Automatica 96 (2018) 32–39

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Stochastic Model Predictive Control with adaptive constraint tightening for non-conservative chance constraints satisfaction ☆

Diego Muñoz-Carpintero [a,*], Guoqiang Hu [b], Costas J. Spanos [c]

[a] Electrical Engineering Department, Universidad de Chile, Av. Tupper 2007, Santiago, Chile
[b] School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798, Singapore
[c] Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720, USA

ARTICLE INFO

Article history:
Received 1 February 2017
Received in revised form 2 April 2018
Accepted 24 May 2018

Keywords:
Predictive control
Stochastic control
Chance constraints
Adaptive control
Stochastic approximation

ABSTRACT

Most stochastic Model Predictive Control (MPC) formulations allow constraint violations via the use of chance constraints, thus increasing control authority and improving performance when compared to their robust MPC counterparts. However, common stochastic MPC methods handle chance constraints conservatively: constraint violations are often smaller than allowed by design, thus limiting the potential improvements in control performance. This is a consequence of enforcing chance constraints overlooking the past behavior of the system and/or of an over tightening of the constraints on the predicted sequences. This work presents a stochastic MPC strategy that uses the observed amount of constraint violations to adaptively scale the tightening parameters, thus eliminating the aforementioned conservativeness. It is proven using Stochastic Approximation that, under suitable conditions, the amount of constraint violations converges in probability when using the proposed method. The effectiveness and benefits of the approach are illustrated by a simulation example.

1. Introduction

Stochastic Model Predictive Control (SMPC) is a type of MPC which explicitly incorporates the stochasticity of the systems in the prediction models and cost functions, and allows occasional constraint violations via the use of chance constraints (Farina, Giulioni, & Scattolini, 2016; Mesbah, 2016). These features help to avoid the conservatism present in robust MPC methods due to unlikely realizations of the uncertainty (Kouvaritakis & Cannon, 2016).

A common type of SMPC approach consists in a deterministic reformulation of the stochastic programs by means of a constraint tightening that accounts for the uncertainty of the predictions. If the predicted nominal sequences (without uncertainty) satisfy the reformulated deterministic constraints, then the full predictions (affected by uncertainty) satisfy the chance constraints (Farina, Giulioni, Magni, & Scattolini, 2015; Kouvaritakis, Cannon, & Muñoz-Carpintero, 2013; Kouvaritakis, Cannon, Raković, & Cheng, 2010; Paulson, Streif, & Mesbah, 2015; Schildbach, Calafiore, Fagiano, & Morari, 2012; Van Hessem, Scherer, & Bosgra, 2001). Some methods assume that the distribution of the uncertainty is known and compute the tightening parameters by finding the inverse of the cumulative distribution function (Kouvaritakis et al., 2013, 2010; Van Hessem et al., 2001), and others use Chebyshev type inequalities to guarantee chance constraint satisfaction for any distribution (Farina, Giulioni, Magni, & Scattolini, 2013; Farina et al., 2015; Paulson et al., 2015). Scenario-based approaches impose the constraints for different scenarios of the uncertainty, thus providing probabilistic guarantees for the satisfaction of the chance constraints (Batina, 2004; Bernardini & Bemporad, 2009; Calafiore & Fagiano, 2013; Schildbach et al., 2012).

These approaches, however, handle chance constraints conservatively. Chance constraints are enforced pointwise-in-time, conditional to the current state and overlooking the past behavior of the system. The consequence of this is that in closed-loop constraints may be violated less than allowed by design (Korda, Gondhalekar, Oldewurtel, & Jones, 2012; Schildbach, Fagiano, Frei, & Morari, 2014), thus limiting the potential improvements in control authority and performance. This is the case even with recent scenario-based approaches that have reduced the conservativeness of early methods by means of sampling-and-discard strategies (Campi & Garatti, 2011; Prandini, Garatti, & Lygeros, 2012). Also, the scenario-based strategy of Schildbach et al. (2014),

☆ This work was supported by the National Research Foundation, Prime Minister's Office, Singapore under the Energy Innovation Research Programme (EIR) for Building Energy Efficiency Grant Call, administered by the Building and Construction Authority (NRF2013EWT-EIRP004-051), and by CONICYT-FONDECYT Postdoctorado-3170040. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Franco Blanchini under the direction of Editor Ian R. Petersen.

