub}}) \geq 1 - \varepsilon_i^{\text{ub}}, \quad i = 1, \dots, n \quad (2)$$

$$\mathbb{P}(x_i(k) \geq x_i^{\text{lb}}) \geq 1 - \varepsilon_i^{\text{lb}}, \quad i = 1, \dots, n \quad (3)$$

$$u^{\text{lb}} \leq u(k) \leq u^{\text{ub}}, \quad (4)$$

for all $k \in \mathbb{N}$. Here, $x^{\text{lb}}, x^{\text{ub}} \in \mathbb{R}^n$ and $u^{\text{lb}}, u^{\text{ub}} \in \mathbb{R}^m$ are upper and lower bounds for the state and the input, while $\varepsilon_i^{\text{ub}}, \varepsilon_i^{\text{lb}} \in (0, 1)$ are risk-tolerance parameters. Furthermore, we assume $x_i^{\text{lb}} \leq 0 \leq x_i^{\text{ub}}$, $i = 1, \dots, n$, and $u_j^{\text{lb}} \leq 0 \leq u_j^{\text{ub}}$, $j = 1, \dots, m$, i.e. $(\bar{x}, \bar{u}) = (0, 0)$ is a feasible equilibrium point.

We consider the class of nonlinear systems that satisfy the following assumption:

Assumption 1. The function $f(\cdot, \cdot)$ is Lipschitz continuous in its first argument, namely, for any x_1, x_2, u :

$$\|f(x_1, u) - f(x_2, u)\|_2 \leq L\|x_1 - x_2\|_2, \quad (5)$$

for some $L \in (0, 1)$. \square

In other words, we require the system dynamics to be contractive in the state, but we allow for an increase of the norm of the state due to the input. Let $i|k$ denote the i -th predicted step at the (closed-loop) iteration k , for $i = 0, \dots, N - 1$, where N denotes the prediction horizon. Then, given an initial value $x(0|k)$, we aim to solve the following finite-horizon stochastic optimal control problem in a receding horizon manner:

$$\min_{\substack{x(i|k), x(N|k), \\ u(i|k), \\ i=0, \dots, N-1}} \mathbb{E} \left[\sum_{i=0}^{N-1} (\|x(i|k)\|_Q^2 + \|u(i|k)\|_R^2) \right] + \mathbb{E} [\|x(N|k)\|_P^2] \quad (6a)$$

$$\text{s.t. } x(0|k) = x(k) \quad (6b)$$

$$x(i+1|k) = f(x(i|k), u(i|k)) + w(i|k) \quad (6c)$$

$$\mathbb{P}(x_j(i|k) \leq x_j^{\text{ub}}) \geq 1 - \varepsilon_j^{\text{ub}}, \quad j = 1, \dots, n \quad (6d)$$

$$\mathbb{P}(x_j(i|k) \geq x_j^{\text{lb}}) \geq 1 - \varepsilon_j^{\text{lb}}, \quad j = 1, \dots, n \quad (6e)$$

$$u^{\text{lb}} \leq u(i|k) \leq u^{\text{ub}} \quad (6f)$$

$$x(N|k) \in \mathcal{X}_f \quad (6g)$$

$$i = 0, \dots, N - 1.$$

Here, the cost function is quadratic, where we assume $Q, P \in \mathbb{R}^{n \times n} \geq 0$, and $R \in \mathbb{R}^{m \times m} > 0$, and \mathcal{X}_f is a terminal set. Problem (6) is computationally intractable due to the chance constraints (6d), (6e), which require the computation of multivariate integrals over the distribution of the state, and due to the propagation of the disturbances through the nonlinear dynamics. Hence, in the next section we propose a tractable reformulation of (6).

3. TRACTABLE REFORMULATION

3.1 Uncertainty Propagation

Propagating random variables through nonlinear functions is one of the main challenges in SNMPC. In this work we employ the unscented transform (UT), which has been originally introduced in the context of filtering (Julier and Uhlmann, 1997), and successfully applied in stochastic control problems (Völz and Graichen, 2015; Liu et al., 2014).

The uncertainty propagation problem is as follows. Given an n -dimensional (not necessarily Gaussian-distributed) random variable x with mean μ and covariance Σ , we wish to compute the statistical moments of the probability distribution of $y = f(x) \in \mathbb{R}^m$, where f denotes a generic nonlinear function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The UT method consists in the following steps:

- 1) Given a freely chosen tuning parameter $\kappa \in \mathbb{R}$, compute $2n + 1$ weights $w^{(i)}$, and $2n + 1$ points $x^{(i)}$ (called *sigma points*), whose sample mean and covariance are respectively μ and Σ :

$$\begin{aligned} x^{(0)} &= \mu & w^{(0)} &= \frac{\kappa}{n + \kappa} \\ x^{(i)} &= \mu + (\sqrt{(n + \kappa)\Sigma})_i & w^{(i)} &= \frac{1}{2(n + \kappa)} \\ x^{(i+n)} &= \mu - (\sqrt{(n + \kappa)\Sigma})_i & w^{(i+n)} &= \frac{1}{2(n + \kappa)}, \end{aligned}$$

for $i = 1, \dots, n$, where $(\sqrt{(n + \kappa)\Sigma})_i$ denotes the i -th column of the matrix $\sqrt{(n + \kappa)\Sigma}$, and $\sqrt{\Sigma}$ is the Cholesky factor of the matrix Σ .

