

Adaptive Stochastic MPC under Unknown Noise Distribution

Charis Stamouli

Anastasios Tsiamis

Manfred Morari

George J. Pappas

University of Pennsylvania, Philadelphia, PA

STAMOULI@SEAS.UPENN.EDU

ATSIAMIS@SEAS.UPENN.EDU

MORARI@SEAS.UPENN.EDU

PAPPASG@SEAS.UPENN.EDU

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

In this paper, we address the stochastic MPC (SMPC) problem for linear systems, subject to chance state constraints and hard input constraints, under unknown noise distribution. First, we reformulate the chance state constraints as deterministic constraints depending only on explicit noise statistics. Based on these reformulated constraints, we design a distributionally robust and robustly stable benchmark SMPC algorithm for the ideal setting of known noise statistics. Then, we employ this benchmark controller to derive a novel robustly stable adaptive SMPC scheme that learns the necessary noise statistics online, while guaranteeing time-uniform satisfaction of the unknown reformulated state constraints with high probability. The latter is achieved through the use of confidence intervals which rely on the empirical noise statistics and are valid uniformly over time. Moreover, control performance is improved over time as more noise samples are gathered and better estimates of the noise statistics are obtained, given the online adaptation of the estimated reformulated constraints. Additionally, in tracking problems with multiple successive targets our approach leads to an online-enlarged domain of attraction compared to robust tube-based MPC. A numerical simulation of a DC-DC converter is used to demonstrate the effectiveness of the developed methodology.

Keywords: Stochastic MPC, Chance Constraints, Online Learning

1. Introduction

Model predictive control (MPC) has significantly matured in recent decades, establishing itself as the primary methodology for optimal control of multivariable constrained systems. However, model uncertainty poses a challenge to system performance under classical MPC frameworks, where a model description derived offline is used during control implementation. To mitigate the effects of uncertainties and improve system performance, various adaptive MPC strategies, which update the system model online, have been developed.

The majority of works on adaptive MPC focus on robust MPC, where hard constraints are enforced [Fukushima et al., 2007; Weiss and Di Cairano, 2014; Tanaskovic et al., 2014; Desaraju et al., 2017]. While well established theoretically, robust MPC can be excessively conservative in practice, given its worst-case treatment of the noise which often accounts for unlikely uncertainty realizations. Stochastic MPC (SMPC) reduces this conservatism and increases control authority by allowing constraint violations via the use of chance constraints. Examples of applications where SMPC has been successfully used in the past include building climate control [Ma et al., 2012], autonomous vehicle control [Carvalho et al., 2014], and wind turbine control [Kou et al., 2016].

In the realm of stochastic model predictive control, several approaches relying on the reformulation of chance constraints into deterministic counterparts have been proposed in both the classical [Korda et al., 2011; Paulson et al., 2019] and the adaptive [Bujarbaruah et al., 2021; Muñoz-Carpintero et al., 2018] regime. The main limitation of these algorithms is the restriction to special noise distributions, since analytic constraint reformulations involve the complex computation of multivariate integrals. The authors in [Kouvaritakis et al., 2010] overcome this challenge through the use of numerical approximations based on the noise distribution. However, since the distribution is itself subject to uncertainty, assuming perfect knowledge of it is unrealistic, and in fact, even small errors can dramatically change the control inputs selected by a model predictive controller [Van Parys et al., 2016]. On the other hand, if the noise distribution is assumed to be unknown, the exact expression of the chance constraints is also unknown.

Recent work in [Santos et al., 2020] addresses the above issue by replacing the true noise distribution with the empirical one and using confidence intervals to bound the uncertainty in the deterministically-reformulated chance constraints. The algorithm learns the reformulated constraints online and ensures satisfaction of the unknown chance constraints at each time step with high probability. Notice though that given such per-time-step high-confidence guarantees, the probability of time-uniform satisfaction of the chance constraints decays with system operation time.

An alternative to chance constraints reformulation that can be employed under unknown noise distribution is provided by scenario-based SMPC [Calafiore and Campi, 2006; Schildbach et al., 2012, 2014; Lorenzen et al., 2017a]. In scenario SMPC formulations, hard constraints are imposed for a finite number of noise samples (scenarios), which are drawn either empirically or through a random number generator. Nevertheless, unlike in [Santos et al., 2020], high-confidence satisfaction of the unknown chance constraints at each time step is ensured only for a sufficiently large number of noise samples.

In this paper, we address the problem of stochastic MPC for linear systems, subject to chance state constraints and hard input constraints, under unknown noise distribution. Our goal is to learn and adapt to the unknown noise online so that control performance is gradually improved, while also maintaining safety at all time steps with high probability. Our contributions are the following:

- i) For the ideal setting of known noise statistics we propose a distributionally robust benchmark SMPC scheme which also guarantees recursive feasibility and robust stability. This is achieved by replacing the chance state constraints with novel distribution-free deterministic constraints relying only on explicit noise statistics, in contrast to [Kouvaritakis et al., 2010; Lorenzen et al., 2017b], where distribution-dependent constraints are employed instead.
- ii) For the actual setting of unknown noise statistics we introduce a novel adaptive SMPC algorithm that estimates the necessary noise statistics and learns the reformulated constraints using noise samples collected online. Recursive feasibility and robust stability are ensured, while control performance is improved over time through adaptation. Additionally, in tracking problems with multiple successive targets our scheme leads to an online-enlarged domain of attraction compared to robust tube-based MPC.
- iii) By exploiting time-uniform confidence intervals based on the empirical noise statistics, we bound the uncertainty in the reformulated constraints uniformly over time with high probability. Note that in contrast to [Santos et al., 2020], this probability is constant and does not decrease with system operation time. To the best of our knowledge, this is the first work in the SMPC literature that guarantees satisfaction of unknown chance state constraints at all times steps jointly with high confidence.

