POLITECNICO DI TORINO
Repository ISTITUZIONALE

Stochastic model predictive control of LPV systems via scenario optimization

Original
Stochastic model predictive control of LPV systems via scenario optimization / Calafiore, G.C., Lorenzo, F.. - In: AUTOMATICA. - ISSN 0005-1098. - STAMPA. - 49:6(2013), pp. 1861-1866. [10.1016/j.automatica.2013.02.060]

Availability:
This version is available at: 11583/2507857 since:

Publisher:
Elsevier

Published
DOI:10.1016/j.automatica.2013.02.060

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

07 November 2026

Stochastic Model Predictive Control of LPV Systems via Scenario Optimization [1]

Giuseppe C. Calafiore [a], Lorenzo Fagiano [b,c]

[a] *Dipartimento di Automatica e Informatica, Politecnico di Torino, Torino, Italy*

[b] *Automatic Control Laboratory, Swiss Federal Institute of Technology, Zurich, Switzerland*

[c] *Department of Mechanical Engineering, University of California at Santa Barbara, Santa Barbara, CA, USA*

Abstract

A stochastic receding-horizon control approach for constrained Linear Parameter Varying discrete-time systems is proposed in this paper. It is assumed that the time-varying parameters have stochastic nature and that the system's matrices are bounded but otherwise arbitrary nonlinear functions of these parameters. No specific assumption on the statistics of the parameters is required. By using a randomization approach, a scenario-based finite-horizon optimal control problem is formulated, where only a finite number M of sampled predicted parameter trajectories ('scenarios') are considered. This problem is convex and its solution is a-priori guaranteed to be probabilistically robust, up to a user-defined probability level p. The p level is linked to M by an analytic relationship, which establishes a tradeoff between computational complexity and robustness of the solution. Then, a receding horizon strategy is presented, involving the iterated solution of a scenario-based finite-horizon control problem at each time step. Our key result is to show that the state trajectories of the controlled system reach a terminal positively invariant set in finite time, either deterministically, or with probability no smaller than p. The features of the approach are illustrated by a numerical example.

Key words: Model Predictive Control, Linear Parameter-Varying Systems, Randomized Algorithms, Scenario Optimization, Random Convex Programs.

1 Introduction

In the last decade, several approaches have been proposed for the design of Model Predictive Control (MPC) laws for Linear Parameter Varying (LPV) systems, see, e.g., [7–9,11–13,16]. The existing techniques have the following common features: they are deterministic algorithms, in the sense that for given state value x and parameter value θ they always provide the same optimal control sequence; they guarantee robust stability and satisfaction of constraints; finally, they assume convexity of the sets Σ containing the time-varying system matrices $A(\theta)$, $B(\theta)$ and affine dependence of the matrices on the parameter θ. Some approaches are also able to reduce conservativeness when a bound on the rate of variation of the parameters is available, see, e.g., [7,12].
However, there exist in practice control problems in which Σ is not convex, and $A(\theta)$, $B(\theta)$ depend nonlinearly on θ. In these cases, the existing approaches cannot be applied directly (they may possibly be applied indirectly, by first overbounding Σ with its convex hull, at the cost of potentially introducing conservatism). In order to cope with this issue, we present an approach for the design of MPC laws for LPV systems, in which only boundedness (but not convexity or even connectedness) of the set Σ is assumed. Following an idea common to stochastic MPC techniques (see, e.g., [6,10,14]), we assume that the time-varying parameters θ have known stochastic nature, and exploit this knowledge in the control design. The characterization of θ is quite general, since not only bounds on its rate of variation, but also complex nonlinear models of its time-evolution can be accounted for. Yet, the problem to be solved at each time step is still convex and of manageable size, and constraint satisfaction and convergence of the state to a terminal set are still achieved, with at least a user-defined probability p. The key point for achieving these features is a shift of paradigm,

Email addresses: giuseppe.calafiore@polito.it (Giuseppe C. Calafiore), fagiano@control.ee.ethz.ch (Lorenzo Fagiano).

[1] This paper was not presented at any IFAC meeting. Corresponding author L. Fagiano. This research received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. PIOF-GA-2009-252284 - Marie Curie project "Innovative Control, Identification and Estimation Methodologies for Sustainable Energy Technologies, and from the Italian Ministry of University and Research, under the grant PRIN 20087W5P2K.

from a *deterministic* algorithm to a *randomized* one, i.e., an algorithm that relies on some random choices, see [2,17] for a thorough survey of randomized methods in control. In particular, we here rely on the solution of a scenario Finite Horizon Optimal Control Problem (FHOCP), in which we do not consider all possible outcomes of parameter values, but only a finite number M of randomly chosen instances of them, named the "scenarios." We provide a precise guideline on how to choose M in order to have the guarantee that the probability of success is indeed at least p, and then we describe a receding-horizon implementation of the scenario FHOCP, named MPCS (MPC via Scenario optimization), and prove its constraint satisfaction and convergence properties. The approach we propose here is similar to the one recently presented in [4], where uncertain linear time invariant systems were considered (i.e., the uncertain parameters were assumed to be constant and not measured for feedback control), and where a non-standard cost function based on the distance of the predicted states from a terminal set was used. Here, we extend the results of [4] to the stochastic LPV framework, and we employ a more standard quadratic cost function. To the best of our knowledge, randomized approaches for MPC in the LPV case have never been studied before.

The paper is organized as follows. The problem formulation is described in Section 2; the scenario FHOCP, the MPCS algorithm and their properties are treated in Sections 3, finally Section 4 contains a numerical example.

2 Problem Setting and Assumptions

Consider the following uncertain, discrete time LPV system:

$$x_{t+1} = A(\theta_t)x_t + B(\theta_t)u_t \tag{1}$$

where $t \in \mathbb{Z}$ is the discrete time variable, $x_t \in \mathbb{R}^n$ is the system state, $u_t \in \mathbb{R}^m$ is the control input, $\theta_t \in \Theta_t \subseteq \mathbb{R}^g$ is the vector of uncertain parameters, and $A(\theta)$, $B(\theta)$ are matrices of suitable dimensions. The (generally time varying) sets Θ_t, containing the values of parameter θ_t at time t, are subsets of a time invariant set $\overline{\Theta}$. Let us consider the following assumptions.