- 2) Evaluate $y^{(i)} = f(x^{(i)})$, $i = 0, \dots, 2n$.
- 3) Then, the mean and the covariance of $f(x)$ are approximated as a weighted sum of the statistical moments of the transformed points y_i :

$$\begin{aligned} \mu_y &\approx \sum_{i=0}^{2n} w^{(i)} y^{(i)}, \\ \Sigma_y &\approx \sum_{i=0}^{2n} w^{(i)} (y^{(i)} - \mu_y)(y^{(i)} - \mu_y)^\top. \end{aligned}$$

Although the UT resembles Monte Carlo approximation methods, an important difference is that we choose the sigma points according to the specific criteria in step 1) and not at random, which leads to a significantly lower number of required samples.

(Julier and Uhlmann, 1997) shows that the UT approximation is accurate up to the second order of the Taylor expansion, arguing that the parameter κ can be chosen to tune the higher-order terms in the approximation, see (Julier and Uhlmann, 1997) for more details. It is worthwhile to remark that the computational cost of this method grows linearly with n , and it is lower than other approaches such as the generalized polynomial chaos expansion, whose complexity grows polynomially in the number of terms employed in the approximation (Paulson et al., 2015). Hence, the UT is the chosen way to propagate the moments of the state distribution throughout the prediction horizon. By denoting $\mu(i|k) = \mathbb{E}[x(i|k)]$, $\Sigma(i|k) = \mathbb{V}[x(i|k)]$, the dynamics of the system are described by:

$$\begin{aligned} \mu(i + 1|k) &= \mathbb{E}[x(i + 1|k)] = \mathbb{E}[f(x(i|k), u(i|k))] & (7) \\ \Sigma(i + 1|k) &= \mathbb{V}[x(i + 1|k)] = \mathbb{V}[f(x(i|k), u(i|k))] + \Sigma_w, & (8) \end{aligned}$$

for $i = 0, \dots, N - 1$, with $x(0|k)$ assumed to be deterministic under the assumption of perfect state measurement.

3.2 Chance Constraints

Let us consider an individual chance constraint of the type (6d), where we drop the indices for ease of notation. In the following we replace it with distributionally robust chance constraint of the form:

$$\inf_{x \sim \mathcal{L}(\mu, \sigma)} \mathbb{P}(x - x^{\text{ub}} \leq 0) \geq 1 - \varepsilon^{\text{ub}}, \quad (9)$$

where x is a scalar random variable, and $\mathcal{L}(\mu, \sigma)$ denotes the family of all possible distributions with mean μ and standard deviation σ . We observe that if (9) is satisfied, then also (6d) is satisfied. The main advantage of considering (9) is that it does not require the exact computation of the distribution of the state x , which might be complicated due to the nonlinear dynamics. For any $\varepsilon^{\text{ub}} \in (0, 1)$, (9) can be equivalently formulated as the following convex con-

straint, in terms of μ and σ , via the Chebyshev inequality (Saw et al., 1984):

$$\sqrt{\frac{1 - \varepsilon^{\text{ub}}}{\varepsilon^{\text{ub}}}} \sqrt{\mathbb{V}[x - x^{\text{ub}}]} + \mathbb{E}[x - x^{\text{ub}}] \leq 0 \quad (10)$$

$$\Leftrightarrow \mu \leq x^{\text{ub}} - \sqrt{\frac{1 - \varepsilon^{\text{ub}}}{\varepsilon^{\text{ub}}}} \sigma. \quad (11)$$

Hence, constraint (11) is a tractable reformulation of the chance constraint (6d). We observe that Chebyshev inequality replaces a chance constraint in x with a convex constraint in μ and σ , and it provides constraint satisfaction at least with the desired probability independently of the actual distribution of the state. As a direct consequence, constraint (11) can be a conservative approximation of the original chance constraint (6d). Similarly, constraints of the type of (6e) can be reformulated as

$$\mu \geq x^{\text{lb}} + \sqrt{\frac{1 - \varepsilon^{\text{lb}}}{\varepsilon^{\text{lb}}}} \sigma. \quad (12)$$

3.3 Cost Function

The cost can then be written in terms of $\mu(i|k)$, $\Sigma(i|k)$ as

$$\begin{aligned} &\mathbb{E} \left[\sum_{i=0}^{N-1} (\|x(i|k)\|_Q^2 + \|u(i|k)\|_R^2) + \|x(N|k)\|_P^2 \right] \\ &= \sum_{i=0}^{N-1} (\|\mu(i|k)\|_Q^2 + \|u(i|k)\|_R^2) + \|\mu(N|k)\|_P^2 \\ &+ \sum_{i=0}^{N-1} \text{tr}(Q\Sigma(i|k)) + \text{tr}(P\Sigma(N|k)), \end{aligned} \quad (13)$$

exploiting the standard probability argument $\mathbb{E}[\|x(i|k)\|_Q^2] = \|\mathbb{E}[x(i|k)]\|_Q^2 + \text{tr}(Q\mathbb{V}[x(i|k)])$.

