Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Randomized Nonlinear MPC for Uncertain Control-Affine Systems with Bounded Closed-Loop Constraint Violations ⋆

Xiaojing Zhang *, **Sergio Grammatico** *, **Kostas Margellos** *, **Paul Goulart** * and **John Lygeros** *

* *Department of Information Technology and Electrical Engineering, ETH Zurich, Physikstrasse 3, 8092 Switzerland.*
E-mail: {xiaozhan, grammatico, margellos, pgoulart, lygeros}@control.ee.ethz.ch.

Abstract: In this paper we consider uncertain nonlinear control-affine systems with probabilistic constraints. In particular, we investigate Stochastic Model Predictive Control (SMPC) strategies for nonlinear systems subject to chance constraints. The resulting non-convex chance constrained Finite Horizon Optimal Control Problems are computationally intractable in general and hence must be approximated. We propose an approximation scheme which is based on randomization and stems from recent theoretical developments on random non-convex programs. Since numerical solvers for non-convex optimization problems can typically only reach local optima, our method is designed to provide probabilistic guarantees for *any* local optimum inside a set of chosen complexity. Moreover, the proposed method comes with bounds on the (time) average closed-loop constraint violation when SMPC is applied in a receding horizon fashion. Our numerical example shows that the number of constraints of the proposed random non-convex program can be up to ten times smaller than those required by existing methods.

Keywords: scenario approach, randomized methods, non-convex optimization, stochastic control, predictive control

1. INTRODUCTION

Model Predictive Control (MPC) is a powerful control design technique for constrained systems [Mayne et al. 2000]. In the presence of uncertainty and/or disturbances, one possible approach is the robust (i.e. worst-case) MPC design [Kothare et al. 1996].

Recently, there has been increasing interest in Stochastic MPC (SMPC) [Cannon et al. 2009, Primbs and Sung 2009, Chatterjee et al. 2011], which presents an alternative paradigm. An essential feature of SMPC, which also raises theoretical and computational challenges, is that the constraints are addressed in a probabilistic sense. The constraints of the Finite Horizon Optimal Control Problem (FHOCP) arising at each time step can be interpreted probabilistically via chance constraints, allowing for a small constraint violation probability [Shapiro et al. 2009]. Unfortunately, chance constrained optimization problems are non-convex in general and require the computation of multi-dimensional integrals. As a consequence, SMPC is computationally intractable in general.

Randomized MPC (RMPC) [Calafiore and Fagiano 2013, Schildbach et al. 2014, Zhang et al. 2014] is a novel methodology to approximate SMPC problems via the scenario approach whenever the underlying optimization problem is *convex*. Interestingly, the results of the scenario approach hold for any fixed, but possibly unknown, probability distribution [Calafiore and Campi 2006]. In RMPC, the chance constraints in the FHOCP are replaced by a finite number of hard constraints, each corresponding to independent realizations of the uncertainty. The number of drawn scenarios, called *sample size*, is chosen so that, with high confidence, the violation probability of the solution of the sampled FHOCP remains below the desired threshold.

RMPC has been indeed effectively exploited for uncertain *linear* systems because feasibility, optimality and sample complexity of random *convex* programs are well characterized [Calafiore and Campi 2006, Campi and Garatti 2008, Calafiore 2010]. On the other hand, to the best of the authors' knowledge, scenario approaches for random *non-convex* programs have not been analyzed to a great extend, nor randomized methods for the stochastic control of uncertain *nonlinear* systems by means of RMPC.

In this paper, we apply recent developments in the field of random *non-convex* programs [Grammatico et al. 2014a,b] to stochastic control of uncertain nonlinear control-affine systems via Randomized Nonlinear MPC (RNMPC). In contrast to existing work, and inspired by [Schildbach et al. 2014], we interpret the constraint violations in closed-loop as average in time, and not pointwise in time. Therefore, we show that the proposed RNMPC algorithm comes with bounds on the (time) average closed-loop

⋆ Research partially supported by Nano-Tera.ch under the project HeatReserves and the KIOS center of excellence 0308/28.

constraint violations. The algorithm is illustrated on a tracking problem for a perturbed unicycle model.

Notation

$\mathbb{Z}[a, b]$ is the integer interval $\{a, a+1, ..., b\} \subseteq \mathbb{Z}$, $\text{conv}(\cdot)$ the convex hull, and $\mathbf{1}_{\mathcal{S}}(\cdot)$ the indicator function over a set $\mathcal{S}$. If $(\Delta, \mathcal{F}, \mathbb{P})$ is a probability space, then $\mathbb{E}$ is the associated expectation and $\mathbb{P}^N$ the N-fold product measure.