Assumption 1 *(Model set) The set* $\Sigma \doteq \{A(\theta),\, B(\theta) : \theta \in \overline{\Theta}\}$ *is bounded.*

Assumption 1 is quite mild as compared with the literature (see e.g. [7–9,11–13,16]), since only boundedness of Σ is required, and not convexity or even connectedness, and there is no restriction on how the parameter θ influences the matrices $A(\theta)$, $B(\theta)$ (e.g., affinely). We will later on ask for another assumption on the system, related to the presence of a (possibly parameter-dependent) affine state feedback law, under which the origin of (1) is stable. The next Assumption characterizes the time-varying parameter θ_t.

Assumption 2 *(Time varying parameters) We assume that the parameter θ_t is measured at each time step t. Also, $\{\theta_t\}_{t=\ldots,-1,0,1,\ldots}$ is assumed to be a strict-sense stationary stochastic process and, for any time instant τ, we denote with $\mathbb{P}_\tau$ the conditional distribution of the forward sequence $\delta_\tau = (\theta_{\tau+1|\tau}, \ldots, \theta_{\tau+N-1|\tau})$, given the past sequence $\mathcal{P}(\tau) \doteq \{\theta_t\}_{t\leq\tau}$, where N is some given integer, and we let Δ_τ be the support set of $\mathbb{P}_\tau$, that is, the set containing the conditional values of δ_τ, given $\mathcal{P}(\tau)$. We assume further that it is possible to obtain sampled values of δ_τ according to $\mathbb{P}_\tau$.*

One common situation arises when $\{\theta_t\}$ is an iid (independent, identically distributed) sequence of random variables, where each variable θ_t has the same distribution $\mathbb{P}_\theta$. In this case, we simply have that $\mathbb{P}_\tau = \mathbb{P}_\theta \times \cdots \times \mathbb{P}_\theta$ (the N-times product measure), and Δ_τ is the N-fold cartesian product of the support set of $\mathbb{P}_\theta$. Another relevant case is when $\{\theta_t\}$ is a Markovian stochastic model produced by a recursion of the type $\theta_{t+1} = \Omega\theta_t + \Upsilon\xi_t$, where $\{\xi_t\}$ is an iid process with marginal $\mathbb{P}_\xi$. Typically, $\mathbb{P}_\xi$ is a standard Normal (Gaussian) with zero mean and unit variance, in which case the conditional distribution $\mathbb{P}_\tau$ of δ_τ is still Normal, with mean given by $[\Omega^\top\ \Omega^{2\top}\ \cdots\ \Omega^{N\top}]^\top\theta_\tau$ and covariance matrix given by the square of a suitable Toeplitz matrix depending on Ω and Υ. As a particular case, such a model can be useful to impose restrictions on the rate of variation of the parameter, by taking $\Omega = I$. Notice that we make no specific assumptions on $\mathbb{P}_t$ and on the support sets Δ_t, which may be unbounded and of any form, as long as Assumption 1 holds. The probability measure $\mathbb{P}_t$ itself can be also not known explicitly, as long as there is some mechanism to obtain samples of δ_t. This can be the case of many application fields, where there exists some complex stochastic model to sample predicted trajectories of the time-varying parameters, but the underlying probability distribution and support set are difficult to compute. We provide an example of this situation in Section 4. We further make the following assumption on the (stochastic) input and state constraints sets.

Assumption 3 *(Convexity of the constraint sets) For any $\theta_t \in \Theta_t$ and any t, $\mathbb{X}(\theta_t) \subseteq \mathbb{R}^n$ and $\mathbb{U}(\theta_t) \subseteq \mathbb{R}^m$ are convex; they contain the origin in their interiors and they are representable by:*

$$\begin{aligned} \mathbb{X}(\theta_t) &= \{x \in \mathbb{R}^n : f_X(x, \theta_t) \preceq 0\} \\ \mathbb{U}(\theta_t) &= \{u \in \mathbb{R}^m : f_U(u, \theta_t) \preceq 0\}, \end{aligned} \tag{2}$$

where $\preceq$ denotes element-wise inequalities, each entry of the functions $f_X : \mathbb{R}^n \times \Theta_t \to \mathbb{R}^r$, $f_U : \mathbb{R}^m \times \Theta_t \to \mathbb{R}^q$ is convex in x and u, respectively, and r, q are suitable integers.

We note that no assumption on the form of the convex sets $\mathbb{X}(\theta_t)$, $\mathbb{U}(\theta_t)$ (e.g. polytopic, ellipsoidal,...), for given $\theta_t \in \Theta$, is made. Indeed, the constraint sets are even allowed to change their form, e.g., from ellipsoids to polytopes, as a function of θ_t. The control problem we consider is to regulate the system state to the origin, subject to satisfaction of the constraints (2) with a probability higher than a user-defined

value p. The latter can be either prescribed by the problem at hand, for example when the input and state constraints are expressed in probability (chance constraints), or it can be interpreted as a relaxation of robust deterministic constraints, such that a small chance of constraint violation is traded off to obtain a tractable problem formulation.

