

A Randomized Approach to Stochastic Model Predictive Control*

Maria Prandini¹, Simone Garatti¹, and John Lygeros²

Abstract—In this paper, we propose a novel randomized approach to Stochastic Model Predictive Control (SMPC) for a linear system affected by a disturbance with unbounded support. As it is common in this setup, we focus on the case where the input/state of the system are subject to probabilistic constraints, i.e., the constraints have to be satisfied for all the disturbance realizations but for a set having probability smaller than a given threshold. This leads to solving at each time t a finite-horizon chance-constrained optimization problem, which is known to be computationally intractable except for few special cases. The key distinguishing feature of our approach is that the solution to this finite-horizon chance-constrained problem is computed by first extracting at random a finite number of disturbance realizations, and then replacing the probabilistic constraints with hard constraints associated with the extracted disturbance realizations only. Despite the apparent naivety of the approach, we show that, if the control policy is suitably parameterized and the number of disturbance realizations is appropriately chosen, then, the obtained solution is guaranteed to satisfy the original probabilistic constraints. Interestingly, the approach does not require any restrictive assumption on the disturbance distribution and has a wide realm of applicability.

I. INTRODUCTION

Model Predictive Control (MPC) is a control design methodology that has been introduced to face infinite-horizon constrained optimal control problems. The key idea of MPC is to find an approximate solution to the original infinite-horizon problem by solving at each sampling time a finite-horizon constrained optimal control problem, and then implementing the control law in accordance to a receding horizon strategy.

The presence of disturbances is quite common in practice. Recently, two different approaches have been introduced to address this issue, namely, robust MPC and stochastic MPC. In robust MPC (see [?], [?], [?], [?], [?], [?] and the references therein), a min-max approach is taken where the control cost is optimized against the worst disturbance realization, while guaranteeing constraint satisfaction. Although successful in many cases, the min-max strategy may lead to conservative results, since the disturbance distribution is not accounted for and all disturbance realizations are treated as equally likely. Indeed, it might be the case that low probability disturbance realizations cause a significant

deterioration in the cost or even the infeasibility of some constraint.

To overcome these limitations of robust MPC, an average cost and probabilistic constraints are typically considered in stochastic MPC (see, e.g., [?], [?], [?], [?], [?], [?]). In this setup, a violation of the constraint is accepted, although this must happen for few disturbance realizations only, having altogether probability no greater than a chosen threshold. This rules out “bad” situations adversely affecting the robust approach. Moreover, probabilistic constraints are the only way to avoid infeasibility of state constraints when the disturbance has unbounded support.

Unfortunately, probabilistic constraints are in general non-convex and more difficult to treat than usual non-probabilistic constraints. The resulting finite-horizon optimization problem with probabilistic constraints belongs, indeed, to the class of the chance-constrained optimization problems, [?], [?], which are known to be hard to solve in general.

In the literature on stochastic MPC, solutions to the chance-constrained problem have been proposed for stochastic linear systems, under the assumption that the disturbance either is a sequence of bounded and i.i.d. (independent and identically distributed) random variables, [?], [?], [?], or has a certain specific distribution, e.g. Gaussian [?]. Our work in this paper aims to provide an alternative scheme for the solution of the finite-horizon optimization problem with probabilistic constraints, which rests on recent developments on the randomized solution to chance-constrained optimization problems, [?]. The advantage with respect to the existing schemes is that no assumption on the disturbance (boundedness, independence, Gaussianity) is required whatsoever.

The paper is organized as follows. We first describe the control problem in Section ??, and then illustrate the randomized resolution scheme in Section ?. The performance of the method is demonstrated in Section ?? on a numerical example. Concluding remarks are provided in Section ?.

II. PROBLEM FORMULATION

Consider a linear system whose state $x_t \in \mathbb{R}^n$ evolves according to the equation

$$x_{t+1} = Ax_t + Bu_t + w_t,$$

where $u_t \in \mathbb{R}^m$ is the control input and $w_t \in \mathbb{R}^n$ is a stochastic disturbance with a possibly unbounded support.

We assume that the entire state vector of the system is known at each time instant and address the problem of designing a state-feedback control policy so as to optimize the system performance over the look-ahead time horizon

*Research supported by the European Commission under the MoVeS project, FP7-ICT-2009-257005.