3.4 Initial Condition

At each closed-loop time step, we solve one instance of the optimal control problem (6) by initializing it with the most recent state measurement $x(k)$. However, this might lead to infeasibility issues as potentially unbounded noise might drive the state of the system arbitrarily far from the state constraint set. Therefore, similarly to (Farina et al., 2013), we observe that the optimal solution $(\mu^*(1|k-1), \Sigma^*(1|k-1))$, obtained at time step $k-1$, results in a feasible initialization at time step k , leading to the following strategies for the initialization:

- S_1 : $\mu(0|k) = x(k)$, $\Sigma(0|k) = 0$,
- S_2 : $\mu(0|k) = \mu^*(1|k-1)$, $\Sigma(0|k) = \Sigma^*(1|k-1)$.

Hence, the variables $(\mu(0|k), \Sigma(0|k))$ are also decision variables of the problem according to the two following alternative choices for the initial constraint:

$$(\mu(0|k), \Sigma(0|k)) \in \{(x(k), 0), (\mu^*(1|k-1), \Sigma^*(1|k-1))\}. \quad (14)$$

Thanks to this strategy, feasibility issues in the initial constraint are eliminated, since $(\mu^*(1|k-1), \Sigma^*(1|k-1))$ is feasible at time step k . Another viable initialization scheme is provided in (Köhler and Zeilinger, 2022), consisting in a convex combination of the two strategies to avoid the need of solving two optimization problems whenever S_1 leads to infeasibility.

3.5 Optimization Problem

At each time-step $k \in \mathbb{N}$ we solve the following problem:

$$\min_{\mu, \Sigma, u} \quad (13) \quad (15a)$$

$$\text{s.t.} \quad (4), (7), (8), (11), (12), (14), \quad (15b)$$

$$\mu(N|k) \in \mathcal{X}_f. \quad (15c)$$

Let $u^*(i|k), i = 0, \dots, N-1$, be the optimal input sequence to problem (15) at time-step k . Then, according to the receding horizon implementation, we apply to the system only the first element, i.e., $u^*(0|k)$, observe the transition of the system to $x(k+1)$, and solve (15) again with a different initialization. We highlight that the tractable reformulation (15) holds for a generic nonlinear system and does not rely on Assumption 1.

Remark 2. Due to the double alternatives for the initial constraint, the applied input u is not, in general, a state feedback, but it is a function of the augmented state $\tilde{x}(k) = (x(k), \mu^*(1|k-1), \Sigma^*(1|k-1))$, i.e., we introduce feedback only when strategy S_1 is selected, resulting in closed-loop constraint satisfaction. When strategy S_2 is selected, chance constraints are satisfied with the desired probability only in prediction. This drawback is present in general in other schemes employing this backup strategy for the initial constraint (Farina et al., 2013).

4. CLOSED-LOOP PROPERTIES

4.1 Preliminaries

We begin by establishing the following result:

Proposition 3. Under Assumption 1 it holds that $\text{tr}(\mathbb{V}[x(k)]) \leq \bar{\Sigma}, \forall k \in \mathbb{N}$, where $\bar{\Sigma} = \frac{1}{1-L^2} \text{tr}(\Sigma_w)$. \square

Proof. We begin by proving that $\text{tr}(\mathbb{V}[f(x(k), u(k))]) \leq L^2 \text{tr}(\mathbb{V}[x(k)])$. Recalling that the covariance of a random variable is translation-invariant, and defining $\tilde{f} = f(x(k), u(k)) - f(\mathbb{E}[x(k)], u(k))$, we have:

$$\text{tr}(\mathbb{V}[f(x(k), u(k))]) = \text{tr}(\mathbb{V}[\tilde{f}]) \quad (16)$$

$$= \mathbb{E}[\tilde{f}^T \tilde{f}] - \mathbb{E}[\tilde{f}]^T \mathbb{E}[\tilde{f}] \quad (17)$$

$$\leq \mathbb{E}[\tilde{f}^T \tilde{f}] = \mathbb{E}[\|f(x(k), u(k)) - f(\mathbb{E}[x(k)], u(k))\|_2^2] \quad (18)$$

$$\leq L^2 \text{tr}(\mathbb{V}[x(k)]). \quad (19)$$

where in (17) we use that: $\text{tr}(\mathbb{V}[X]) = \mathbb{E}[X^T X] - \mathbb{E}[X]^T \mathbb{E}[X]$ and in (19) we exploit Assumption 1 and $\text{tr}(\mathbb{V}[X]) = \mathbb{E}[\|X - \mathbb{E}[X]\|_2^2]$, for any random variable X . By iteratively applying the dynamics of the covariance (8), we obtain:

$$\text{tr}(\mathbb{V}[x(k+1)]) \leq L^{2(k+1)} \text{tr}(\mathbb{V}[x(0)]) + \text{tr}(\Sigma_w) \sum_{j=0}^k L^{2j}.$$