As usual in MPC for LPV systems (see e.g. [7]), our approach relies on the existence of a convex positively invariant terminal set $\mathbb{X}_f$ and an associated affine state-feedback control law $u = K_f(\theta_t)\, x$, possibly depending on the parameter θ_t, that renders the origin of (1) robustly asymptotically stable:

Assumption 4 *(Terminal set and terminal control law) A set $\mathbb{X}_f$, containing the origin in its interior, and a linear state feedback terminal control law $u = K_f(\theta_t)\, x$, $K_f \in \mathbb{R}^{m\times n}$, exist for system (1), such that $\mathbb{X}_f = \{x : f_{X_f}(x) \preceq 0\}$, where $f_{X_f} : \mathbb{R}^n \to \mathbb{R}^l$ has convex components and l is a suitable integer, and $\forall \theta_t \in \Theta_t$, $\forall x_t \in \mathbb{X}_f$, $\forall t$, $A(\theta_t)x_t + B(\theta_t)K_f(\theta_t)\, x_t \in \mathbb{X}_f$; $f_X(x_t, \theta_t) \preceq 0$, $f_U(K_f(\theta_t)\, x_t, \theta_t) \preceq 0$. The origin of the closed loop system with the feedback law $u = K_f(\theta_t)$ is asymptotically stable.*

The terminal set and terminal control law in Assumption 4 are required to always be inside the state constraint set and satisfy input constraints, respectively. Possible methods to compute $\mathbb{X}_f$ and K_f rely on quadratic stability and robust state feedback laws for LPV systems (see, e.g., [7,15]). As an example, by tolerating some conservativeness in the computation of the terminal set, one can embed the set Σ in a polytope, and apply the procedure used in [15] to derive a static state feedback gain K_f and an associated matrix Q_f, such that the level sets of the Lyapunov function $x^\top Q_f x$, $Q_f = Q_f^\top \succ 0$, are robustly positively invariant ellipsoids for the closed loop system, and the origin is asymptotically stable. Finally, since the state and input constraint sets are assumed to contain the origin in their interiors (see Assumption 3) one can find an ellipsoid $\mathbb{X}_f = \{x : x^\top Q_f x \le \rho\}$, for a sufficiently small $\rho > 0$, such that constraints are satisfied for all $x \in \mathbb{X}_f$, all $\theta_t \in \Theta_t$ and all t.

3 MPC for LPV systems via Scenario Optimization

3.1 The scenario finite-horizon optimal control problem

Let $N \in \mathbb{N}$ be a finite control horizon, chosen by the control designer, let $t \ge 0$ be the current time instant and let x_t be the system state observed at time t. We consider the predicted evolution of (1) for N steps forward, under a control law determined at the current time t: $u_{j|t} \doteq K_f(\theta_{j|t})x_{j|t} + v_{j|t}$, $j = 0, \dots, N-1$, where $u_{j|t}$ is the predicted input at time $t+j$ computed at time t, $x_{0|t} = x_t$, $\theta_{0|t} = \theta_t$ and, for $j = 1, \dots, N$, $x_{j|t} = A_{cl}(\theta_{j|t})x_{j-1|t} + B(\theta_{j|t})v_{j-1|t}$, and $A_{cl}(\theta_{j|t}) = A(\theta_{j|t}) + B(\theta_{j|t})K_f(\theta_{j|t})$. Here, $v_{j|t}$, $j = 0, \dots, N-1$, are control corrections at time steps $t+j$, computed at time t, which we collect in the vector $\mathcal{V}_t = [v_{0|t}^\top \ \cdots \ v_{N-1|t}^\top]^\top \in \mathbb{R}^{Nm}$. For given initial state x_t and sequence $\mathcal{V}_t$, let us define the following (stochastic) cost function:

$$J(x_t, \delta_t; \mathcal{V}_t) \doteq \sum_{j=0}^{N-1} x_{j|t}^\top Q x_{j|t} + u_{j|t}^\top R u_{j|t} \tag{3}$$

where $Q = Q^\top \succ 0$, $R = R^\top \succ 0$ are weighting matrices chosen by the control designer. Now, let us consider a finite number M of randomly extracted *scenarios* of δ_t at time t, i.e., $\delta_t^{(1)}, \dots, \delta_t^{(M)}$, which we collect in the *multisample* $\omega_t \doteq (\delta_t^{(1)}, \dots, \delta_t^{(M)})$. The probability distribution of $\omega_t \in \Delta_t^M$ is given by $\mathbb{P}_t^M$. Based on ω_t, for a given state x_t we can formulate the scenario-based FHOCP, as follows (see also [4]):

$$\mathcal{P}(x_t, \omega_t) : \min_{\mathcal{V}_t, z_t, q_t \ge 0} \ z_t + \alpha q_t \quad \text{s.t.:} \tag{4a}$$

$$J(x_t, \delta_t^{(i)}; \mathcal{V}_t) \le z_t; \ \ i = 1, \dots, M \tag{4b}$$

$$f_X(x_{j|t}^{(i)}, \theta_{j|t}^{(i)}) - \mathbf{1} q_t \preceq 0; \ \ {\scriptstyle j = 1, \dots, N-1, \ i = 1, \dots, M} \tag{4c}$$

$$f_U(u_{j|t}^{(i)}, \theta_{j|t}^{(i)}) - \mathbf{1} q_t \preceq 0; \ \ {\scriptstyle j = 1, \dots, N-1, \ i = 1, \dots, M} \tag{4d}$$

$$f_{X_f}(x_{N|t}^{(i)}) - \mathbf{1} q_t \preceq 0; \ \ i = 1, \dots, M.. \tag{4e}$$

In (4), the slack variable $q_t \ge 0$ is used to guarantee feasibility of the optimization problem, by transforming the hard constraints of Assumption 3 into soft ones; the weighting scalar $\alpha > 0$ is chosen by the control designer; finally $\mathbf{1}$ denotes a column vector of appropriate length, containing all ones. We denote with $\mathcal{V}_t^*(x_t, \omega_t) = \{v_{0|t}^*, \dots, v_{N-1|t}^*\}$, $z_t^*(x_t, \omega_t)$ and $q_t^*(x_t, \omega_t)$ an optimal solution to problem (4). We note that, once the multisample ω_t has been extracted, the scenario FHOCP is a convex optimization problem, which can be solved efficiently also with a large number M of samples, even when the system's matrices and the constraints are not convex w.r.t. the time varying parameters θ_t.