We exhibit the benefits of the proposed scheme through numerical simulations. All proofs can be found in [Stamouli et al., 2022, Appendix A].

2. Problem Formulation

Consider a discrete-time linear time-invariant system with additive noise, described by the following state-space model:

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

where $x_t \in \mathbb{R}^n$ is the state, $u_t \in \mathbb{R}^m$ is the control input, and A , B are system matrices of appropriate dimensions. The initial state x_0 is assumed to be deterministic and fixed. Moreover, we assume that the process noise $w_t \in \mathbb{R}^n$ is i.i.d. with unknown probability distribution supported on the hypercube $\mathcal{W} = \{w \in \mathbb{R}^n \mid \|w\|_\infty \leq a_w\}$, where a_w is a known positive noise bound.

Assumption 1 *The noise distribution is unknown. The system matrices A , B and the noise bound a_w are known. The noise mean is also known and equal to zero, that is, $\mu_w := \mathbb{E}w_t = 0$.*

Remark 1 *An extension for unknown noise mean μ_w is presented in [Stamouli et al., 2022, Appendix B]. Note that in practice the noise bound a_w could be a rough estimate of the noise components' distribution support that we know from system structure or have obtained from observations. Our results can easily be extended to the case of general polytopes \mathcal{W} that contain the origin.*

The state is assumed to be subject to chance constraints of the form:

$$\mathbb{P}(g_j^T x_t \leq h_j) \geq 1 - \alpha_j, \quad (2)$$

for $j = 1, \dots, n_x$, whereas the control input should satisfy hard constraints of the form:

$$u_t \in \mathcal{U}, \quad (3)$$

for all $t \geq 0$, where $\alpha_j \in (0, 1)$ is the maximum admissible probability of violating the constraint $g_j^T x_t \leq h_j$, and $\mathcal{U} \subseteq \mathbb{R}^m$ is a polytope containing the origin. The probabilities α_j are specified directly from application requirements. Note that given the assumption that the noise distribution is unknown, the exact expression of the chance state constraints (2) is also unknown. Hence, herein we refer to (2) as *unknown* chance state constraints. Meanwhile, we consider the control performance objective determined by the cost function:

$$\ell(x_t, u_t) = x_t^T Q x_t + u_t^T R u_t,$$

where $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$ are positive-definite matrices. Notice that using the state measurements we can sequentially and online collect the noise samples w_0, \dots, w_{t-1} , as we control the system up to time step t .

The goal of this work is to develop, under the above assumptions, a recursively feasible and robustly stable adaptive SMPC algorithm for system (1). In particular, the algorithm should respect the input constraints (3) as well as certain state constraints which are designed online so that:

- i) control performance is improved over time, as more noise samples are collected and better estimates of explicit noise statistics are obtained;
- ii) time-uniform satisfaction of (2) (i.e., satisfaction of (2) at all time steps) is guaranteed with high confidence.

The benchmark SMPC scheme introduced in Section 4 for the ideal case of known noise statistics is used as reference for evaluating the control performance of our adaptive SMPC scheme (see Section 5) for unknown noise statistics.

3. Robust Tube-based MPC

In this section, we briefly review the classical robust tube-based MPC [Mayne et al., 2005], which we leverage in the design of our benchmark SMPC algorithm to provide recursive feasibility and robust stability guarantees (see Section 4 for details). The main difference between our framework and robust tube-based MPC lies in the formulation of the state constraints. Particularly, the authors in [Mayne et al., 2005] consider hard state constraints of the form:

$$x_t \in \mathcal{X} = \{x \in \mathbb{R}^n \mid g_j^T x \leq h_j, j = 1, \dots, n_x\} \quad (4)$$

instead of chance state constraints of the form (2). In addition, they only assume that the noise distribution is bounded, thus adopting a worst-case noise model, in contrast to our approach where a stochastic noise model is employed instead.