Since we assume initial feasibility, we have: $\mathbb{V}[x(0)] = 0$. The term $\sum_{j=0}^k L^{2j}$ is a geometric series truncated after $k+1$ terms, and as $L \in (0, 1)$, it is upper-bounded by $\frac{1}{1-L^2}$. Hence, the following holds, for all $k \in \mathbb{N}$:

$$\text{tr}(\mathbb{V}[x(k+1)]) \leq \frac{1}{1-L^2} \text{tr}(\Sigma_w) =: \bar{\Sigma}. \quad (20)$$

This concludes the proof. \blacksquare

Next, we introduce the following assumption that is instrumental for the construction of the terminal set:

Assumption 4. Let $\bar{\mathcal{X}} = \{\mu : x_i^{\text{lb}} + \sqrt{\frac{1-\varepsilon_i^{\text{lb}}}{\varepsilon_i^{\text{lb}}}} \sqrt{\bar{\Sigma}} \leq \mu_i \leq x_i^{\text{ub}} - \sqrt{\frac{1-\varepsilon_i^{\text{ub}}}{\varepsilon_i^{\text{ub}}}} \sqrt{\bar{\Sigma}}, i = 1, \dots, n\}$, where $\bar{\Sigma}$ is the upper bound given in Proposition 3. Then, we assume there exists $\alpha \in \mathbb{R}_{\geq 0}$ such that the terminal set defined as $\mathcal{X}_f = \{\mu(N|k) \in \bar{\mathcal{X}} : \|\mu(N|k)\|_2^2 \leq \alpha + \frac{L^2}{1-L^2} \bar{\Sigma}, k \in \mathbb{N}\}$ is nonempty. \square

In addition, the following property is established.

Proposition 5. Let $\pi_f : \mathcal{X}_f \rightarrow \mathbb{R}^m$ be a terminal controller. If Assumption 1 holds, the terminal set \mathcal{X}_f described in Assumption 4 is invariant under the terminal controller $\pi_f(x(N|k)) = 0$. \square

Proof. Under Assumption 1, we can make use of the bound on the variance provided by Proposition 3. By means of Jensen's inequality, for any $\mu(N|k) \in \mathcal{X}_f$, it holds that:

$$\|\mu(N+1|k)\|_2^2 = \|\mathbb{E}[f(x(N|k), 0) - f(0, 0)]\|_2^2 \quad (21)$$

$$\leq \mathbb{E}[\|f(x(N|k), 0) - f(0, 0)\|_2^2] \quad (22)$$

$$\leq L^2 \mathbb{E}[\|x(N|k)\|_2^2] \quad (23)$$

$$= L^2 (\|\mu(N|k)\|_2^2 + \text{tr}(\Sigma(N|k))) \quad (24)$$

$$\leq L^2 \alpha + \left(L^2 \frac{L^2}{1-L^2} + L^2 \right) \bar{\Sigma} \quad (25)$$

$$\leq \alpha + \frac{L^2}{1-L^2} \bar{\Sigma}, \quad (26)$$

where in (25)-(26) we exploit that $\mu(N|k) \in \mathcal{X}_f$, the bound $\text{tr}(\Sigma(N|k)) \leq \bar{\Sigma}$, and $L < 1$. \blacksquare

We have now all the ingredients to evaluate feasibility and stability properties of the closed-loop system.

4.2 Recursive Feasibility

Theorem 6. Assume that at time-step $k=0$ problem (15) is feasible for a given initial condition $x(0)$ (i.e., strategy S_1 is applied at $k=0$). In addition, let Assumptions 1 and 4 hold. Then, the MPC optimization problem (15) is recursively feasible. \square

Proof. At a given time step k , let us consider the following candidate solution, constructed by shifting the optimal solution at time step $k-1$ and completing it by means of the terminal controller $\pi_f(x(N|k)) = 0$:

$$\begin{aligned} & \{\mu^*(1|k), \dots, \mu^*(N|k), \mu(N+1|k)\} \\ & \{\Sigma^*(1|k), \dots, \Sigma^*(N|k), \Sigma(N+1|k)\} \\ & \{u^*(1|k), \dots, u^*(N|k), 0\}, \end{aligned} \quad (27)$$

where $\mu(N+1|k)$ and $\Sigma(N+1|k)$ are the mean and the covariance of the last predicted state $x(N+1|k) = f(x^*(N|k), 0) + w(N|k)$, namely: $\mu(N+1|k) := \mathbb{E}[x(N+1|k)]$, $\Sigma(N+1|k) := \mathbb{V}[x(N+1|k)]$.