* Corresponding author.
E-mail addresses: dimunoz@ing.uchile.cl (D. Muñoz-Carpintero), gqhu@ntu.edu.sg (G. Hu), spanos@berkeley.edu (C.J. Spanos).

https://doi.org/10.1016/j.automatica.2018.06.026
0005-1098/© 2018 Published by Elsevier Ltd.

which reduces conservatism by abandoning the pointwise-in-time enforcement of chance constraints in favor of a re-interpretation as time-average constraints, ignores the past story of the system. Furthermore, even the pointwise-in-time enforcement of constraints may be itself conservative. For instance, constraint tightening using inequalities Chebyshev type inequalities is conservative for any distribution except for the worst case (Farina et al., 2015; Paulson et al., 2015).

Some approaches have been proposed to eliminate the conservativeness in constraint satisfaction in SMPC by explicitly exploiting the observed constraint violation level. A multi-layer setting with guaranteed recursive feasibility is proposed in Korda, Gondhalekar, Oldewurtel, and Jones (2014) and Korda et al. (2012). In this approach, the state moves through the layers depending on the observed amount of constraint violations, until this amount converges to the desired value in the top layer. In Oldewurtel, Sturzenegger, Esfahani, Andersson, Morari, and Lygeros (2013), an adaptive scheme that uses the past constraint violations to adaptively scale the tightening parameters is proposed. However, the conclusion for the convergence of the time-average of constraint violations is argued intuitively, rather than with a rigorous proof. Additionally, it is assumed that the control policy can decide at an arbitrary instant, with probability 1, if a constraint is violated or not. This assumption is too strong since there is no mechanism to enforce this by the MPC optimization.

The work presented in this paper also considers an SMPC strategy that exploits the observed constraint violations to adaptively scale the tightening parameters to eliminate the aforementioned conservativeness. The setting is similar to that of Oldewurtel et al. (2013), with mild differences in the adaptive law that is necessary for proving convergence. The main contribution of this work is a rigorous analysis of the SMPC strategy with adaptive scaling to eliminate the aforementioned conservativeness, but without the strong assumption of Oldewurtel et al. (2013), and using instead other assumptions that are commonly seen in practice. In this analysis, the tools of Stochastic Approximation (SA) (Kushner & Vázquez-Abad, 1996; Kushner & Yin, 2003) are used to prove convergence of the amount of constraint violations and the scaling of the tightening parameters to the desired level. The analysis also includes a discussion on the conditions required to obtain the proof of convergence.

We highlight that, unlike the method of Korda et al. (2014, 2012), our strategy does not use controllable sets (just like the method of Oldewurtel et al., 2013). Instead, it can be built upon many SMPC designs as it only requires the addition of a simple update law. Thus it is well suited to large systems, and is easy to implement, making it attractive for industrial adoption.

The structure of the paper is as follows. Section 2 introduces basic notations, definitions, and relevant tools of stochastic processes and SA. Section 3 defines the problem setting, a usual formulation of SMPC and discusses a motivation for this work. Section 4 presents the proposed SMPC strategy with adaptive tightening and Section 5 analyzes its convergence properties, and discusses possible extensions, the relevance of the main conditions and compares the strategy considered here with other methods. Finally, Section 6 presents numerical examples and Section 7 provides concluding remarks.

2. Mathematical preliminaries

2.1. Basic definitions and notation

The sets of non-negative, positive integers and non-negative reals are denoted by $\mathbb{N}, \mathbb{N}_+$ and $\mathbb{R}_+$. For a sequence $\mathbf{x} = \{x_k : k \in \mathbb{N}\}$, $x_{j|k}$ denotes the prediction of x_{k+j} made at time k. The Minkowski sum of sets $X, Y \subseteq \mathbb{R}^n$ is given by $X \oplus Y = \{x + y : x \in X, y \in Y\}$. Given a closed and convex set $\mathbf{X}$, the normal cone of $\mathbf{X}$ at x is $C_{\mathbf{X}}(x) = \{z : z^T(y - x) \le 0, \ \forall y \in \mathbf{X}\}$. (Note that if x is in the interior of $\mathbf{X}$ then trivially $C_{\mathbf{X}}(x) = \{0\}$, and $C_{\mathbf{X}}(x) = \emptyset$ if $x \notin \mathbf{X}$.)