Let \bar{x}_t denote the nominal (noise-free) state, $\bar{u}_t = u_t(\bar{x}_t)$ the corresponding nominal input, and $e_t = x_t - \bar{x}_t$ the difference between the actual and the nominal state. We select gains $K, \bar{K} \in \mathbb{R}^{m \times n}$ such that the system matrices $A_{cl} := A + BK, \bar{A}_{cl} := A + B\bar{K}$ are stable. For any vector z , we use $z_{t+k|t}$ to denote the k -step ahead prediction of z made at time step t by some model predictive controller (i.e., the prediction of z_{t+k} made at time step t). At each time step t the robust tube-based MPC solves the following finite horizon optimal control problem:

$$\begin{aligned} \min_{\substack{\bar{x}_{t|t}, \dots, \bar{x}_{t+N|t} \\ \bar{u}_{t|t}, \dots, \bar{u}_{t+N-1|t}}} & \bar{x}_{t+N|t}^T P \bar{x}_{t+N|t} + \sum_{k=0}^{N-1} \ell(\bar{x}_{t+k|t}, \bar{u}_{t+k|t}) \\ \text{s.t.} & \bar{x}_{t+k+1|t} = A \bar{x}_{t+k|t} + B \bar{u}_{t+k|t} \\ & \bar{x}_{t+k|t} \in \bar{\mathcal{X}}_R, k = 1, \dots, N-1 \\ & \bar{u}_{t+k|t} \in \mathcal{U} \ominus K\mathcal{E}, k = 0, \dots, N-1 \\ & \bar{x}_{t+N|t} \in \bar{\mathcal{X}}_{f,R} \\ & x_t \in \bar{x}_{t|t} \oplus \mathcal{E}, \end{aligned} \quad (5)$$

where $\bar{x}_{t+k|t}, \bar{u}_{t+k|t}$ denote the prediction of $\bar{x}_{t+k}, \bar{u}_{t+k}$ made at time step t , respectively, \mathcal{E} is the minimum robust invariant set [Raković et al., 2004, Definition 3] for the error system:

$$e_{t+k+1|t} = A_{cl} e_{t+k|t} + w_{t+k|t}, \quad (6)$$

and $\bar{\mathcal{X}}_R = \mathcal{X} \ominus \mathcal{E}$. The set $\bar{\mathcal{X}}_{f,R}$ is the maximum invariant set [Kerrigan and Maciejowski, 2000, Definition 2.2] for the system $\bar{x}^+ = \bar{A}_{cl} \bar{x}$ that satisfies the constraints $\bar{\mathcal{X}}_{f,R} \subseteq \bar{\mathcal{X}}_R$ and $\bar{K} \bar{\mathcal{X}}_{f,R} \subseteq \mathcal{U} \ominus K\mathcal{E}$, and P is the unique positive definite solution of the Lyapunov equation:

$$\bar{A}_{cl}^T P \bar{A}_{cl} - P = -Q - \bar{K}^T R \bar{K}. \quad (7)$$

After we obtain the optimal sequences $\{\bar{x}_{t|t}^*, \dots, \bar{x}_{t+N|t}^*\}$ and $\{\bar{u}_{t|t}^*, \dots, \bar{u}_{t+N-1|t}^*\}$ of (5), the tube-based state-feedback control policy:

$$u_t(x_t) = K \left(x_t - \bar{x}_{t|t}^*(x_t) \right) + \bar{u}_{t|t}^*(x_t) \quad (8)$$

is applied to system (1). Notice in (8) that the first term on the right-hand side keeps the actual state close to the nominal state, whereas the second term steers the nominal state to the origin. At the

next time step we repeat the same process, thus yielding a receding horizon strategy. Employing (6), one can write the halfspace inequalities determining the nominal state constraints of (5) as follows:

$$g_j^T \bar{x}_{t+k|t} \leq h_j - \gamma_j - \|g_j\|_1 a_w, \quad (9)$$

where $\gamma_j = \max_{e \in \mathcal{E}} g_j^T A_{cl} e$, for all $j = 1, \dots, n_x$. In the next section, stochasticity is employed to relax these constraints, thus reducing the conservatism of robust MPC [Mayne et al., 2005].

4. Tube-based SMPC under Known Noise Statistics

In this section, we consider the ideal setting where the noise covariance $\Sigma_w := \mathbb{E}w_t w_t^T$ is known and introduce a stochastic model predictive controller with polytopic state constraints that depend only on Σ_w and are sufficient for the original state constraints (2). Herein, we refer to this controller as *benchmark SMPC*. This is because the performance of this controller will be used as reference for evaluating the performance of our adaptive SMPC scheme, which is introduced in the next section for the actual setting where Σ_w is unknown. More specifically, inspired by the idea of constraint tightening employed in robust tube-based MPC, we derive a less conservative robustly stable MPC algorithm, by imposing suitable mixed probabilistic/worst-case constraint tightening to (2) based on Σ_w . As we show next, our approach leads to a relaxation of the nominal state constraints (9) of the robust tube-based MPC, where worst-case constraint tightening was instead applied to (4).

Lemma 1 *Given system (1) and a prediction horizon of length N , at each time step t the k -step ahead prediction of the actual state $x_{t+k|t}$ satisfies:*

$$\mathbb{P}(g_j^T x_{t+k|t} \leq h_j) \geq 1 - \alpha_j,$$

if $x_t \in \bar{x}_{t|t} \oplus \mathcal{E}$ and the k -step ahead prediction of the nominal state $\bar{x}_{t+k|t}$ satisfies:

$$g_j^T \bar{x}_{t+k|t} \leq h_j - \gamma_j - f(\alpha_j) \sqrt{g_j^T \Sigma_w g_j}, \quad (10)$$

with γ_j as in (9) and $f(\alpha_j) = \sqrt{\frac{1-\alpha_j}{\alpha_j}}$, for all $j = 1, \dots, n_x$ and $k = 1, \dots, N$.