We now prove that the candidate solution (27) is feasible for the optimal control problem solved at time step $k+1$. First of all, choosing $\pi_f(x(N|k)) = 0$ results in a feasible terminal controller, since by assumption 0 satisfies the input constraints. In addition, $(\mu^*(1|k), \Sigma^*(1|k))$ satisfies the initial constraint, according to strategy S_2 . As (27) is feasible at time step k , at time step $k+1$ the first N

terms in (27) satisfy the chance constraints, as well as the dynamics of the system. We prove now that the last predicted state in (27) satisfies the chance constraint (11). Since $\mu(N|k) \in \mathcal{X}_f \subseteq \mathcal{X}$, we know that, for $j = 1, \dots, n$:

$$\mu_j(N+1|k) \leq x_j^{\text{ub}} - \sqrt{\frac{1 - \varepsilon_j^{\text{ub}}}{\varepsilon_j^{\text{ub}}}} \sqrt{\bar{\Sigma}}. \quad (28)$$

The trace of the covariance matrix of state $x(N+1|k)$ is always bounded by $\bar{\Sigma}$ thanks to Proposition 3. This implies that also the single variances are bounded by $\bar{\Sigma}$, namely $\sigma_j^2(N+1|k) \leq \bar{\Sigma}$, leading to:

$$\mu_j(N+1|k) \leq x_j^{\text{ub}} - \sqrt{\frac{1 - \varepsilon_j^{\text{ub}}}{\varepsilon_j^{\text{ub}}}} \sigma_j(N+1|k), \quad (29)$$

which proves that $(\mu(N+1|k), \Sigma(N+1|k))$ satisfies the chance constraint (11). A similar procedure can be derived for chance constraint (12). ■

4.3 Stability

We require the following assumption:

Assumption 7. The cost matrices $Q, P \in \mathbb{R}^{n \times n}$ satisfy: $P \geq \rho(P)L^2I_n + Q$. □

Theorem 8. Assume that at time-step $k=0$ problem (15) is feasible for a given initial condition $x(0)$ (i.e., strategy S_1 is applied at $k=0$). In addition, let Assumptions 1, 7 hold. Then, the closed-loop system under $u(k) = u^*(0|k)$ satisfies:

$$\lim_{T \rightarrow \infty} l_{\text{avg}}(T) \leq \text{tr}(P\Sigma_w), \quad (30)$$

where $l_{\text{avg}}(T) = \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E} [\|x(k)\|_Q^2 + \|u(k)\|_R^2]$. □

Proof. We prove now that, under Assumptions 1, 7, the cost function of the closed-loop system exhibits a Lyapunov-like decrease condition in the augmented state $\tilde{x}(k) = (x(k), \mu^*(1|k-1), \Sigma^*(1|k-1))$, which then implies (30). Since the initial constraint is a decision variable according to (14), it is sufficient to prove that the required decrease condition holds for the initialization strategy S_2 , since, whenever feasible, the cost associated to strategy S_1 will be lower. Hence, we will make use of the candidate solution (27) in the following proof, which is a feasible, albeit suboptimal, solution of the SNMPC problem. Let $J^*(\tilde{x}(k))$ be the cost function associated to the optimal solution obtained at time step k . By denoting the candidate sequence (27) at time step k as $\tilde{x}_c(k)$, its associated cost is:

$$\begin{aligned} J(\tilde{x}_c(k)) &= \sum_{i=1}^N (\|\mu^*(i|k)\|_Q^2 + \|u^*(i|k)\|_R^2) \\ &\quad + \|\mu(N+1|k)\|_P^2 \\ &\quad + \sum_{i=1}^N \text{tr}(Q\Sigma^*(i|k)) + \text{tr}(P\Sigma(N+1|k)). \end{aligned}$$

The corresponding cost decrease is:

$$\begin{aligned} J(\tilde{x}_c(k+1)) - J^*(\tilde{x}(k)) &= \\ &= -\|\mu^*(0|k)\|_Q^2 - \text{tr}(Q\Sigma^*(0|k)) - \|u^*(0|k)\|_R^2 \\ &\quad + \|\mu^*(N|k)\|_Q^2 + \|\mu(N+1|k)\|_P^2 - \|\mu^*(N|k)\|_P^2 \\ &\quad - \text{tr}((P-Q)\Sigma^*(N|k)) + \text{tr}(P\Sigma(N+1|k)). \end{aligned} \quad (31)$$

Recalling that w is zero-mean, in (31) we have:

$$\begin{aligned} &\|\mu(N+1|k)\|_P^2 + \text{tr}(P\Sigma(N+1|k)) \\ &= \mathbb{E}[\|f(x^*(N|k), 0) + w(N|k)\|_P^2] \\ &= \mathbb{E}[\|f(x^*(N|k), 0)\|_P^2 + \|w(N|k)\|_P^2] \\ &\leq \rho(P)L^2 \mathbb{E}[\|x^*(N|k)\|_2^2] + \|w(N|k)\|_P^2 \\ &\leq \rho(P)L^2 \mathbb{E}[\|x^*(N|k)\|_2^2] + \text{tr}(P\Sigma_w). \end{aligned}$$