Now, denote with $d = mN + 2$ the number of decision variables in problem $\mathcal{P}(x_t, \omega_t)$, let $p \in (0, 1)$ be a given desired probability with which the constraints (4b)-(4e) shall be satisfied, let $\beta \in (0, 1)$ be a given, small probability level (say, $\beta = 10^{-9}$), finally let M be chosen such that

$$\Phi(p, d, M) \le \beta, \tag{5}$$

with $\Phi(p, d, M) \doteq \sum_{j=0}^{d-1} \binom{M}{j} (1-p)^j p^{M-j}$. Then, the solution to problem $\mathcal{P}(x_t, \omega_t)$ enjoys the properties stated in the following proposition (for further comments, see also the discussion in [4] on a similar result for the case of linear time-invariant systems).

Proposition 3.1 *(Finite horizon robustness) With probability larger than $1-\beta$ it holds that the computed control sequence $\mathcal{V}_t^*$:*
a) *steers the state of system (1) to the terminal set $\mathbb{X}_f$ in N steps, with probability at least p and constraint violation q_t^*, i.e.: $\mathbb{P}_t\{\delta_t : f_{X_f}(x_{t+N},\delta_t) - \mathbf{1}q_t^* \preceq 0\} \geq p$;*
b) *Satisfies all state constraints with probability at least p and constraint violation q_t^*, i.e.: $\mathbb{P}_t\{\delta_t : f_X(x_{t+j},\delta_t) - \mathbf{1}q_t^* \preceq 0,\ \forall j \in [1,N]\} \geq p$;*
c) *Satisfies all input constraints with probability at least p and constraint violation q_t^*, i.e.: $\mathbb{P}_t\{\delta_t : f_U(u_{t+j}^*,\delta_t) - \mathbf{1}q_t^* \preceq 0,\ \forall j \in [0,N-1]\} \geq p$.* ■

The inequality (5) provides a precise guideline on how to tune the value of M, for given values of p and β. A suitable value of M can be derived either by inverting numerically equation (5), or by using existing results on explicit bounds for the sample complexity. For example, eq. (5.2) in [3] provides

$$M \geq \frac{2}{(1-p)}\left(\log(\beta^{-1}) + m\,N + 1\right), \qquad (6)$$

(see also [1] for more details and tighter, although slightly more involved, explicit bounds). We remark that the value of M needed to satisfy condition (5) grows at most logarithmically with β^{-1}, as it can be evinced by (6). Hence, the parameter β can be fixed by the designer to a very low level without increasing significantly the complexity of the resulting scenario FHOCP. More insights on this aspect and on the related "certainty equivalence" principle can be found in [4]. We will now present a receding-horizon algorithm that embeds the scenario problem and allows to take advantage of the measure of the state x_t and of the parameter θ_t, at each time step.

3.2 MPC for LPV systems via Scenario optimization (MPCS)

We next use the following notation: "$*$" variables, such as $z_t^*, q_t^*, \mathcal{V}_t^* = \{v_{0|t}^*,\dots,v_{N-1|t}^*\}$, denote the optimal solution of the scenario optimization problem $\mathcal{P}(x_t,\omega_t)$ at time t, given x_t; "$\sim$" variables, $\tilde{z}_t, \tilde{q}_t, \tilde{\mathcal{V}}_t$, denote, respectively, two scalar values and a sequence of N vectors of dimension m, as defined in the algorithm below; finally plain variables, $z_t, q_t, \mathcal{V}_t$, denote the running values of the variables z, q and of the sequence $\mathcal{V} = \{v_{0|t},\dots,v_{N-1|t}\}$ in the algorithm. The first entry in $\mathcal{V}_t$, namely $v_{0|t}$, is the actual control correction that is applied to the system (1) at time t.

3.2.1 MPCS Algorithm

(Initialization) Choose a desired reliability level $p \in (0,1)$ and "certainty equivalence" level $\beta \in (0,1)$ (say, $\beta = 10^{-9}$, or $\beta = 10^{-12}$). Choose an integer M satisfying (5). Choose $\varepsilon \in (0,1]$ (see [4] for the meaning of ε and for tuning guidelines). Given an initial state x_0, extract ω_0 according to $\mathbb{P}_0^M$, solve problem $\mathcal{P}(x_0,\omega_0)$ and obtain the optimal control sequence $\mathcal{V}_0^* = \{v_{0|0}^*, v_{1|0}^* \dots, v_{N-1|0}^*\}$, and the optimal objective z_0^* and constraint violation q_0^*. Set $z_0 = z_0^*, q_0 = q_0^*$, $\mathcal{V}_0 = \mathcal{V}_0^*$, and apply to the system the control action $u_0 = K_f x_0 + v_{0|0}$.
(1) Let $t := t+1$, observe x_t, and set
$\tilde{\mathcal{V}}_t = \{v_{1|t-1},\dots,v_{N-1|t-1},0\}$,
$\tilde{z}_t = \max\left(0, z_{t-1} - \varepsilon x_{t-1}^\top Q x_{t-1}\right)$, $\tilde{q}_t = q_{t-1}$.
(2) Extract the multisample ω_t according to $\mathbb{P}_t^M$, and solve problem $\mathcal{P}(x_t,\omega_t)$. Let $(\mathcal{V}_t^*, z_t^*, q_t^*)$ be the obtained optimal solution.
(3) Evaluate the following collectively exhaustive and mutually exclusive cases:
(3.a) If $z_t^* > \tilde{z}_t$ and $\tilde{z}_t < x_t^\top Q x_t$, then set $\mathcal{V}_t = \tilde{\mathcal{V}}_t$; $z_t = 0$; $q_t = \tilde{q}_t$;
(3.b) If $z_t^* > \tilde{z}_t$ and $\tilde{z}_t \geq x_t^\top Q x_t$, then set $\mathcal{V}_t = \tilde{\mathcal{V}}_t$; $z_t = \tilde{z}_t$; $q_t = \tilde{q}_t$;
(3.c) If $z_t^* \leq \tilde{z}_t$, then set $\mathcal{V}_t = \mathcal{V}_t^*$; $z_t = z_t^*$; $q_t = q_t^*$;
(4) Apply the control input $u_t = K_f\,x_t + v_{0|t}$, and then go to (1). ■

The next theorem provides the main result of this paper, which concerns the guaranteed properties of the closed loop system obtained by applying Algorithm 3.1. We note that, by virtue of Assumption 4, under the terminal control law $u = K_f\,x$ the origin of system (1) is robustly asymptotically stable with region of attraction equal to $\mathbb{X}_f$, and constraints are robustly satisfied for all $x \in \mathbb{X}_f$. Therefore, only the convergence of the state trajectories to $\mathbb{X}_f$ and the satisfaction of constraints for $x \notin \mathbb{X}_f$ are of interest here.