We want to employ the constraints (10) to develop a simple stochastic tube-based MPC algorithm that is less conservative than the robust tube-based MPC. This means that (10) should be looser than (9), that is, we would like to have $f(\alpha_j) \sqrt{g_j^T \Sigma_w g_j} \leq \|g_j\|_1 a_w$, for all $\alpha_j \in (0, 1)$. However, it is clear from the form of $f(\cdot)$ that this inequality may not be satisfied for small values of α_j . Hence, in order to prevent excessive conservatism, we replace the constraints (10) with:

$$g_j^T \bar{x}_{t+k|t} \leq h_j - \gamma_j - \min \left\{ \|g_j\|_1 a_w, f(\alpha_j) \sqrt{g_j^T \Sigma_w g_j} \right\}. \quad (11)$$

Remark 2 *Notice in (11) that the first member of the minimum (robust) will be active for small values of α_j and the second one (stochastic) otherwise. This is reasonable, because selecting very small maximum admissible probabilities of constraint violation α_j is approximately equivalent to imposing hard state constraints, which are guaranteed to be satisfied due to the robust member of the minimum.*

Let $\bar{\mathcal{X}}_S$ denote the polytope defined by the halfspace inequalities (11), and $\bar{\mathcal{X}}_{f,S}$ the maximum invariant set for the system $\bar{x}^+ = \bar{A}_{cl}\bar{x}$ that satisfies the constraints $\bar{\mathcal{X}}_{f,S} \subseteq \bar{\mathcal{X}}_S$ and $\bar{K}\bar{\mathcal{X}}_{f,S} \subseteq \mathcal{U} \ominus K\mathcal{E}$. At each time step t our benchmark SMPC solves the following optimal control problem:

$$\begin{aligned}
 \min_{\substack{\bar{x}_{t|t}, \dots, \bar{x}_{t+N|t} \\ \bar{u}_{t|t}, \dots, \bar{u}_{t+N-1|t}}} & \bar{x}_{t+N|t}^T P \bar{x}_{t+N|t} + \sum_{k=0}^{N-1} \ell(\bar{x}_{t+k|t}, \bar{u}_{t+k|t}) \\
 \text{s.t.} & \bar{x}_{t+k+1|t} = A\bar{x}_{t+k|t} + B\bar{u}_{t+k|t} \\
 & \bar{x}_{t+k|t} \in \bar{\mathcal{X}}_S, \quad k = 1, \dots, N-1 \\
 & \bar{u}_{t+k|t} \in \mathcal{U} \ominus K\mathcal{E}, \quad k = 0, \dots, N-1 \\
 & \bar{x}_{t+N|t} \in \bar{\mathcal{X}}_{f,S} \\
 & x_t \in \bar{x}_{t|t} \oplus \mathcal{E}
 \end{aligned} \tag{12}$$

and then the state-feedback control law (8) is applied to system (1).

Theorem 1 (Benchmark SMPC) *Suppose the noise covariance Σ_w is known and the optimization problem (12) is feasible at time step $t = 0$. Then, (12) is feasible and ensures satisfaction of the original chance state constraints (2) at every time step $t \geq 0$. Moreover, system (1) in closed-loop with the model predictive controller defined by (12) and (8) asymptotically converges to the set \mathcal{E} for all noise realizations.*

Remark 3 *The problem of stochastic MPC with stability guarantees was addressed in [Kouvaritakis et al., 2010], where distribution-dependent reformulated state constraints were employed. In contrast, the reformulated state constraints (11) of our benchmark SMPC scheme are distribution-free in the sense that they guarantee satisfaction of the original state constraints (2) for all noise distributions with zero mean and covariance Σ_w . Employing these constraints, we were able to derive a distributionally robust SMPC algorithm, which also enjoys stability guarantees.*

5. Adaptive Tube-based SMPC under Unknown Noise Statistics

In the previous section, we introduced a simple tube-based SMPC algorithm assuming that the noise covariance Σ_w is known. We are now ready to present an adaptive SMPC algorithm for unknown Σ_w that learns the nominal state constraints (11) of the benchmark SMPC online and guarantees their satisfaction at all time steps with high probability. The main idea is to use noise samples to estimate the term $g_j^T \Sigma_w g_j$ of (11) online, and provide explicit confidence intervals for these estimates so that (11) is always satisfied. As more samples are gathered, the confidence intervals become smaller, thus relaxing the constraints and enhancing the control performance over time.

We estimate Σ_w using the standard sample covariance, which is defined as $\hat{\Sigma}_w^t = \frac{1}{t} \sum_{i=0}^{t-1} w_i w_i^T$, at each time step $t \geq 1$. The noise samples w_0, \dots, w_{t-1} can be computed as $w_i = x_{i+1} - Ax_i - Bu_i$ for all $t \geq 1$. To obtain a confidence interval for the unknown term $g_j^T \Sigma_w g_j$ of (11) based on $\hat{\Sigma}_w^t$, we employ the boundedness of the noise, which allows us to derive the following lemma.