¹M. Prandini and S. Garatti are with the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20113 Milano, Italy {prandini, sgaratti}@elet.polimi.it

²J. Lygeros is with ETH Zurich, Automatic Control Laboratory, ETL I 22, Physikstrasse 3, CH-8092 Zurich, Switzerland lygeros@control.ee.ethz.ch

$[\tau, \tau + M]$, while satisfying some constraint on the input/state along the same horizon. Here, τ represents the current time instant and $M > 0$ the prediction horizon.

More precisely, let $\mu_{\tau+i} : \mathbb{R}^{n \times i} \rightarrow \mathbb{R}^m$, $i \in \{0, 1, \dots, M-1\}$ denote a measurable function mapping the state from time τ to time $\tau+i$ into the control input to be applied at time $\tau+i$, i.e.,

$$u_{\tau+i} = \mu_{\tau+i}(x_\tau, x_{\tau+1}, \dots, x_{\tau+i}).$$

Our goal is to design a state-feedback control policy, i.e., a sequence of functions $\{\mu_\tau, \mu_{\tau+1}, \dots, \mu_{\tau+M-1}\}$, so as to minimize the average quadratic cost

$$\mathbb{E} \left[\sum_{i=1}^M x_{\tau+i}^T Q_i x_{\tau+i} + \sum_{i=0}^{M-1} u_{\tau+i}^T R_i u_{\tau+i} \right], \quad (1)$$

with Q_i and R_i symmetric positive definite matrices of appropriate dimensions, subject to the probabilistic constraint:

$$\mathbb{P}\{f(x_{\tau+1}, \dots, x_{\tau+M}, u_\tau, \dots, u_{\tau+M-1}) \leq 0\} \geq 1 - \varepsilon, \quad (2)$$

where $f : \mathbb{R}^{n \times M + m \times (M-1)} \rightarrow \mathbb{R}^q$ is a q -dimensional measurable function and the inner inequality appearing in equation (??) is interpreted componentwise. In the above expressions, \mathbb{P} denotes the probability distribution of the disturbance process $w_\tau, w_{\tau+1}, \dots, w_{\tau+M-1}$ affecting the system evolution along the time horizon $[\tau, \tau + M]$. $\mathbb{E}[\cdot]$ is the expectation associated to \mathbb{P} .

Remark 1 (probabilistic constraints): In the probabilistic constraint (??), condition $f(x_{\tau+1}, \dots, x_{\tau+M}, u_\tau, \dots, u_{\tau+M-1}) \leq 0$ is not required to hold for all possible disturbance realizations. Parameter $\varepsilon \in (0, 1)$ entering (??) quantifies the extent to which the constraint on the state and input defined through f can be violated. This sort of compromise is usually unavoidable in a stochastic setting. Indeed, if the disturbance has unbounded support, a hard constraint (i.e. a constraint that must hold for all possible disturbance realizations with no exceptions) on the input would generally lead to a too conservative solution, whereas a hard constraint on the state would generally be unfeasible. \square

Remark 2 (input/state constraint function): In typical cases, function f defining the constraint in (??) is used to force a saturation on the input or to pose a safety constraint on the state. If e.g. f takes the form:

$$f(x_{\tau+1}, \dots, x_{\tau+M}, u_\tau, \dots, u_{\tau+M-1}) = \begin{bmatrix} \sup_{i=0, \dots, M-1} \|u_{\tau+i}\|_\infty - \bar{u} \\ \sup_{i=1, \dots, M} \|x_{\tau+i}\|_\infty - \bar{x} \end{bmatrix}$$

then, (??) represents an input saturation and a safety constraint that should hold jointly on a set of disturbance realizations whose probability is at least $1 - \varepsilon$. \square