Theorem 3.1 *(Properties of Scenario MPC) Let $v_{0|t}$, $t = 0,1,\dots$ denote the sequence of control actions produced by the MPCS Algorithm, and consider the closed loop system obtained by applying to (1) the control law $u_t = K_f x_t + v_{0|t}$. Let $x_0 \notin \mathbb{X}_f$. Then:*

(a) *With probability larger than $1-\beta$, at all time steps $t = 0,1,\dots$, the probability that the state and input constraints are satisfied with constraint violation q_t is at least p, that is, for $t = 0,1,\dots$, it holds that*

$$\mathbb{P}_t\{\delta_t : f_X(x_{t+1},\delta_t) - \mathbf{1}q_t \preceq 0 \cap f_U(u_t,\delta_t) - \mathbf{1}q_t \preceq 0\} \geq p.$$

(b) *The MPCS Algorithm either: **(i)** makes the state trajectory converge to the terminal set in finite time, i.e., $x_{t+\overline{N}} \in \mathbb{X}_f$, for some $\overline{N} < \infty$, or **(ii)** there exists a finite time t^* such that, with probability larger than $1-\beta$, the forward control sequence $\{v_{0|t^*}, v_{0|t^*+1},\dots v_{0|t^*+N-1}\}$ drives the state of the closed-loop system to the terminal set at time t^*+N-1, with probability at least p and constraint violation q_{t^*}.* ■

Proof 3.1 *See the Appendix.*

4 Numerical example

We consider system (1) with $A(\theta_t) = \begin{bmatrix} \theta_{t,1}\log(\theta_{t,2}) & \theta_{t,3} + e^{\theta_{t,4}}\sin(\theta_{t,5}) \\ 0 & \theta_{t,6} \end{bmatrix}$, $B(\theta) = \begin{bmatrix} \theta_{t,3} + e^{\theta_{t,4}}\cos(\theta_{t,5}) \\ 0.5\theta_{t,6} \end{bmatrix}$, where $\theta_{t,i}$ is the $i-$th component of the parameter vector θ_t. We also consider the following parameter-dependent constraints on the input and state variables:

$$\mathbb{X}(\theta_t = \left\{ x \in \mathbb{R}^2 : \begin{bmatrix} 1 & 0.1\theta_{t,8} \\ 0.1\theta_{t,8} & 1 \\ -1 & -0.1\theta_{t,8} \\ -0.1\theta_{t,8} & -1 \end{bmatrix} x \leq \begin{bmatrix} 2 \\ 0.12 \\ 2 \\ 2 \end{bmatrix} \right\},$$

$\mathbb{U}(\theta_t) = \{u \in \mathbb{R} : |u| \leq 2 + 0.1\theta_{t,7}\}$. The parameter vector θ_t has 300 components, characterized as shown in Table 1.

Table 1
Description of the time-varying parameters in the example (u.d. stands for uniformly distributed, n.d. for normally distributed, $\overline{\lambda}(A)$ is the maximum absolute value of the eigenvalues of A).

$\theta_{t,1}$	u.d. in [0.9,1.1]
$\theta_{t,2}$	u.d. in [2,2.6] if $\theta_{t,1} < 0.95$ n.d. and saturated in the interval [2.8,3] if $\theta_{t,1} \geq 0.95$
$\theta_{t,3}$	u.d. in [0.95,1.1] if $\theta_{t,10} > 1$, u.d. in [1.2,1.3] else
$\theta_{t,4}$	u.d. in [-10,-6] if $t > 2$ and $\theta_{t-3,9} + \theta_{t-2,9} < 1.7$ and $\theta_{t-2,9} + \theta_{t-1,9} > 0.2$; u.d. in [-5,-2] else
$\theta_{t,5}$	u.d. in [0,$\frac{\pi}{2}$] if $\sum_{i=11}^{300} \theta_{t,i} \geq 10$, u.d. in [$\pi$,2$\pi$] else
$\theta_{t,6}$	u.d. in [1,1.2] if $t > 0$ and $\overline{\lambda}(A(\theta_{t-1})) < 0.8$ u.d. in [0.8,1.2] else
$\theta_{t,7}$	u.d. in [-1,1]
$\theta_{t,8}$	u.d. in [-1,1]
$\theta_{t,9}$	u.d. in [0,1]
$\theta_{t,10}$	n.d.
$\theta_{t,i}$	with $i = 11, 300$; distributed according to logistic functions with mean i^{-1} and variance $\frac{pi}{3}i^2$