Lemma 2 *Fix a failure probability $\delta \in (0, 1)$. With probability at least $1 - \delta$, the sample covariance $\hat{\Sigma}_w^t$ satisfies:*

$$g_j^T \Sigma_w g_j \leq g_j^T \hat{\Sigma}_w^t g_j + r_{tj}(\delta), \tag{13}$$

where $r_{tj}(\delta) = \|g_j\|_1^2 a_w^2 \sqrt{0.5t^{-1} \log((\pi^4 t^2 n_x^2)/(36\delta))}$, uniformly over all $t \geq 1$ and $j = 1, \dots, n_x$.

Notice in (13) that the bound $g_j^T \widehat{\Sigma}_w^t g_j + r_{tj}(\delta)$, where δ is a design parameter determining its confidence, converges to $g_j^T \Sigma_w g_j$ as fast as $\mathcal{O}\left(\sqrt{t^{-1} \log t}\right)$. From Lemma 2 we deduce that a sufficient condition for the constraints (11) of the benchmark SMPC to be satisfied for all $t \geq 1$ with probability at least $1 - \delta$ is the following:

$$g_j^T \bar{x}_{t+k|t} \leq h_j - \gamma_j - \min \left\{ \|g_j\|_1 a_w, f(\alpha_j) \sqrt{g_j^T \widehat{\Sigma}_w^t g_j + r_{tj}(\delta)} \right\}. \quad (14)$$

Given that it suffices to update the nominal state constraint set only when the constraints (14) become looser, we define $\bar{\mathcal{X}}_S^t$ as the intersection of the halfspaces determined by the inequalities:

$$g_j^T \bar{x}_{t+k|t} \leq h_j - \gamma_j - \min \left\{ \|g_j\|_1 a_w, \min_{1 \leq \tau \leq t} f(\alpha_j) \sqrt{g_j^T \widehat{\Sigma}_w^\tau g_j + r_{\tau j}(\delta)} \right\}, \quad (15)$$

for each $t \geq 1$. Moreover, we define $\bar{\mathcal{X}}_{f,S}^t$ as the maximum invariant set for the system $\bar{x}^+ = \bar{A}_{cl} \bar{x}$ that satisfies the constraints $\bar{\mathcal{X}}_{f,S}^t \subseteq \bar{\mathcal{X}}_S^t$ and $\bar{K} \bar{\mathcal{X}}_{f,S}^t \subseteq \mathcal{U} \ominus K\mathcal{E}$, for each $t \geq 1$. Since the first noise sample is computed at time step $t = 1$, we set $\bar{\mathcal{X}}_S^0 = \bar{\mathcal{X}}_R$ and $\bar{\mathcal{X}}_{f,S}^0 = \bar{\mathcal{X}}_{f,R}$. At each time step t our adaptive SMPC algorithm solves the following finite horizon optimal control problem:

$$\begin{aligned} \min_{\substack{\bar{x}_{t|t}, \dots, \bar{x}_{t+N|t} \\ \bar{u}_{t|t}, \dots, \bar{u}_{t+N-1|t}}} & \bar{x}_{t+N|t}^T P \bar{x}_{t+N|t} + \sum_{k=0}^{N-1} \ell(\bar{x}_{t+k|t}, \bar{u}_{t+k|t}) \\ \text{s.t.} & \bar{x}_{t+k+1|t} = A \bar{x}_{t+k|t} + B \bar{u}_{t+k|t} \\ & \bar{x}_{t+k|t} \in \bar{\mathcal{X}}_S^t, \quad k = 1, \dots, N-1 \\ & \bar{u}_{t+k|t} \in \mathcal{U} \ominus K\mathcal{E}, \quad k = 0, \dots, N-1 \\ & \bar{x}_{t+N|t} \in \bar{\mathcal{X}}_{f,S}^t \\ & x_t \in \bar{x}_{t|t} \oplus \mathcal{E} \end{aligned} \quad (16)$$

and then the state-feedback controller (8) is applied to system (1).

Theorem 2 (Adaptive SMPC) *Fix a failure probability $\delta \in (0, 1)$ and suppose the optimization problem (16) is feasible at time step $t = 0$. Then, (16) is feasible at every time step $t \geq 0$, and ensures satisfaction of the unknown original chance state constraints (2) uniformly over all $t \geq 0$ with probability at least $1 - \delta$. Moreover, system (1) in closed-loop with the model predictive controller defined by (16) and (8) asymptotically converges to the set \mathcal{E} for all noise realizations.*

Remark 4 *Notice that the difference between the robust MPC, the benchmark SMPC and the adaptive SMPC lies in the nominal state and terminal constraints. We know that $\widehat{\Sigma}_w^t$ is an asymptotically consistent estimate of Σ_w , and we have $r_{tj}(\delta) = \mathcal{O}\left(\sqrt{t^{-1} \log t}\right) \xrightarrow{t \rightarrow \infty} 0$, for all $j = 1, \dots, n_x$. Hence, the nominal state constraints (15) of (16) converge to the nominal state constraints (11) of (12) as $t \rightarrow \infty$. Moreover, from our theoretical analysis we know that $\bar{\mathcal{X}}_R = \bar{\mathcal{X}}_S^0 \subseteq \dots \subseteq \bar{\mathcal{X}}_S^t$, for all $t \geq 1$, and $\bigcup_{t=0}^{\infty} \bar{\mathcal{X}}_S^t \subseteq \bar{\mathcal{X}}_S$ with high probability. The above observations are illustrated in Figure 1. We conclude that both the benchmark and the adaptive SMPC are less conservative than the robust MPC, and in addition, the adaptive SMPC ensures satisfaction of the benchmark SMPC's nominal state constraints at all time steps with high confidence, while learning them online.*