Note that in the standard LQG setting, the problem of minimizing the cost (??) without the constraint (??) can be solved analytically, and the optimal solution is characterized by μ_{τ_i} functions that are affine in the state. In presence of constraints, or for an arbitrary distribution of the disturbance, the problem of finding the optimal state-feedback policy becomes quite challenging. One can then look for a suboptimal

solution by parameterizing the $\mu_{\tau+i}$ functions so as to obtain an optimization problem with a finite number of optimization variables. Inspired by the structure of the LQG solution, one can choose $\mu_{\tau+i}$ to be affine in the state, [?], [?], [?]. In that case, the resulting control input turns also out to be affine in the disturbance samples, since the disturbance can be perfectly reconstructed from the state measurements according to

$$w_{\tau+i} = x_{\tau+i+1} - Ax_{\tau+i} - Bu_{\tau+i}.$$

As a result, the control input can be directly parameterized as an affine function of the disturbance

$$u_{\tau+i} = \gamma_i + \sum_{j=0}^{i-1} \theta_{i,j} w_{\tau+j}, \quad (3)$$

with $\gamma_i \in \mathbb{R}^m$ and $\theta_{i,j} \in \mathbb{R}^{m \times n}$, which makes the closed-loop control system to be designed equivalent to an open-loop control system with a feedforward disturbance compensator, [?].

To illustrate the advantages obtained by adopting this parametrization, we first need to introduce some compact notations for the system evolution along the reference time horizon $[\tau, \tau + M]$.

If we define the following vectors of state, input and disturbance signals:

$$\mathbf{x}_+ = \begin{bmatrix} x_{\tau+1} \\ x_{\tau+2} \\ \vdots \\ x_{\tau+M} \end{bmatrix} \quad \mathbf{u} = \begin{bmatrix} u_\tau \\ u_{\tau+1} \\ \vdots \\ u_{\tau+M-1} \end{bmatrix} \quad \mathbf{w} = \begin{bmatrix} w_\tau \\ w_{\tau+1} \\ \vdots \\ w_{\tau+M-1} \end{bmatrix}$$

then, it is easy to show that

$$\begin{aligned} \mathbf{x}_+ &= \mathbf{F}x_\tau + \mathbf{G}\mathbf{u} + \mathbf{H}\mathbf{w} \\ \mathbf{u} &= \mathbf{\Gamma} + \mathbf{\Theta}\mathbf{w}, \end{aligned}$$

where matrices \mathbf{F} , \mathbf{G} and \mathbf{H} are given by

$$\mathbf{F} = \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^M \end{bmatrix} \quad \mathbf{G} = \begin{bmatrix} B & 0_{n \times m} & \cdots & 0_{n \times m} \\ AB & B & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0_{n \times m} \\ A^{M-1}B & \cdots & AB & B \end{bmatrix}$$

$$\mathbf{H} = \begin{bmatrix} I_{n \times n} & 0_{n \times n} & \cdots & 0_{n \times n} \\ A & I_{n \times n} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0_{n \times n} \\ A^{M-1} & \cdots & A & I_{n \times n} \end{bmatrix},$$

whereas $\mathbf{\Gamma}$ and $\mathbf{\Theta}$ contains the control law parameters and are given by

$$\mathbf{\Gamma} = \begin{bmatrix} \gamma_0 \\ \gamma_1 \\ \vdots \\ \gamma_{M-1} \end{bmatrix} \quad \mathbf{\Theta} = \begin{bmatrix} 0_{m \times n} & 0_{m \times n} & \cdots & 0_{m \times n} \\ \theta_{1,0} & 0_{m \times n} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0_{m \times n} \\ \theta_{M-1,0} & \cdots & \theta_{M-1,M-2} & 0_{m \times n} \end{bmatrix}.$$

If we set

$$\mathbf{Q} = \begin{bmatrix} Q_1 & \cdots & 0_{n \times n} \\ \vdots & \ddots & \vdots \\ 0_{n \times n} & \cdots & Q_M \end{bmatrix} \quad \mathbf{R} = \begin{bmatrix} R_0 & \cdots & 0_{m \times m} \\ \vdots & \ddots & \vdots \\ 0_{m \times m} & \cdots & R_{M-1} \end{bmatrix}$$