In view of this, and recalling that the candidate sequence can be suboptimal, it holds that:

$$\begin{aligned} &J^*(\tilde{x}(k+1)) - J^*(\tilde{x}(k)) \\ &\leq J(\tilde{x}_c(k+1)) - J^*(\tilde{x}(k)) \\ &\leq -\|\mu^*(0|k)\|_Q^2 - \text{tr}(Q\Sigma^*(0|k)) - \|u^*(0|k)\|_R^2 \\ &\quad + \|\mu^*(N|k)\|_Q^2 + \rho(P)L^2 \mathbb{E}[\|x^*(N|k)\|_2^2] + \text{tr}(P\Sigma_w) \\ &\quad - \|\mu^*(N|k)\|_P^2 - \text{tr}((P-Q)\Sigma^*(N|k)) \\ &= -\|\mu^*(0|k)\|_Q^2 - \text{tr}(Q\Sigma^*(0|k)) - \|u^*(0|k)\|_R^2 + \text{tr}(P\Sigma_w) \\ &\quad + \mathbb{E}[\|x^*(N|k)\|_Q^2] - \mathbb{E}[\|x^*(N|k)\|_{P-\rho(P)L^2I_n}^2]. \end{aligned}$$

Recalling that $P - \rho(P)L^2I_n \geq Q$ due to Assumption 7, and by the monotonicity of the expectation, we have:

$$\begin{aligned} &J^*(\tilde{x}(k+1)) - J^*(\tilde{x}(k)) \quad (32) \\ &\leq -\|\mu^*(0|k)\|_Q^2 - \text{tr}(Q\Sigma^*(0|k)) - \|u^*(0|k)\|_R^2 + \text{tr}(P\Sigma_w) \\ &= -\mathbb{E}[\|x^*(0|k)\|_Q^2] - \|u^*(0|k)\|_R^2 + \text{tr}(P\Sigma_w) \\ &= -\mathbb{E}[\|x(k)\|_Q^2 + \|u(k)\|_R^2] + \text{tr}(P\Sigma_w), \quad (33) \end{aligned}$$

where the last line follows since we assume perfect state measurement. Given a closed-loop horizon of length T , summing (32) and (33) over k results in:

$$J^*(\tilde{x}(T)) - J^*(\tilde{x}(0)) \quad (34)$$

$$= \sum_{k=0}^{T-1} (J^*(\tilde{x}(k+1)) - J^*(\tilde{x}(k))) \quad (35)$$

$$\leq -\sum_{k=0}^{T-1} (\mathbb{E}[\|x(k)\|_Q^2 + \|u(k)\|_R^2]) + T\text{tr}(P\Sigma_w), \quad (36)$$

dividing (34), (36) by T and taking the limit for T to infinity we get:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=0}^{T-1} \mathbb{E} [\|x(k)\|_Q^2 + \|u(k)\|_R^2] \leq \text{tr}(P\Sigma_w). \quad (37)$$

This concludes the proof. ■

Remark 9. The quantity l_{avg} is a typical tool to quantify stability in SMPC (Hewing and Zeilinger, 2018; Chaouach et al., 2022). If $\lim_{T \rightarrow \infty} l_{\text{avg}}(T)$ is bounded, we know that the quantity $\sum_{k=0}^T \mathbb{E} [\|x(k)\|_Q^2 + \|u(k)\|_R^2]$ grows at most linearly in T . Hence, $\mathbb{E} [\|x(k)\|_Q^2 + \|u(k)\|_R^2]$ converges to a finite value. This is an index of stability property of the system. Furthermore, note that Theorem 8 establishes that a contraction in the system dynamics translates into a contraction of the cost function despite the presence of unbounded noise.

5. NUMERICAL EXAMPLE

In this section, we carry out a numerical example, implemented in MATLAB with CasADi (Andersson et al., 2019) and IPOPT as solver (Biegler and Zavala, 2009), a primal-dual interior point method. All the simulations have been




Fig. 1. The closed-loop trajectories, tested on 500 samples, converge to a neighbourhood of the origin.

performed on a Macbook Pro (Apple Silicon M1 pro, 32GB RAM). We consider the following oscillator, adapted from (Dashkovskiy, 2019):

$$\begin{cases} x_1(k+1) = 0.5x_1(k) + 0.5x_2(k) + w_1(k) \\ x_2(k+1) = 0.5\sin(x_1(k)) + 0.5x_2(k) + u(k) + w_2(k) \end{cases} \quad (38)$$

We consider hard input constraints as $-2 \leq u(k) \leq 2$, and chance constraints of the form $\mathbb{P}(x_2(k) \leq 0.5) \geq 1 - \varepsilon$, with $\varepsilon = 0.1$. The disturbances $w(k) \in \mathbb{R}^2$ are Gaussian-distributed, with mean $[0 \ 0]^\top$ and covariance matrix $\Sigma_w = \begin{bmatrix} 0.0025 & 0.0005 \\ 0.0005 & 0.0025 \end{bmatrix}$. The nonlinearity is introduced by the sin function, which is Lipschitz continuous with Lipschitz constant 1. Assumption 1 is satisfied with $L = 0.5 \cdot \sqrt{2} \approx 0.707$. From Proposition 3, we have $\bar{\Sigma} = 0.01$, where $\bar{\Sigma}$ is an upper bound on $\text{tr}(\mathbb{V}[x(k)])$ and thus on the single variances of the state components. The set \mathcal{X} is accordingly designed as $\mathcal{X} = \{\mu \in \mathbb{R}^2 : \mu_2 \leq 0.2\}$, following Assumption 4. Hence, the terminal set $\mathcal{X}_f = \{\mu(N|k) : \|\mu(N|k)\|_2 \leq 0.2\} \subseteq \mathcal{X}$ is invariant according to Proposition 5. Finally, we choose $Q = 0.1I_2, P = I_2, R = 0.1I_2$, which satisfy Assumption 7. We consider a closed-loop simulation starting from the initial condition $(-3, -4)$, and we set $N = 5$. We employ the UT to update the mean and the covariance of the state in the prediction horizon as described in Section 3.1, setting $\kappa = 1$.