Note that, herein, conservatism is related to the size of the feasible set of each MPC problem in the sense that the larger the feasible set, the smaller the cost and thus the better the performance of a controller. For this reason, the performance of the adaptive SMPC is expected to be worse than the one of the benchmark SMPC and converge to it as $t \rightarrow \infty$.

Remark 5 (Tracking of multiple successive targets) Up to now, we only studied the problem of tracking the origin $x = 0$. However, our adaptive SMPC scheme can easily be extended to track non-zero targets. This can be particularly useful in a tracking scenario with multiple successive targets. More specifically, notice that every time we rerun our adaptive SMPC to track a new target, the feasible set of problem (16) is larger and larger than the one of the robust MPC problem (5), which implies an online enlargement of the domain of attraction. In Section 6, we validate this observation through an illustrative example.

6. Case Study: DC-DC Converter

In this section, we illustrate the efficacy of the proposed adaptive SMPC strategy using a simplified model of a DC-DC converter [Lorenzen et al., 2015]. Particularly, we consider a linearized DC-DC converter model of the form (1) with:

$$A = \begin{bmatrix} 1 & 0.0075 \\ -0.143 & 0.996 \end{bmatrix}, B = \begin{bmatrix} 4.798 \\ 0.115 \end{bmatrix},$$

and zero-mean Gaussian noise with covariance $\Sigma = 0.25\mathbb{I}$ truncated at $a_w = 1$. Suppose the state is subject to:

$$\begin{aligned} \mathbb{P}([1 \ 0] x_t \leq 9.6) &\geq 0.2, \quad \mathbb{P}([-1 \ 0] x_t \leq 9.6) \geq 0.2, \\ \mathbb{P}([0 \ 1] x_t \leq 9.6) &\geq 0.2, \quad \mathbb{P}([0 \ -1] x_t \leq 9.6) \geq 0.2, \end{aligned}$$

and the input to hard constraints given by $|u_t| \leq 5$. We wish to minimize the cost for $Q = \text{diag}(1, 10)$ and $R = 1$ over a prediction horizon of length $N = 9$. The gain matrices were selected as $K = [-0.3429 \ 0.8629]$ and $\bar{K} = [-0.2858 \ 0.4910]$, and the failure probability of time-uniform satisfaction of (11) for the adaptive SMPC as $\delta = 0.1$.

Figure 2 shows the nominal state and terminal constraint sets of the robust MPC, the benchmark SMPC, and the adaptive SMPC for 10, 100 and 1000 noise samples. Notice that $\bar{\mathcal{X}}_R \subseteq \bar{\mathcal{X}}_S^{10} \subseteq \bar{\mathcal{X}}_S^{100} \subseteq \bar{\mathcal{X}}_S^{1000} \subseteq \bar{\mathcal{X}}_S$ and $\bar{\mathcal{X}}_{f,R} \subseteq \bar{\mathcal{X}}_{f,S}^{10} \subseteq \bar{\mathcal{X}}_{f,S}^{100} \subseteq \bar{\mathcal{X}}_{f,S}^{1000} \subseteq \bar{\mathcal{X}}_{f,S}$, which indicates that the robust MPC is more conservative than the adaptive SMPC, which is more conservative than the benchmark SMPC. Hence, the time-uniform high-confidence guarantees of our adaptive SMPC satisfying the unknown chance state constraints (2) are validated. Moreover, we observe that the sets $\bar{\mathcal{X}}_S^{1000}$ and $\bar{\mathcal{X}}_{f,S}^{1000}$ of the adaptive SMPC are very close to the sets $\bar{\mathcal{X}}_S$ and $\bar{\mathcal{X}}_{f,S}$, respectively, of the benchmark SMPC, in accordance with our theoretical result that the nominal state and terminal constraint sets of the adaptive SMPC asymptotically converge to the corresponding sets of the benchmark SMPC.

To demonstrate the benefits of our adaptive SMPC scheme, we compare it with the robust MPC and the benchmark SMPC in a two-target tracking scenario for initial state $x_0 = [6 \ 4]^T$. More specifically, suppose we want to employ the DC-DC converter in a DC motor drive to track $x_{\text{target}}^{(1)} =$




Fig. 1: High-confidence sketch illustration of the subset relationship between $\bar{\mathcal{X}}_R$, $\bar{\mathcal{X}}_S$, $\bar{\mathcal{X}}_S^i$, and $\bar{\mathcal{X}}_S^t$, for some $t \geq 2$. The set $\bar{\mathcal{X}}_S^t$ converges to $\bar{\mathcal{X}}_S$ as $t \rightarrow \infty$.