We employ the following terminal control law and terminal set satisfying Assumption 4: $K_f = [-0.8057 \; -0.8543]$, $\mathbb{X}_f = \{x \in \mathbb{R}^2 : x^T Q_f x \leq 0.01\}$, where $Q_f = \begin{bmatrix} 1.5452 & 0.1865 \\ 0.1865 & 0.9792 \end{bmatrix}$. We designed the MPCS law with $N = 10$, $Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, $R = 1$ and $\alpha = 10^4$, $\beta = 10^{-8}$ and $\varepsilon = 0.01$. By setting desired guaranteed probabilities $p = 0.3$, $p = 0.5$ and $p = 0.95$ for the design, we obtained from (5) values of $M = 42$, 77 and 840, respectively. We carried out $N_{\text{trials}} = 50,000$ Monte Carlo simulations, starting from the state value $x_0 = [-1, \; -2.5]^\top$, which is outside the state constraints and whose corresponding uncorrected input, i.e., $K_f x_0 = 2.94$, is also outside the input constraint set. Indeed this initial condition is not feasible for the deterministic robust counterpart of the scenario problem, hence for some extractions of ω_0 the constraint violation q_0^* is not negligible. In the Monte Carlo simulations, the empirical probability of success $\hat{p}$ has been computed as $\hat{p} = (N_{\text{trials}} - N_{\text{failures}})/N_{\text{trials}}$, where N_{failures} is the number of simulations in which some of the constraints were not satisfied. We computed $\hat{p}$ both for the finite horizon solution, i.e., the sequence $\mathcal{V}_0^*$ obtained at $t = 0$ and applied without re-optimizing at the next time steps, and for the receding horizon solution given by the MPCS algorithm. In particular, we considered as a failure, for a given trial, any constraint violation, i.e., input and/or state constraint violations and the failure in reaching the terminal set either within N steps (for the finite horizon solution) or within $N + 10$ steps for the receding horizon approach, to approximately take into account the convergence results of Theorem 3.1. The outcome of the Monte Carlo simulations is reported in Table 2.

Table 2
Empirical probabilities $\hat{p}$ of constraint satisfaction and convergence to the terminal set, for three different values of p, and $\beta = 10^{-8}$.

	Finite Horizon	MPCS algorithm
$p = 0.3\,(M = 42)$	0.619	0.973
$p = 0.5\,(M = 77)$	0.814	0.981
$p = 0.95\,(M = 840)$	0.989	0.996

The obtained results confirm that the receding horizon implementation, by re-optimizing the control corrections at each time step, yields higher probabilities w.r.t. the finite horizon solution. It can be also noted that some conservativeness exists in the result, in the sense that the empirical probabilities are larger (better) than the theoretical bound, especially for the receding horizon implementation. In this regard, we note that the bound (5) on the violation probability is actually exact for the case of fully supported problems, i.e., random convex programs where the number of support constraints is equal to the number of optimization variables with probability one (see, e.g., [3] and [5]). In other cases, the bound may not be tight in general. In the case of MPC, it is quite unlikely to have a fully supported problem, due to the fact that the control objective is to drive the states well inside the state constraint set, towards the origin, hence the optimal solution is typically supported only by the input and state constraints for the very first time steps. Therefore, the conservativeness comes from the settings of the example itself rather than from poor tightness of the bounds. In other cases, e.g. if the control objective is such that the optimal system's operation is close to the boundary of the state or input constraints, the empirical violation probability might be closer to the bound. As it regards the MPCS algorithm, another important aspect that contributes to outperforming the theoretical bound is the use of state feedback through

the receding horizon strategy, as one can intuitively expect.

We performed the computations in this example using a standard laptop under Yalmip and Matlab®. The average computational times (including also the definition of the constraints in Yalmip) were of about 1.8 s for the case of 42 samples, 2 s for the case of 77 samples, and 10 s for the case of 840 samples. However, we did not carry out any effort to speed-up the computations, since this was not the primary purpose of this work. The important features of the proposed approach, as far as computational aspects are concerned, are the convexity of the scenario problem and the presence of a well precise structure of the constraints. Considering these features, and given the recent developments in the field of efficient on-line MPC (see, e.g., [18] and the references therein) and the continuous improvements in terms of computational power/cost of hardware, the computational times of the method appear to pose little issues in perspective.

Appendix

Proof of Theorem 3.1. The proof of statements (a) and (b)-(ii) follows a line similar to that of Theorem 4.1 in [4], whereas a different reasoning should be made for statement (b)-(i). We next report a full proof of Theorem 3.1.

Proof of statement (a). At time $t = 0$, Proposition 3.1 guarantees with probability larger than $1 - \beta$ that the first control correction satisfies the constraints on u_0 and x_1 with probability no less than p and constraint violation $q_0 = q_0^*$. At any generic time step $t \geq 1$, the variables $(\tilde{\mathcal{V}}_t, \tilde{z}_t, \tilde{q}_t)$ are computed. Then, two cases may occur. If $z_t^* \leq \tilde{z}_t$, then case **(3.c)** is detected, and the first element $v_{0|t}^*$ of the optimal sequence $\mathcal{V}_t^*$ is applied to the system. Being this sequence the solution of a scenario optimization problem, with probability larger than $1 - \beta$ the probability of satisfying state and input constraints is no less than p, with constraint violation $q_t = q_t^*$. If, on the other hand, $z_t^* > \tilde{z}_t$, then we are either in case **(3.a)** or **(3.b)**, and in both cases the element $v_{k|t-k}^*$, for some $k \in [1, N-1]$, is applied to the system. Being this value part of the solution sequence $\mathcal{V}_{t-k}^*$, with corresponding constraint violation q_{t-k}^*, again the probability of satisfying state and input constraints is no less than p, with constraint violation $q_t = \tilde{q}_t = q_{t-k}^*$.