2. PROBLEM DESCRIPTION

2.1 Stochastic Control Problem

We consider a discrete-time uncertain control affine system

$$x^+ = f(x, \delta) + g(x, \delta)u, \tag{1}$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ the control input, $\delta \in \Delta \subseteq \mathbb{R}^d$ the uncertainty, and $f : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^n$ and $g : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^{n\times m}$ are measurable functions. We assume that δ is a random variable defined on a probability space $(\Delta, \mathcal{F}, \mathbb{P})$. Neither the support set Δ, nor the probability measure $\mathbb{P}$ need to be known explicitly. However, we assume that i.i.d. samples $\delta^{(1)}, \delta^{(2)}, \dots$ are available, and that they are independent in time. Measurability issues are glossed over in this paper, and we refer to [Grammatico et al. 2014a] for such technical details.

Let $\epsilon \in (0, 1)$ be the admissible average closed-loop constraint violation probability, and $\mathbb{X} \subset \mathbb{R}^n$ and $\mathbb{U} \subset \mathbb{R}^m$ two compact convex sets which contain the origin in their relative interior. Let $x(0)$ be a fixed initial state, and $x(t)$ and $u(t)$ the state and input at time t, respectively. Let $\kappa : \mathbb{R}^n \to \mathbb{R}^m$ be a feedback control law which, with $u(t) := \kappa(x(t))$, generates an input sequence $\{u(0), u(1), \dots\}$, resulting in the closed-loop state sequence $\{x(1), x(2), \dots\}$. We are interested in constructing a law κ satisfying the following property.

Definition 1. A feedback control law $\kappa : \mathbb{R}^n \to \mathbb{U}$ is called *probabilistically feasible* if, for all $T > 0$, the closed-loop state sequence generated by κ satisfies

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \mathbf{1}_{\mathbb{X}}\left(x(t)\right)\right] \geq 1 - \epsilon.$$

□

Note that for the sake of simplicity, there is a slight abuse of notation in the above inequality, since the dependence of $x(t)$ on the random input sequence $\{\delta(0), \delta(1), \dots, \delta(t-1)\}$ is not made explicit, and we actually have $\mathbb{E}^T$ in place of $\mathbb{E}$. One way of constructing such a probabilistically feasible control law is by means of Stochastic MPC (SMPC), where at each time step a chance constrained finite horizon optimal control problem (ccFHOCP) has to be solved.

2.2 Stochastic MPC

Let $\boldsymbol{\delta} := (\delta_0, \delta_1, \dots, \delta_{N-1}) \in \Delta^N$ be a random disturbance sequence of length N, and $\mathbf{u} := (u_0, u_1, ..., u_{N-1}) \in \mathbb{U}^N$ the planned input sequence. If $x(t)$ is the current state at time t, we denote by $\phi(i; x(t), \mathbf{u}, \boldsymbol{\delta})$ the *predicted* state at time $t+i$, evolved according to (1) subject to $\mathbf{u}$ and $\boldsymbol{\delta}$. Let $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ be a convex stage cost, and $\ell_f : \mathbb{R}^n \to \mathbb{R}$ be a convex terminal cost. Then, we consider the stochastic cost function

$$J(x(t), \mathbf{u}, \boldsymbol{\delta}) := \ell_f\left(\phi(N; x(t), \mathbf{u}, \boldsymbol{\delta})\right) + \sum_{i=0}^{N-1} \ell\left(\phi(i; x(t), \mathbf{u}, \boldsymbol{\delta}), u_i\right). \tag{2}$$

Since J is a random variable, typical methods to obtain a deterministic cost function rely on using $J(x(t), \mathbf{u}, \mathbb{E}^N[\boldsymbol{\delta}])$ or $\mathbb{E}^N\left[J(x(t), \mathbf{u}, \boldsymbol{\delta})\right]$.