then, the control cost (??) can be expressed as follows

$$\begin{aligned} J(\Gamma, \Theta) &= \mathbb{E} [\mathbf{x}_+^T \mathbf{Q} \mathbf{x}_+ + \mathbf{u}^T \mathbf{R} \mathbf{u}] \\ &= (\mathbf{F}x_\tau + \mathbf{G}\Gamma)^T \mathbf{Q} (\mathbf{F}x_\tau + \mathbf{G}\Gamma) \\ &\quad + 2(\mathbf{F}x_\tau + \mathbf{G}\Gamma)^T \mathbf{Q} (\mathbf{H} + \mathbf{G}\Theta) \cdot \mathbb{E} [\mathbf{w}] \\ &\quad + \text{tr} [(\mathbf{H} + \mathbf{G}\Theta)^T \mathbf{Q} (\mathbf{H} + \mathbf{G}\Theta) \cdot \mathbb{E} [\mathbf{w}\mathbf{w}^T]] \\ &\quad + \Gamma^T \mathbf{R} \Gamma + 2\Gamma^T \mathbf{R} \Theta \mathbb{E} [\mathbf{w}] + \text{tr} [\Theta^T \mathbf{R} \Theta \cdot \mathbb{E} [\mathbf{w}\mathbf{w}^T]], \end{aligned}$$

which shows that is a convex function of the control policy parametrization (Γ, Θ) .

The optimization problem to be solved at time τ is then given by the following chance-constraint optimization program:

$$\begin{aligned} \min_{\Gamma, \Theta} J(\Gamma, \Theta) \quad \text{subject to:} \quad (4) \\ \mathbb{P}\{f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}, \Gamma + \Theta\mathbf{w}) \leq 0\} \geq 1 - \varepsilon \end{aligned}$$

Note that the probabilistic constraint is generally not convex, even for a convex function $f(\cdot)$, which makes the chance-constraint problem (??) hard to solve. Approximate solutions resting on a convex over-approximation of the probabilistic constraint can be worked out for certain disturbance distributions, like the Gaussian one (see, e.g., [?] and the references therein). In the next section, we describe an alternative method that is of quite general applicability since it does not require any specific assumption on the disturbance distribution.

III. THE SCENARIO-BASED SOLUTION

The scenario approach, [?], [?], [?], [?], [?], is a recent paradigm for computing approximate solutions to general chance-constrained problems

$$\begin{aligned} \min_{\alpha \in \mathbb{R}^{n_\alpha}} \ell(\alpha) \quad (5) \\ \text{subject to: } \mathbb{P}_\omega \{\varphi(\alpha, \omega) \leq 0\} \geq 1 - \varepsilon, \end{aligned}$$

at relatively low computational effort.

In problem (??), α is an n_α -dimensional optimization variable, whereas ω is the stochastic uncertainty parameter with probability distribution \mathbb{P}_ω .

The only assumption the scenario approach relies on is the convexity of the cost $\ell(\alpha)$ and the convexity of $\varphi(\alpha, \omega)$ with respect to the optimization variable α only (the dependence on the stochastic parameter ω can be arbitrary). In our context $\alpha = (\Gamma, \Theta)$ and $\ell(\alpha) = J(\Gamma, \Theta)$ which is convex (indeed, it is quadratic) in α given the affine parametrization of the control policy in (??). As for the second assumption, we have that

$$\varphi(\alpha, \omega) = f(\mathbf{x}_+, \mathbf{u}) = f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}, \Gamma + \Theta\mathbf{w}),$$

where the stochastic parameter ω is the noise vector \mathbf{w} . Given again the affine parametrization of \mathbf{x}_+ and \mathbf{u} , it is enough that

$f(\mathbf{x}_+, \mathbf{u})$ is convex in \mathbf{x}_+, \mathbf{u} . This latter condition is verified in standard problems, e.g. for the function f given in Remark ??.

On the algorithmic side, the scenario approach builds on a very intuitive idea: a number, say N , of realizations of the noise vector \mathbf{w} , say $\mathbf{w}^{(i)}$ for $i = 1, 2, \dots, N$, are extracted according to the underlying probability measure \mathbb{P} and optimization is performed by taking into account this finite number of instances of \mathbf{w} only. To be precise, letting ε' be a user chosen parameter such that $0 \leq \varepsilon' < \varepsilon$, the Scenario Algorithm aims at finding a solution that violates the condition

$$f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}^{(i)}, \Gamma + \Theta\mathbf{w}^{(i)}) \leq 0$$

$\lfloor \varepsilon' N \rfloor$ times out of N , i.e. with an *empirical* probability equal to ε' ($\lfloor \cdot \rfloor$ denotes the integer part). See Algorithm ??.