Proof of statement (b). Each run of the MPCS Algorithm may have one of two possible behaviors, depending on whether or not there exists a finite time $t > 0$ such that $z_t^* > \tilde{z}_t$ and $\tilde{z}_t < x_t^\top Q x_t$, that is, whether or not the situation in step **(3.a)** is ever satisfied. We then name $\mathcal{A}$ the situation when condition in step **(3.a)** is met at some finite $t > 0$, and $\bar{\mathcal{A}}$ the complementary situation when this condition is not satisfied at any finite time, that is when $z_t^* \leq \tilde{z}_t$ or $\tilde{z}_t \geq x_t^\top Q x_t$ holds for all $t > 0$.
(b)-(i). Consider the situation $\bar{\mathcal{A}}$, and let us define the ellipsoid $\mathbb{X}_{Q_f} \doteq \{x \in \mathbb{R}^n : x^\top Q_f x \leq \rho\}$ for some $\rho > 0$, such that $Q_f = Q_f^\top \succ 0$ and $\mathbb{X}_{Q_f} \subset \mathbb{X}_f$. Such a set can always be found, since the origin is assumed to be contained in the interior of the terminal set $\mathbb{X}_f$. At a generic time t, at step **(3)** of the MPCS algorithm, if $z_t^* > \tilde{z}_t$, then, since it is assumed that we are in situation $\bar{\mathcal{A}}$, it must be $\tilde{z}_t \geq x_t^\top Q x_t$, thus case **(3.b)** occurs, and the values $\mathcal{V}_t = \tilde{\mathcal{V}}_t$ and $z_t = \tilde{z}_t$ are set. Now, recalling that $\tilde{z}_t = \max(0, z_{t-1} - \varepsilon x_{t-1}^\top Q x_{t-1})$, two cases may occur: either $\tilde{z}_t = 0$ or $\tilde{z}_t = z_{t-1} - \varepsilon x_{t-1}^\top Q x_{t-1} > 0$. If $\tilde{z}_t = 0$, we have $0 = \tilde{z}_t \geq x_t^\top Q x_t$, i.e., $x_t^\top Q x_t = 0$, which would imply that $x_t = 0$ and that the terminal set has been reached, since the origin is contained in its interior. Otherwise, if $\tilde{z}_t = z_{t-1} - \varepsilon x_{t-1}^\top Q x_{t-1} > 0$, then we have:

$$z_t \geq x_t^\top Q x_t \geq \beta x_t^\top Q_f x_t, \tag{7}$$

where $\beta = \frac{\underline{\lambda}(Q)}{\overline{\lambda}(Q_f)}$ and $\underline{\lambda}(\cdot)$, $\overline{\lambda}(\cdot)$ are the minimum and maximum eigenvalues, respectively. Moreover, we have

$$z_t = z_{t-1} - \varepsilon x_{t-1}^\top Q x_{t-1} \leq z_{t-1} - \varepsilon\beta x_{t-1}^\top Q_f x_{t-1},$$

hence

$$z_t < z_{t-1} - \varepsilon\beta\rho. \tag{8}$$

Note that the inequality in (8) holds true as long as $x_{t-1}^\top Q x_{t-1} > \rho$, i.e., as long as the state x_{t-1} is outside the ellipsoid $\mathbb{X}_{Q_f}$. Indeed this condition is satisfied if $x_{t-1} \notin \mathbb{X}_f$, since $\mathbb{X}_{Q_f} \subset \mathbb{X}_f$. On the other hand, if at step **(3)** of the MPCS algorithm it happens that $z_t^* \leq \tilde{z}_t$, then case **(3.c)** occurs, and the optimal values $\mathcal{V}_t^*$ and z_t^* are retained, i.e., $z_t = z_t^*$, $\mathcal{V}_t = \mathcal{V}_t^*$. In this case, it can be noted that the inequalities (7)-(8) still hold true. Then, as long as the case $z_t^* \leq (z_{t-1} - \varepsilon d(x_{t-1}, \mathbb{X}_f))$ or $\tilde{z}_t \geq d(x_t, \mathbb{X}_f)$ holds true as assumed, (8) can be applied recursively starting from $z_0 = z_0^*$ (i.e., the worst-case cost computed at time $t = 0$), and, by considering (7), we have:

$$x_{t+\overline{N}}^\top Q_f x_{t+\overline{N}} \leq \frac{z_{t+\overline{N}}}{\beta} \leq \frac{z_0}{\beta} - \overline{N}\varepsilon\rho. \tag{9}$$

Therefore, for a finite $\overline{N} > \frac{z_0}{\beta\varepsilon\rho} - \frac{1}{\varepsilon}$, we have that $x_{t+\overline{N}}^\top Q_f x_{t+L} \leq \rho$, hence $x_{t+\overline{N}} \in \mathbb{X}_{Q_f} \subset \mathbb{X}_f$, i.e., the state has reached the terminal set. This proves the statement **(b)-(i)**.
(b)-(ii). Let us next analyze what happens in case $\mathcal{A}$. Let $\bar{t} > 0$ be the time instant at which the case $z_t^* > \tilde{z}_t$ and $\tilde{z}_t < x_t^\top Q x_t$ is met for the first time, and let $t^* < \bar{t}$ be the last time at which case $z_t^* \leq \tilde{z}_t$ was satisfied, that is the last time previous to $\bar{t}$ when an optimal command sequence was retained, together with its constraint violation q_t^*, according to case **(3.c)** of the MPCS algorithm; let $\ell = \bar{t} - t^* \geq 1$. According to step **(3.a)** of the MPCS algorithm, we set

$$\mathcal{V}_{\bar{t}} = \tilde{\mathcal{V}}_{\bar{t}},\ z_{\bar{t}} = 0,\ q_{\bar{t}} = \tilde{q}_{\bar{t}}. \tag{10}$$