A probability space is defined by the triple $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the sample space, $\mathcal{F}$ is a σ-algebra of Ω and $\mathbb{P}$ is a probability measure on $(\Omega, \mathcal{F})$. The expected value of the random variable x is denoted by $\mathbb{E}\{x\}$. Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a set $\mathbf{X}$, an $\mathbf{X}$-valued discrete time stochastic process is a sequence of random variables $\mathbf{x} = \{x_k \in \mathbf{X} : k \in \mathbf{T}\}$ on Ω, where $\mathbf{T}$ is a countable time-set. Let $\{\mathcal{F}_k : k \in \mathbf{T}\}$ be a filtration, i.e. a sequence of sub-σ-algebras of $\mathcal{F}$ such that $\mathcal{F}_k \subseteq \mathcal{F}_{k+1}$ for all $k \in \mathbf{T}$. The expectation of a random variable y conditioned on the variables that are $\mathcal{F}_k$-measurable is denoted by $\mathbb{E}_k\{y\}$ (this can be interpreted as the expectation conditioned on the values of the variables at time k).

Definition 1 (*Convergence in Probability*). The sequence of random variables $\mathbf{x} = \{x_k \in \mathbf{X} : k \in \mathbf{T}\}$ is said to converge in probability to a random variable x if $\lim_{k\to\infty} \mathbb{P}\{|x_k - x| \ge \epsilon\} = 0$ for all $\epsilon > 0$, and is denoted by $x_k \xrightarrow{p} x$.

Definition 2 (*Martingales*). A sequence of random variables $\mathbf{x} = \{x_k \in \mathbf{X} : k \in \mathbf{T}\}$, where x_k is $\mathcal{F}_k$-measurable for the filtration $\{\mathcal{F}_k : k \in \mathbf{T}\}$, is a martingale if $\mathbb{E}_k\{x_{k+1}\} = x_k$ with probability 1 for all k. The sequence of random variables $\mathbf{d} = \{d_k : k \in \mathbf{T}\}$, where $d_k = x_{k+1} - x_k$, is called a martingale difference and satisfies $\mathbb{E}_k\{d_k\} = 0$ with probability 1.

2.2. Stochastic approximation

Stochastic Approximation (SA) (Kushner & Yin, 2003) is a root-finding or optimization method for unknown functions $g(\theta)$ from which one only has access to noisy measurements. The idea is to recursively adapt the parameter θ_k using the noisy measurements y_k of $g(\theta_k)$, so that θ_k converges to θ^*, which is the root or the optimal argument of $g(\theta)$. SA can take different forms depending on the type of problem it deals with (optimization or root-finding), the nature of the noise affecting the system (correlated, non-correlated, driven by a Markov process, among others) or the type of convergence that can be obtained (weak convergence or with probability 1). The basics of the variant of SA that is relevant for our work are described next.

A root-finding problem is considered in this work, where the goal is to find a parameter θ^* that satisfies $g(\theta^*) = 0$. Here, y_k are the noisy measurements of $g(\theta_k)$, such that $g(\theta)$ is the limit of the time average of the expected values of y_k, for a fixed $\theta_k := \theta$, when $k \to \infty$. Also, a Markovian type of uncertainty is considered. Associated with the noisy measurements there is a dynamic process with state ξ_k, which is also known as a memory process (Kushner & Yin, 2003), that evolves in a Markovian way (i.e., the probability of ξ_{k+1} is a function of the value of ξ_k, thus Markovian, and θ_k). This memory process defines the noise of the measurements y_k, such that they satisfy $\mathbb{E}_k\{y_k\} = \tilde{g}(\theta_k, \xi_k)$, for some function $\tilde{g}(\cdot,\cdot)$ (Kushner & Vázquez-Abad, 1996). Thus, $\tilde{g}(\cdot,\cdot)$, and θ_k, ξ_k can be understood as the function and the internal states that define the expected value of y_k. This setting will be formalized in Theorem 3 (conditions (i), (iv) and (v)).

For finding θ^* in the bounded set $\Theta = \{\theta : \underline{\theta} \le \theta \le \bar{\theta}\}$, with $\theta^* \in \Theta$, consider the algorithm

$$\theta_{k+1} = \Pi_\Theta \left[\theta_k + \epsilon_k y_k\right] \tag{1}$$

where $\Pi_\Theta[\cdot]$ denotes the projection to Θ and $\{\epsilon_k : k \in \mathbf{T}\}$ is the step-size sequence. The conditions for the convergence of θ_k to θ^* with (1) are studied in Kushner and Vázquez-Abad (1996) and Kushner and Yin (2003), where a weak convergence result