Figure 1 shows 500 closed-loop experiments. We notice that the controller is able to stabilize the closed-loop system to the equilibrium in $(0, 0)$, despite mild oscillations due to the additive disturbances. In particular, the theoretical bound (30) is reached after approximately 1500 time steps. Every time step in which $x_2(k)$ exceeds the constraint bound, strategy S_2 is chosen. As we pointed out in Remark 2, when S_2 is selected we have chance constraints satisfaction only in prediction, which might lead to an empirical violation rate larger than the desired one. However, in this example we notice that the empirical constraint violation amounts to a maximum value over time of 2%, significantly smaller than the theoretical violation rate $\varepsilon = 0.1$. This reflects the conservatism introduced by the Chebyshev inequality, which is a sufficient condition for chance constraint satisfaction.

We also showcase the effectiveness of the UT to propagate the stochastic disturbances through the dynamics of system (38). Figure 2 shows the open-loop predicted state trajectories resulting from the solution of three different SMPC schemes, each of which employs a different propagation method. In particular, we compare the computation




Fig. 2. Comparison of the open-loop trajectories in a prediction horizon of $N = 10$.

of the mean and the variance of the state x_2 of system (38), over a prediction horizon $N = 10$. To better visualize the error propagation, the values of covariance matrix of the additive disturbance are increased: $\tilde{\Sigma}_w = 4 \cdot \Sigma_w$. The initial condition is $(-3, -4)$ and it is deterministic for all the three methods. The first scheme is the proposed method based on the UT. The second scheme (LIN) is based on the linearization of the dynamics (38) around the equilibrium $(0, 0)$. This gives: $A = \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}, B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, and the mean and the covariance of the state are updated according to:

$$\begin{aligned} \mu^{\text{LIN}}(i+1) &= A\mu^{\text{LIN}}(i) + Bu(i), \\ \Sigma^{\text{LIN}}(i+1) &= A\Sigma^{\text{LIN}}(i)A^\top + \tilde{\Sigma}_w. \end{aligned}$$

The third scheme performs a Monte Carlo (MC) approximation of the mean and the covariance of the state with $M = 1000$ samples for each time step; hence it is considered the ground-truth. For each $i = 0, \dots, N-1$, the MC approximation consists in the following equations:

$$\begin{aligned} \mu_{\text{MC}}(i) &= \frac{1}{M} \sum_{j=1}^M x^{(j)}(i), \\ \Sigma_{\text{MC}}(i) &= \frac{1}{M} \sum_{j=1}^M (x^{(j)}(i) - \mu_{\text{MC}})(x^{(j)}(i) - \mu_{\text{MC}})^\top, \end{aligned}$$

where each $x^{(j)}(i)$ follows the dynamics (38), $i = 0, \dots, N-1$, $j = 1, \dots, M$. Table 1 reports the time required by the solver to compute the three predicted trajectories in Figure 2, as well as the errors of mean and variance of the state x_2 of UT and LIN compared to MC. We observe that the propagation error in the prediction horizon is very small for the UT, and its accuracy is comparable to an MC approximation, being at the same time much more computationally efficient than MC. On the other side, the scheme based on linearization leads to a larger error in the propagation of the dynamics, and it becomes accurate only close to the equilibrium. Similar results can be derived for the state x_1 , which are omitted in the interest of space. Since our numerical results show that the UT is very accurate, the theoretical results for recursive feasibility and stability, which assume exact propagation of the disturbances, are practically not compromised.

6. CONCLUSION

This paper presents a provably-stable and recursively feasible MPC framework for a class of stochastic nonlinear systems subject to possibly unbounded additive disturbances. Assuming a contractive property in the system

	UT	LIN	MC
Avg. solver time [sec.]	0.32	0.081	257.9
$\max_{i=0,\dots,N} \mu_2(i) - \mu_2^{\text{MC}}(i) $	0.06	1.14	0
$\max_{i=0,\dots,N} \sigma_2(i) - \sigma_2^{\text{MC}}(i) $	0.0029	0.015	0

Table 1. Solver time and approximation errors for the schemes in Figure 2.

dynamics, we propose a computationally-efficient design of the terminal cost and of the terminal set leading to the sought closed-loop properties. Robustifying stability and feasibility properties with respect to the approximation error in the mean and the covariance matrix, as well as the development of methods to ensure closed-loop chance constraints are relevant future work. Furthermore, we are also interested in broadening the class of nonlinear systems for which we can guarantee closed-loop properties under tractable design of the terminal set and the terminal cost.