Algorithm 1 The Scenario Algorithm (SA)

1: Solve problem

$$\begin{aligned} \min_{\Gamma, \Theta} J(\Gamma, \Theta) \quad \text{subject to:} \\ f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}^{(i)}, \Gamma + \Theta\mathbf{w}^{(i)}) \leq 0, \\ i \in \{1, 2, \dots, N\}. \end{aligned}$$

Store the solution.

2: Let i run over $1, 2, \dots, N$ and find the constraints violated by the stored solution, i.e. find the indexes i such that

$$f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}^{(i)}, \Gamma + \Theta\mathbf{w}^{(i)}) > 0.$$

Let these indexes be j_1, j_2, \dots, j_L . If L is greater than or equal to $\lfloor \varepsilon' N \rfloor$, then **halt** the algorithm and return the stored solution.

3: Find the active constraints for the stored solution, i.e. the indexes i such that

$$f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}^{(i)}, \Gamma + \Theta\mathbf{w}^{(i)}) = 0.$$

Let these indexes be i_1, i_2, \dots, i_q .

4: **For** $k = 1, 2, \dots, q$

Solve problem

$$\begin{aligned} \min_{\Gamma, \Theta} J(\Gamma, \Theta) \quad \text{subject to:} \\ f(\mathbf{F}x_\tau + \mathbf{G}\Gamma + (\mathbf{H} + \mathbf{G}\Theta)\mathbf{w}^{(i)}, \Gamma + \Theta\mathbf{w}^{(i)}) \leq 0, \\ i \in \{1, 2, \dots, N\} / \{i_k, j_1, j_2, \dots, j_L\}. \end{aligned}$$

If the obtained cost is better than the cost of the stored solution, then store the last computed solution.

End For

5: **Goto** 2

Algorithm ?? comes to termination whenever there is at each step at least one active constraint such that, if eliminated, then the cost J improves. This condition is assumed here for granted and is usually satisfied unless for degenerate situations.

which corresponds to:

$$\Theta = \begin{bmatrix} \theta_1 & & & & & \\ \vdots & \ddots & & & & \\ \theta_k & \ddots & \ddots & & & \\ & \ddots & \ddots & \ddots & & \\ & & & \theta_k & \cdots & \theta_1 \end{bmatrix}.$$

In this case $r = mM + mnk$

4.

$$u_{\tau+i} = \gamma_i,$$

i.e. $\Theta = 0$. In this case $r = mM$. This last control policy is that with the minimum number of optimization variables, but corresponds to an open loop scheme. It can be combined with a fixed linear state-feedback controller,

$$u_{\tau+i} = \gamma_i + \bar{K}x_{\tau+i}, \quad \bar{K} \text{ fixed},$$

to improve performance, [?], [?].

IV. A NUMERICAL EXAMPLE

In this section, we illustrate the proposed approach to the approximate solution of the chance-constrained optimization problem (??) on a numerical example. The example is taken from [?], and the performance of our approach is compared with that of the methods for the over-approximation of the probabilistic constraint presented in [?].




Fig. 1. Mechanical system. $\bar{l}_1, \bar{l}_2, \bar{l}_3, \bar{l}_4$ represents the nominal position of masses when forces are all equal to 0.

Consider the mechanical system reported in Figure ??, where d_i , $i = 1, 2, 3, 4$, represent the mass displacements from an equilibrium position, whereas u_1 , u_2 and u_3 are forces acting on the masses. The state of the system is given by the mass displacements and their derivatives: $x = [d_1, d_2, d_3, d_4, \dot{d}_1, \dot{d}_2, \dot{d}_3, \dot{d}_4]^T$. We set all masses and stiffness constants equal to 1, i.e., $m_1 = m_2 = m_3 = m_4 = 1$ and $k_1 = k_2 = k_3 = k_4 = 1$, and consider a discrete time model of the system $x_{t+1} = Ax_t + Bu_t + w_t$, where the additive noise w_t is an i.i.d. sequence of Gaussian random variables with zero mean and variance $\sigma^2 I_{8 \times 8}$, $\sigma = 0.02$. Matrices A and B are obtained by time discretization of a standard mechanical model under the assumption that the control action is piecewise constant over the intervals $[t, t+1)$.