The SMPC algorithm proceeds now in a receding horizon fashion as follows. At every time step t, it first solves the following ccFHOCP:

$$\begin{aligned} \min_{\mathbf{u}\in\mathbb{U}^N} \ & \mathbb{E}^N\left[J(x(t), \mathbf{u}, \boldsymbol{\delta})\right] \\ \text{s.t.} \ & \mathbb{P}^N\left[\phi\left(i; x(t), \mathbf{u}, \boldsymbol{\delta}\right) \in \mathbb{X}\right] \geq 1 - \epsilon \quad \forall i \in \mathbb{Z}[1, N]. \end{aligned} \tag{3}$$

Given the optimal planned input $\mathbf{u}^\star$, SMPC then applies the input $u(t) = \kappa(x(t)) := u_0^\star$. The whole process is then repeated at time $t + 1$. We assume that the ccFHOCP in (3) is feasible at every time instant. Then it will be clear later on that the control law constructed by certain SMPC algorithms is probabilistically feasible, see e.g. [Schildbach et al. 2014, Theorem 15].

In general, the above ccFHOCP problem requires multi-dimensional integrals and is computationally intractable. In the following section, we describe three different randomized algorithms which approximately solve the ccFHOCP problem in (3), while producing a probabilistically feasible control law κ.

3. RANDOMIZED MPC

One way to approximate (3) is by sampling the uncertainty sequence $\boldsymbol{\delta}$, and replacing the chance constraints and objective function with a finite number of uncertainty realizations. This method is known as Randomized MPC (RMPC) or Scenario-based MPC [Calafiore and Fagiano 2013, Schildbach et al. 2014, Zhang et al. 2014]. A variation thereof, which uses a combination of robust and randomized methods, is described in [Zhang et al. 2013a]. These methods are based on the scenario approach introduced in [Calafiore and Campi 2006, Campi and Garatti 2008, Calafiore 2010]. The scenario approach is a way to approximate chance constraints whenever the underlying problem is convex, i.e. the cost and the constraints are convex functions for any fixed uncertainty.

We first review the general RMPC setup: let (S_0, S_1, S_2) be non-negative integers, and assume that $(S_0 + S_1 + S_2)$ full-horizon samples are drawn independently according to $\mathbb{P}^N$. We define three sets of multisamples $\boldsymbol{\omega}_0 := \{\boldsymbol{\delta}^{(1)}, \dots, \boldsymbol{\delta}^{(S_0)}\}$, $\boldsymbol{\omega}_1 := \{\boldsymbol{\delta}^{(S_0+1)}, \dots, \boldsymbol{\delta}^{(S_0+S_1)}\}$, and $\boldsymbol{\omega}_2 := \{\boldsymbol{\delta}^{(S_0+S_1+1)}, \dots, \boldsymbol{\delta}^{(S_0+S_1+S_2)}\}$. The collection $\boldsymbol{\omega}_0$ is used to empirically approximate the cost function J; the samples $\boldsymbol{\omega}_1$ are used to enforce the state constraint for the first predicted step $i = 1$; and finally, the samples in $\boldsymbol{\omega}_2$ are used to enforce the constraints for the remaining predicted stages $i = 2, 3, \dots, N$.

Consider the following sampled counterpart of (3):

$$\text{RMPC}[\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \boldsymbol{\omega}_2]: \tag{4}$$
$$\begin{cases} \min_{\mathbf{u}\in\mathbb{U}^N} & \sum_{\boldsymbol{\delta}^{(j)}\in\boldsymbol{\omega}_0} J(x(t),\mathbf{u},\boldsymbol{\delta}^{(j)}) \\ \text{s.t.} & \phi(1;x(t),\mathbf{u},\boldsymbol{\delta}^{(k)}) \in \mathbb{X} \quad \forall \boldsymbol{\delta}^{(k)} \in \boldsymbol{\omega}_1, \\ & \phi(i;x(t),\mathbf{u},\boldsymbol{\delta}^{(l)}) \in \mathbb{X} \quad \forall i \in \mathbb{Z}[2,N],\ \forall \boldsymbol{\delta}^{(l)} \in \boldsymbol{\omega}_2. \end{cases}$$

In the SMPC framework, the RMPC problem in (4) must be solved at each time step with an updated initial state $x(t)$. For the receding horizon approach to run adequately in closed-loop, we postulate the following assumption which holds throughout this paper, adopted from [Schildbach et al. 2014, Assumption 5].

Assumption 1. (Recursive feasibility). At every time instance t, the RMPC problem in (4) admits a feasible solution for every realization of $(\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \boldsymbol{\omega}_2)$ almost surely.

The main challenge in RMPC is to find the sample sizes S_0, S_1, and S_2 such that the closed-loop violation probability is below a desired level ϵ. We next discuss three different methods to solve RMPC in (4).