Thus, at step **(4)** of the algorithm, the control move $u_{\bar{t}} = K_f x_{\bar{t}} + v_{0|\bar{t}}$ is applied to the system at time $\bar{t}$, where $v_{0|\bar{t}} = v_{\ell|t^*}^*$, i.e., $v_{0|\bar{t}}$ is the optimal correction predicted for time $t^* + \ell = \bar{t}$, computed at time t^*. At time step $t = \bar{t}+1$, the state variable $x_{\bar{t}+1}$ is observed and $(\tilde{V}_{\bar{t}+1}, \tilde{z}_{\bar{t}+1}, \tilde{q}_{\bar{t}+1})$ are computed as $\tilde{z}_{\bar{t}+1} = \max(0, z_{\bar{t}} - \varepsilon x_{\bar{t}}^\top Q x_t)$, $\tilde{q}_{\bar{t}+1} = q_{\bar{t}}$, $\tilde{\mathcal{V}}_{\bar{t}+1} = \{v_{1|\bar{t}}, \ldots, v_{N-1|\bar{t}}, 0\} = \{v_{\ell+1|t^*}^*, v_{\ell+2|t^*}^*, \ldots, v_{N-1|t^*}^*, 0, \ldots, 0\}$. Since (10) holds, it must be $\tilde{z}_{\bar{t}+1} = 0$. Then, $z_{\bar{t}+1}^*$, $q_{\bar{t}+1}^*$ and $\mathcal{V}_{\bar{t}+1}^*$ are computed at step **(2)**, and we notice that, by definition, $z_{\bar{t}+1}^* \geq 0$. Therefore, at step **(3)** of the algorithm either *(i)* case **(3.a)** $\{z_{\bar{t}+1}^* > (z_{\bar{t}} - \varepsilon x_t^\top Q x_t)$ and $\tilde{z}_{\bar{t}+1} < x_t^\top Q x_t\}$ is detected again, or *(ii)* one of cases **(3.b)** or **(3.c)** are detected, which would imply, respectively, $0 = \tilde{z}_{\bar{t}+1} \geq x_t^\top Q x_t$, or $0 \leq x_t^\top Q x_t \leq z_{\bar{t}+1}^* = \tilde{z}_{\bar{t}+1} = 0$. In either sub-case of *(ii)* we would have $x_t^\top Q x_t = 0$, so that convergence to the terminal set would be achieved. Consider then case *(i)*: the values

$\mathcal{V}_{\bar{t}+1} = \tilde{\mathcal{V}}_{\bar{t}+1}$, $z_{\bar{t}+1} = 0$ and $q_{\bar{t}+1} = q_{\bar{t}}$ are set in the algorithm, and the control move $u_{\bar{t}+1} = K_f x_{\bar{t}+1} + v^*_{\ell+1|t^*}$ is applied to the system. Now, the same circumstances actually reproduce for all time steps $t = \bar{t} + k$, $k \geq 0$, so that the optimal input sequence $\mathcal{V}^*_{t^*}$, computed at time t^* by solving a scenario FHOCP, is the one actually next applied to the system, and the related constraint violation q^*_t is retained for all $t \geq t^*$. Thus, in case $\mathcal{A}$, there exists a finite time t^* such that the sequence $\mathcal{V}^*_{t^*}$ is applied to the system for all subsequent instants $t = t^* + k$, $k = 0, \ldots, N-1$. Now, the sequence $\mathcal{V}^*_{t^*}$ is the result of the solution of the scenario FHOCP $\mathcal{P}(x_{t^*}, \omega_{t^*})$: we can hence claim, in virtue of Proposition 3.1, that with probability larger than $1 - \beta$ this sequence will satisfy the problem constraints and reach the terminal set within the time window from t^* to $t^* + N$, with probability at least p and constraint violation q^*_t. This proves the statement **(b)-(ii)**.

References

[1] T. Alamo, R. Tempo, and A. Luque. *On the Sample Complexity of Probabilistic Analysis and Design Methods*, pages 39–50. Perspectives in Mathematical System Theory, Control and Signal Processing. Springer-Verlag, London, 2010.

[2] G. Calafiore, F. Dabbene, and R. Tempo. Research on probabilistic methods for control system design. *Automatica*, 47:1279–1293, 2011.

[3] G.C. Calafiore. Random convex programs. *SIAM Journal on Optimization*, 20:3427–3464, 2010.

[4] G.C. Calafiore and L. Fagiano. Robust model predictive control via scenario optimization. *IEEE Transactions on Automatic Control*, 58(1):219–224, 2013.

[5] M.C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

[6] M. Cannon, B. Kouvaritakis, S.V. Raković, and Q. Cheng. Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Transactions on Automatic Control*, 56(1):194–200, 2011.

[7] A. Casavola, D. Famularo, and G. Franzè. A feedback min-max MPC algorithm for LPV systems subject to bounded rates of change of parameters. *IEEE Transactions on Automatic Control*, 47(7):1147–1153, 2002.

[8] A. Casavola, D. Famularo, and G. Franzè. Predictive control of constrained nonlinear systems via LPV linear embeddings. *International Journal of Robust and Nonlinear Control*, 13:281–294, 2003.

[9] A. Casavola, M. Giannelli, and E. Mosca. Min-max predictive control strategies for input-saturated polytopic uncertain systems. *Automatica*, 36:125–133, 2000.

[10] Eugenio Cinquemani, Mayank Agarwal, Debasish Chatterjee, and John Lygeros. Convexity and convex approximations of discrete-time stochastic control problems with constraints. *Automatica*, 47(9):2082 – 2087, 2011.

[11] M. Jungers, R. C.L.F. Oliveira, and P. L.D. Peres. MPC for LPV systems with bounded parameter variations. *International Journal of Control*, 84(1):24–36, 2011.

[12] D. Li and Y. Xi. The feedback robust MPC for LPV systems with bounded rates of parameter changes. *IEEE Transactions on Automatic Control*, 55(2):503–507, 2010.

[13] Y. Lu and Y. Arkun. Quasi-min-max MPC algorithms for LPV systems. *Automatica*, 36:527–540, 2000.

[14] Maaz Mahmood and Prashant Mhaskar. Lyapunov-based model predictive control of stochastic nonlinear systems. *Automatica*, 48(9):2271 – 2276, 2012.

[15] W. Mao. Robust stabilization of uncertain time-varying discrete systems and comments on “an improved approach for constrained robust model predictive control”. *Automatica*, 39:1109–1112, 2003.

[16] B. Pluymers, J.A. Rossiter, J.A.K. Suykens, and B. De Moor. Interpolation based MPC for LPV systems using polyhedral invariant sets. In *American Control Conference*, pages 810–815, Portland, OR, 2005.

[17] R. Tempo, G. Calafiore, and F. Dabbene. *Randomized Algorithms for Analysis and Control of Uncertain Systems - With Applications, II Ed.* Springer, 2013.

[18] P. Zometa, M. Kögel, T. Faulwasser, and R. Findeisen. Implementation aspects of model predictive control for embedded systems. In *American Control Conference*, Montreal, Canada, 2012.