We consider the case when $\tau = 0$ and the cost to be minimize is given by

$$\mathbb{E} \left[\sum_{i=1}^M x_i^T Q x_i + \sum_{i=0}^{M-1} u_i^T R u_i \right],$$

with a prediction horizon of length $M = 5$ and weight matrices $Q = \begin{bmatrix} I_{4 \times 4} & 0_{4 \times 4} \\ 0_{4 \times 4} & 0_{4 \times 4} \end{bmatrix}$ (penalizing displacements but not their derivatives) and $R = I_{3 \times 3}$.

The initial state is set to $x_0 = [0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0]^T$ and the state and input are subject to the probabilistic constraint

$$\mathbb{P}\{|x_t^{(k)}| \leq 0.45, |u_{t-1}^{(k)}| \leq 0.3, k = 1, 2, 3, |x_t^{(4)}| \leq 1, \\ t = 1, 2, \dots, M\} \geq 1 - \varepsilon,$$

where $v^{(k)}$ denotes the k -th component of vector v .

By adopting the control input parametrization in equation (??) and setting $\varepsilon = 0.1$, $\varepsilon' = 0$ and $\beta = 10^{-5}$, we get $N = 3432$ disturbance realizations to extract. The corresponding scenario-optimization was solved by running YALMIP over SeDuMi, [?]. The performance achieved by the randomized method was $J(\Gamma_N^*, \Theta_N^*) = 1.328$, which is quite close to the best performance (1.325) achieved in [?] by over-approximating the probabilistic constraint. Our method, however, can be applied in full generality, while the approaches in [?] relies on the assumption of Gaussianity of the disturbance w_t .

As for the computing time, the scenario approach is much slower than other approaches. The computing time, however, can be decreased by reducing the number of parameters to be optimized. In the case when $\sigma = 0.02$, the noise contribution is actually so small that the disturbance feedback term Θw can be removed without a significant deterioration of the cost. Indeed, by setting Θ equal to zero and optimizing over Γ only, we obtained a cost $J(\Gamma_N^*) = 1.342$ with a computing time that was 100 times smaller than the previous case. The number of realizations to extract was $N = 451$ this time.

By discarding constraints (i.e. $\varepsilon' \neq 0$), performance in terms of control cost can be improved both when $\Theta \neq 0$ (i.e. with the feedback disturbance term) and when $\Theta = 0$ (i.e. without the feedback disturbance term), at the price of an additional computational effort. Depending on the application context, this can be critical for the on-line receding horizon implementation of the approach.

V. CONCLUSIONS

In this paper, we take a first step towards the application of randomized methods for the solution of chance-constrained optimization problems in stochastic MPC.

The focus of this paper is on the finite-horizon constrained optimization problem that has to be solved at each time step τ to determine the control input to be applied at τ , according to the receding-horizon strategy.

A main aspect that still needs to be addressed concerns the stability of the resulting MPC control scheme, which is not easy to prove for the unbounded disturbance case. Notable exceptions are represented by [?], [?], [?], [?],

where mean square stability is proven for the case when the open-loop system is stable and hard bounds are imposed on the control input. Interestingly, in [?] the hard input bounds are imposed by adopting a feedback control law that is affine in the reconstructed and pre-saturated disturbance samples. Since convexity of the cost and of the constraint function is preserved by this control law parametrization, the technique for proving stability adopted in [?] could be possibly extended to our randomized setting by adopting hard constraint for the control input and probabilistic constraint for the state only. This is subject of ongoing research activity.