Fig. 2: (a) Comparison of the nominal state constraint sets of the robust MPC ($\bar{\mathcal{X}}_R$), the benchmark SMPC ($\bar{\mathcal{X}}_S$), and the adaptive SMPC for 10, 100 and 1000 noise samples ($\bar{\mathcal{X}}_S^{10}$, $\bar{\mathcal{X}}_S^{100}$, $\bar{\mathcal{X}}_S^{1000}$). (b) Comparison of the terminal constraint sets of the robust MPC ($\bar{\mathcal{X}}_{f,R}$), the benchmark SMPC ($\bar{\mathcal{X}}_{f,S}$), and the adaptive SMPC for 10, 100 and 1000 noise samples ($\bar{\mathcal{X}}_{f,S}^{10}$, $\bar{\mathcal{X}}_{f,S}^{100}$, $\bar{\mathcal{X}}_{f,S}^{1000}$).

$[0 \ 0]^T$ until time step $t = 100$, when the target changes to $x_{\text{target}}^{(2)} = [-0.1169 \ 3.9983]^T$. In order to track the two consecutive desired states, we sequentially apply each MPC algorithm twice: i) with target $x_{\text{target}}^{(1)}$, for $t = 0, \dots, 99$ (Phase 1), and ii) with target $x_{\text{target}}^{(2)}$, for $t = 100, 101, \dots$ (Phase 2). We will evaluate the performance of our adaptive SMPC scheme using the cumulative cost $C(t_0, t_0 + T) = \sum_{t=t_0}^{t_0+T-1} \ell(x_t, u_t)$, where t_0 is the initial time, T the running time, and x_t, u_t the state and the input, respectively, of the system in closed-loop with a model predictive controller.

In simulations we observed that robust MPC is feasible in Phase 1 and infeasible in Phase 2, in contrast to our benchmark and adaptive SMPC schemes, which are feasible in both phases. This indicates the enlarged domain of attraction of problems (12) and (16) compared to the one of problem (5) in Phase 2, by the beginning of which 100 noise samples have been gathered online (see Remark 5). Figure 3 illustrates the state evolution of the system under the benchmark and the adaptive SMPC schemes for 30 noise sequences $\{w_0, \dots, w_{118}\}$, as well as part of the boundary of the hard state constraint set \mathcal{X} . We observe that the closed-loop trajectories of the two schemes are more and more similar over time (in Phase 2, at the beginning of which 100 noise samples are available, they are almost indistinguishable), and in each phase the system converges to a small neighborhood of the corresponding target, thus verifying our robust stability guarantees. Notice also that the hard constraints $x_{t,1} \leq 9.6$ and $x_{t,1} \geq -9.6$ are violated in Phase 1 and 2, respectively.

	Violation of $x_{t,1} \leq 9.6$ in Phase 1 (%)	Violation of $x_{t,1} \geq -9.6$ in Phase 2 (%)	$C(0, 6)$	$C(100, 106)$
Benchmark SMPC	6.24	5.60	515.73	788.19
Adaptive SMPC	0.32	3.92	518.62	788.50

Table 1: Average constraint violation and cumulative cost over 10^3 noise sequences.

Simulating the system with the benchmark and the adaptive SMPC for 10^3 noise sequences, we computed the average constraint violations and cumulative costs shown in Table 1. Note that these values were calculated only for the first 7 time steps of each phase, which was the time needed for the system to converge. Notice that in both phases the average constraint violation is larger and the




Fig. 3: State evolution of the DC-DC converter system under the benchmark SMPC (blue) and the adaptive SMPC (green) for initial state $x_0 = [6 \ 4]^T$ and 30 noise sequences. The target state is $x_{\text{target}}^{(1)} = [0 \ 0]^T$ until $t = 100$, when it is changed to $x_{\text{target}}^{(2)} = [-0.1169 \ 3.9983]^T$. The trajectories move clockwise from $(6, 4)$ to $(0, 0)$, and then to $(-0.1169, 3.9983)$. The red lines belong to the boundary of the hard state constraint set \mathcal{X} . Robust MPC is not depicted as it becomes infeasible when the target changes to $x_{\text{target}}^{(2)}$.

average cumulative cost is smaller in the case of benchmark SMPC. This can be explained by the fact that in each phase the benchmark SMPC problem (12) has looser nominal state and terminal constraints than the adaptive SMPC problem (16), and hence a larger feasible set, which allows the benchmark controller to violate the hard state constraints more often and achieve smaller cost than the adaptive one. However, in Phase 2, by the beginning of which 100 noise samples have been collected online, the difference in the average constraint violation and cumulative cost of the benchmark and the adaptive SMPC is much smaller than in Phase 1. This was expected given our theoretical remark that the performance of the adaptive controller improves over time, as more noise samples are gathered and better estimates of the constraints (11) are obtained, gradually converging to the performance of the benchmark controller (see Remark 4).

7. Conclusion and Future Work

In this paper, we studied the problem of stochastic MPC (SMPC) for constrained linear systems under unknown noise distribution. We developed the first robustly stable adaptive SMPC scheme that guarantees time-uniform satisfaction of unknown chance state constraints with high probability, while learning their deterministic counterparts and thus improving control performance online.