3.1 Randomized Linear MPC

One method to deal with the ccFHOCP in (3) is to linearize the dynamics and solve a linearized version of (4), see e.g. [Zhang et al. 2013b] for such an application. The nonlinear dynamics in (1) are approximated around some fixed $\bar{x}$ by $A(\bar{x},\delta)(x-\bar{x}) + B(\bar{x},\delta)u$, rendering the RMPC problem in (4) convex. In this case, the standard scenario approach for convex programs can be applied.

More precisely, let $\phi_{\text{lin}}(i; x(t), \mathbf{u}, \boldsymbol{\delta})$ be the predicted state of the linearized dynamics, and $J_{\text{lin}}(x(t), \mathbf{u}, \boldsymbol{\delta})$ the corresponding cost. Then we can consider the Randomized Linear MPC (RLMPC):

$$\text{RLMPC}[\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \boldsymbol{\omega}_2]: \tag{5}$$
$$\begin{cases} \min_{\mathbf{u}\in\mathbb{U}^N} & \sum_{\boldsymbol{\delta}^{(j)}\in\boldsymbol{\omega}_0} J_{\text{lin}}(x(t),\mathbf{u},\boldsymbol{\delta}^{(j)}) \\ \text{s.t.} & \phi_{\text{lin}}(1;x(t),\mathbf{u},\boldsymbol{\delta}^{(k)}) \in \mathbb{X} \quad \forall \boldsymbol{\delta}^{(k)} \in \boldsymbol{\omega}_1, \\ & \phi_{\text{lin}}(i;x(t),\mathbf{u},\boldsymbol{\delta}^{(l)}) \in \mathbb{X} \quad \forall i \in \mathbb{Z}[2,N],\ \forall \boldsymbol{\delta}^{(l)} \in \boldsymbol{\omega}_2. \end{cases}$$

Consider now Algorithm 1.

Algorithm 1 Randomized Linear MPC (RLMPC)

Require: $\epsilon \in (0,1)$; $S_0, S_2 \in \mathbb{Z}[1,\infty)$.
1: Let
$$S_1 := \left\lceil \frac{m}{\epsilon} - 1 \right\rceil.$$
2: Extract $S_0 + S_1 + S_2$ i.i.d. samples according to $\mathbb{P}^N$.
3: Compute the optimal point for RLMPC in (5).
4: Apply the first control input $u(t) = \kappa(x(t)) := u_0^\star$.

It follows from [Schildbach et al. 2014, Theorem 15] that the control law κ constructed by Algorithm 1 (RLMPC) is probabilistically feasible.

Remark 1. Out of the three sample sizes S_0, S_1, S_2, only S_1 is required to take some minimum value when constructing a probabilistically feasible policy. The variables S_0 and S_2 are allowed to take any values and can be viewed as tuning variables. However, as illustrated in Section 5.1, we show that the choice of S_2 indeed influences the closed-loop cost. □

3.2 Randomized Nonlinear MPC

The performance of the controller based on the RLMPC algorithm greatly depends on the accuracy of the linearization step. In full generality, however, the linearization process is not appropriate. Examples are the discrete non-holonomic integrator, which is not controllable if linearized around the origin, or the unicycle model presented in Section 5. For such systems, the nonlinear dynamics in (1) must be kept, at the cost that a non-convex RMPC problem must be solved. One way to determine the appropriate sample sizes is based on the sampling method for random non-convex optimization problems as described in [Grammatico et al. 2014a,b] and presented next.

Let $M \geq m+1$ be an integer, and consider M different vectors $c_1, c_2, ..., c_M \in \mathbb{R}^m$. Extract $S_1 + S_2 + S_3$ samples for $\boldsymbol{\omega}_0$, $\boldsymbol{\omega}_1$, and $\boldsymbol{\omega}_2$. For each c_i, $i \in \mathbb{Z}[1,M]$, consider the following ith auxiliary scenario program SP_i:

$$\text{SP}_i[\boldsymbol{\omega}_1]: \tag{6}$$
$$\begin{cases} \min_{v\in\mathbb{U}} & c_i^\top v \\ \text{s.t.} & f(x(t),\delta^{(k)}) + g(x(t),\delta^{(k)})v \in \mathbb{X} \quad \forall \delta^{(k)} \in \omega_1, \end{cases}$$

where ω_1 is $\boldsymbol{\omega}_1$ with each element restricted to its first uncertainty component, i.e. $\omega_1 := \left\{\delta_0^{(S_0+1)}, \ldots, \delta_0^{(S_0+S_1)}\right\}$. Given $v_i^\star$ being the optimizer of $\text{SP}_i[\boldsymbol{\omega}_1]$, we define

$$\mathbb{U}_M(\boldsymbol{\omega}_1) := \text{conv}\left(\{v_1^\star, \ldots, v_M^\star\}\right) \subseteq \mathbb{U}. \tag{7}$$