REFERENCES

- [1] T. Alamo, R. Tempo, and A. Luque. On the sample complexity of randomized approaches to the analysis and design under uncertainty. In *American Control Conference (ACC)*, Baltimore, Maryland, 2010.
- [2] I. Batina. *Model predictive control for stochastic systems by randomized algorithms*. PhD thesis, Technische Universiteit Eindhoven, 2004.
- [3] A. Bemporad, F. Borrelli, and M. Morari. Min-max control of constrained uncertain discrete-time linear systems. *IEEE Transactions on Automatic Control*, 48(9):1600–1606, 2003.
- [4] A. Bemporad and M. Morari. Robust model predictive control: a survey. In *Robustness in Identification and Control*, volume 245 of *Lecture Notes in Control and Information Sciences*, pages 207–226. Springer, 1999.
- [5] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solution of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.
- [6] G. Calafiore and M.C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.
- [7] G. Calafiore and M.C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.
- [8] G.C. Calafiore and L. Fagiano. Robust model predictive control via random convex programming. In *Proc. 50th IEEE Conf. on Decision and Control and European Control Conference (CDC-ECC)*, pages 1910–1915, Dec. 2011.
- [9] 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.
- [10] M.C. Campi and S. Garatti. A sampling-and-discarding approach to chance-constrained optimization: Feasibility and optimality. *Journal of Optimization Theory and Applications*, 148(2):257–280, 2011.
- [11] M.C. Campi, S. Garatti, and M. Prandini. The scenario approach for systems and control design. *Annual Reviews in Control*, 33(2):149–157, 2009.
- [12] M. Cannon, B. Kouvaritakis, and X. Wu. Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, 45:167–172, 2009.
- [13] M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty. *IEEE Transactions on Automatic Control*, 54(7):1626–1632, 2009.
- [14] D. Chatterjee, J. Lygeros, and P. Hokayem. Stochastic Receding Horizon Control With Bounded Control Inputs: A Vector Space Approach. *IEEE Transactions on Automatic Control*, 56:2704 – 2710, November 2011.
- [15] D. Chatterjee, F. Ramponi, P. Hokayem, and J. Lygeros. On mean square boundedness of stochastic linear systems with bounded controls. *Systems and Control Letters*, 61:375–380, 2012.
- [16] E. Cinquemani, M. Agarwal, D. Chatterjee, and J. Lygeros. Convexity and convex approximations of discrete-time stochastic control problems with constraints. *Automatica*, 47(9):2082–2087, 2011.
- [17] D. Muñoz de la Peña, T. Alamo, A. Bemporad, and F. Camacho. A decomposition algorithm for feedback min-max model predictive control. *IEEE Transactions on Automatic Control*, 51(10):1688–1692, 2006.
- [18] P.J. Goulart, E.C. Kerrigan, and J.M. Maciejowski. Optimization over state feedback policies for robust control with constraints. *Automatica*, 42(4):523–533, April 2006.
- [19] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, Feb 2011.
- [20] P. Hokayem, D. Chatterjee, and J. Lygeros. On Stochastic Receding Horizon Control with Bounded Control Inputs. In *IEEE Conference on Decision and Control*, pages 6359–6364, December 2009.
- [21] P. Hokayem, E. Cinquemani, D. Chatterjee, F. Ramponi, and J. Lygeros. Stochastic receding horizon control with output feedback and bounded controls. *Automatica*, 48(1):77–88, January 2012.
- [22] M. Kothare, V. Balakrishnan, and M. Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10):1361–1379, 1996.
- [23] J. Löfberg. *Minimax approaches to robust model predictive control*. PhD thesis, Linköping studies in Science and Technology, 2003.
- [24] J. Löfberg. Yalmip: a toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [25] A. Prékopa. *Stochastic programming*. Kluwer, Boston, MT, USA, 1995.
- [26] A. Prékopa. Probabilistic programming. In A. Ruszczyński and A. Shapiro, editors, *Stochastic Programming*, volume 10 of *handbooks in operations research and management science*, London, UK, 2003. Elsevier.
- [27] J.A. Primbs and C.H. Sung. Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Transactions on Automatic Control*, 54(2):221–230, 2009.
- [28] D.M. Raimondo, D.Limon, M. Lazar, L. Magni, and E.F. Camacho. Min-max model predictive control of nonlinear systems: A unifying overview on stability. *European Journal of Control*, 15:521, 2009.
- [29] P.O.M. Scokaert and D.Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43:1136–1142, 1998.
- [30] D.H. van Hessem and O.H. Bosgra. Stochastic closed-loop model predictive control of continuous nonlinear chemical processes. *J. Proc. Control*, 16(3):225–241, 2006.