Moving forward, our paper opens up several research directions. First, in order to reduce the conservatism introduced by the use of a tube-based control law with a precomputed stabilizing gain, we could try different parametrizations of the feedback control policy (e.g., a disturbance feedback parametrization with the gain as an optimization variable). Moreover, to mitigate possible conservatism resulting from the use of distributionally robust reformulated state constraints, we could try a new reformulation based on the quantile instead of the noise statistics, and provide time-uniform confidence intervals for the empirical quantile to design the new adaptive model predictive controller. Finally, our adaptive SMPC scheme could be used to analyze the regret of stochastic MPC under unknown noise distribution, which is an open problem.

Acknowledgement

This work was supported by the AFOSR under grant FA9550-19-1-0265 (Assured Autonomy in Contested Environments).

References

- M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli. Adaptive Stochastic MPC Under Time-Varying Uncertainty. *IEEE Transactions on Automatic Control*, 66(6):2840–2845, 2021.
- G. C. Calafiore and M. C. Campi. The Scenario Approach to Robust Control Design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.
- A. Carvalho, Y. Gao, S. Lefevre, and F. Borrelli. Stochastic predictive control of autonomous vehicles in uncertain environments. In *International Symposium on Advanced Vehicle Control*, pages 712–719, 2014.
- V. R. Desaraju, A. E. Spitzer, and N. Michael. Experience-driven Predictive Control with Robust Constraint Satisfaction under Time-Varying State Uncertainty. In *Proceedings of Robotics: Science and Systems*, 2017.
- H. Fukushima, T.-H. Kim, and T. Sugie. Adaptive model predictive control for a class of constrained linear systems based on the comparison model. *Automatica*, 43:301–308, 2007.
- E. C. Kerrigan and J. M. Maciejowski. Invariant sets for constrained nonlinear discrete-time systems with application to feasibility in model predictive control. In *39th IEEE Conference on Decision and Control*, pages 4951–4956, 2000.
- 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 1245–1251, 2011.
- P. Kou, D. Liang, and L. Gao. Stochastic model predictive control for wind turbines with doubly fed induction generators. In *IEEE Power and Energy Society General Meeting*, pages 1–5, 2016.
- B. Kouvaritakis, M. Cannon, S. V. Raković, and Q. Cheng. Explicit use of probabilistic distributions in linear predictive control. In *UKACC International Conference on Control*, pages 1–6, 2010.
- M. Lorenzen, F. Allgöwer, F. Dabbene, and R. Tempo. An improved constraint-tightening approach for Stochastic MPC. In *American Control Conference*, pages 944–949, 2015.
- M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer. Stochastic MPC with offline uncertainty sampling. *Automatica*, 81:176–183, 2017a.
- M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer. Constraint-Tightening and Stability in Stochastic Model Predictive Control. *IEEE Transactions on Automatic Control*, 62(7):3165–3177, 2017b.
- Y. Ma, S. Vichik, and F. Borrelli. Fast stochastic MPC with optimal risk allocation applied to building control systems. In *IEEE Conference on Decision and Control*, pages 7559–7564, 2012.

- 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.
- D. Muñoz-Carpintero, G. Hu, and C. J. Spanos. Stochastic Model Predictive Control with adaptive constraint tightening for non-conservative chance constraints satisfaction. *Automatica*, 96:32–39, 2018.
- J. A. Paulson, T. L. M. Santos, and A. Mesbah. Mixed stochastic-deterministic tube MPC for offset-free tracking in the presence of plant-model mismatch. *Journal of Process Control*, 83:102–120, 2019.
- S. V. Raković, E. C. Kerrigan, K. I. Kouramas, and D. Q. Mayne. Invariant approximations of robustly positively invariant sets for constrained linear discrete-time systems subject to bounded disturbances. Technical Report CUED/F-INFENG/TR.473, Department of Engineering, University of Cambridge, 2004.
- T. L. M. Santos, V. M. Cunha, and A. Mesbah. Stochastic Model Predictive Control with Adaptive Chance Constraints based on Empirical Cumulative Distributions. In *IFAC*, pages 11257–11263, 2020.
- G. Schildbach, G. C. Calafiore, L. Fagiano, and M. Morari. Randomized Model Predictive Control for Stochastic Linear Systems. In *IEEE American Control Conference*, pages 417–422, 2012.
- G. Schildbach, L. Fagiano, C. Frei, and M. Morari. The Scenario approach for Stochastic Model Predictive Control with bounds on closed-loop constraint violations. *Automatica*, 50(12):3009–3018, 2014.
- C. Stamouli, A. Tsiamis, M. Morari, and G. J. Pappas. Adaptive Stochastic MPC under Unknown Noise Distribution. *arXiv preprint arXiv:2204.01107*, 2022.
- M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari. Adaptive receding horizon control for constrained MIMO systems. *Automatica*, 50(12):3019–3029, 2014.
- B. P. G. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari. Distributionally Robust Control of Constrained Stochastic Systems. *IEEE Transactions on Automatic Control*, 61(2):430–442, 2016.
- A. Weiss and S. Di Cairano. Robust dual control MPC with guaranteed constraint satisfaction. In *IEEE Conference on Decision and Control*, pages 6713–6718, 2014.