ACKNOWLEDGEMENTS

Part of this research has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101018826 - ERC Advanced Grant CLar-iNet). In addition, we wish to thank Marcello Farina and Lotfi Chaouach for the fruitful discussions.

REFERENCES

- Aminzare, Z. and Sontagy, E.D. (2014). Contraction methods for nonlinear systems: A brief introduction and some open problems. In *53rd Conference on Decision and Control*, 3835–3847.
- Andersson, J.A., Gillis, J., Horn, G., Rawlings, J.B., and Diehl, M. (2019). CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1), 1–36.
- Bemporad, A. and Morari, M. (1999). Robust model predictive control: A survey. In *Robustness in identification and control*, 207–226. Springer.
- Biegler, L.T. and Zavala, V.M. (2009). Large-scale nonlinear programming using IPOPT: An integrating framework for enterprise-wide dynamic optimization. *Computers & Chemical Engineering*, 33(3), 575–582.
- Buehler, E.A., Paulson, J.A., and Mesbah, A. (2016). Lyapunov-based stochastic nonlinear model predictive control: Shaping the state probability distribution functions. In *2016 American Control Conference (ACC)*, 5389–5394.
- Chaouach, L.M., Fiacchini, M., and Alamo, T. (2022). Stochastic model predictive control for linear systems affected by correlated disturbances. *IFAC-PapersOnLine*, 55(25), 133–138.
- Dashkovskiy, S. (2019). Practical examples of ISS systems. *IFAC-PapersOnLine*, 52(16), 1–6.
- Farina, M., Giulioni, L., Magni, L., and Scattolini, R. (2013). A probabilistic approach to model predictive control. In *52nd Conference on Decision and Control*, 7734–7739.
- Farina, M., Giulioni, L., and Scattolini, R. (2016). Stochastic linear model predictive control with chance constraints—a review. *Journal of Process Control*, 44, 53–67.
- Hewing, L. and Zeilinger, M.N. (2018). Stochastic model predictive control for linear systems using probabilistic reachable sets. In *2018 57th Conference on Decision and Control (CDC)*, 5182–5188.
- Julier, S.J. and Uhlmann, J.K. (1997). New extension of the Kalman filter to nonlinear systems. In *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068, 182–193.
- Köhler, J., Andina, E., Soloperto, R., Müller, M.A., and Allgöwer, F. (2019). Linear robust adaptive model predictive control: Computational complexity and conservatism. In *2019 58th Conference on Decision and Control (CDC)*, 1383–1388.
- Köhler, J., Soloperto, R., Müller, M.A., and Allgöwer, F. (2020). A computationally efficient robust model predictive control framework for uncertain nonlinear systems. *IEEE Transactions on Automatic Control*, 66(2), 794–801.
- Köhler, J. and Zeilinger, M.N. (2022). Recursively feasible stochastic predictive control using an interpolating initial state constraint. *IEEE Control Systems Letters*, 6, 2743–2748.
- Lazar, M. and Tetteroo, M. (2018). Computation of terminal costs and sets for discrete-time nonlinear MPC. *IFAC-PapersOnLine*, 51(20), 141–146.
- Liu, C., Gray, A., Lee, C., Hedrick, J.K., and Pan, J. (2014). Nonlinear stochastic predictive control with unscented transformation for semi-autonomous vehicles. In *2014 American Control Conference*, 5574–5579.
- McAllister, R.D. and Rawlings, J.B. (2021). Stochastic exponential stability of nonlinear stochastic model predictive control. In *2021 60th Conference on Decision and Control (CDC)*, 880–885.
- Mesbah, A. (2016). Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6), 30–44.
- Paulson, J.A., Streif, S., and Mesbah, A. (2015). Stability for receding-horizon stochastic model predictive control. In *2015 American Control Conference (ACC)*, 937–943.
- Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model Predictive Control: Theory, Computation, and Design*, volume 2. Nob Hill Publishing Madison, WI.
- Russo, G., di Bernardo, M., and Slotine, J.J. (2011). Contraction theory for systems biology. *Design and Analysis of Biomolecular Circuits: Engineering Approaches to Systems and Synthetic Biology*, 93–114.
- Saw, J.G., Yang, M.C., and Mo, T.C. (1984). Chebyshev inequality with estimated mean and variance. *The American Statistician*, 38(2), 130–132.
- Völz, A. and Graichen, K. (2015). Stochastic model predictive control of nonlinear continuous-time systems using the unscented transformation. In *2015 European Control Conference (ECC)*, 3365–3370.
- Weissel, F., Huber, M.F., and Hanebeck, U.D. (2009). Stochastic nonlinear model predictive control based on gaussian mixture approximations. In *Informatics in control, automation and robotics*, 239–252. Springer.
- Yu, S., Maier, C., Chen, H., and Allgöwer, F. (2013). Tube MPC scheme based on robust control invariant set with application to Lipschitz nonlinear systems. *Systems & Control Letters*, 62(2), 194–200.