The set $\mathbb{U}_M(\boldsymbol{\omega}_1)$ has the property that, loosely speaking, any optimizer belonging to it is probabilistically feasible. Precisely, consider the following approximation of RMPC in (4), called Randomized Nonlinear MPC (RNMPC):

$$\text{RNMPC}[\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \boldsymbol{\omega}_2]: \tag{8}$$
$$\begin{cases} \min_{\mathbf{u}\in\mathbb{U}^N} & \sum_{\boldsymbol{\delta}^{(j)}\in\boldsymbol{\omega}_0} J(x(t),\mathbf{u},\boldsymbol{\delta}^{(j)}) \\ \text{s.t.} & u_0 \in \mathbb{U}_M(\boldsymbol{\omega}_1), \\ & \phi(i;x(t),\mathbf{u},\boldsymbol{\delta}^{(l)}) \in \mathbb{X} \quad \forall i \in \mathbb{Z}[2,N],\ \forall \boldsymbol{\delta}^{(l)} \in \boldsymbol{\omega}_2. \end{cases}$$

Let the integer $\zeta \leq m$ be the so-called Helly's dimension of $\text{SP}_i[\boldsymbol{\omega}_1]$ in (6) [Calafiore 2010, Definition 3.1], which can always be upper bounded by m. Moreover, we define $\Phi(\nu,\zeta,S) := \sum_{i=0}^{\zeta-1}\binom{S}{i}\nu^i(1-\nu)^{S-i}$. In view of (6)–(8), consider now Algorithm 2.

Algorithm 2 Randomized Nonlinear MPC (RNMPC)

Require: : $\epsilon \in (0,1)$; $S_0, S_2 \in \mathbb{Z}[1,\infty)$; $M \in \mathbb{Z}[m+1,\infty)$; $c_1, \ldots, c_M \in \mathbb{R}^m$.
1: Compute the smallest integer S_1 such that
$$\int_0^1 \min\left\{1, \binom{M}{m+1}\Phi\left(\nu,(m+1)\zeta,S_1\right)\right\} d\nu \leq \epsilon \tag{9}$$
2: Extract $S_0 + S_1 + S_2$ i.i.d. samples according to $\mathbb{P}^N$.
3: For all $i \in \mathbb{Z}[1,M]$, solve $\text{SP}_i[\boldsymbol{\omega}_1]$ in (6) and construct $\mathbb{U}_M(\boldsymbol{\omega}_1)$ as in (7).
4: Compute a feasible point (e.g. local minimum) for RNMPC in (8).
5: Apply the first control input $u(t) = \kappa(x(t)) := u_0^\star$.

We then have the following result whose proof is given in the Appendix.

Theorem 1. The control law κ generated by Algorithm 2 (RNMPC) is probabilistically feasible. □

Similar to Algorithm 1 (RLMPC), S_0 and S_1 can be treated as tuning variables in Algorithm 2. Moreover, the user defined integer M determines how "complex" the set $\mathbb{U}_M(\boldsymbol{\omega}_1)$, see Section 4.

Algorithm 2 (RNMPC), together with Theorem 1, can be viewed as an extension of the RLMPC algorithm to the class of uncertain *nonlinear* control-affine systems. The key difference between Algorithm 2 (RNMPC) and Algorithm 1 (RLMPC), however, is that the probabilistic guarantees in RLMPC hold for the global optimizer, whereas in RNMPC *any feasible* solution comes with the guarantees. Even though the latter property renders the solution somewhat more conservative, it is of great practical importance. This is because in general one cannot hope to solve for the global optimum in a non-convex optimization problem, but only for a local optimum.

3.3 Randomized MPC via Statistical Learning Theory

The third method we look at is called RMPC-SLT. It is similar to the RNMPC method described in the previous section, but solves the non-convex RMPC problem in (4) directly. It is based on the Vapnik-Chervonenkis (VC) theory of statistical learning, see e.g. [Alamo et al. 2009]. More precisely, let us define

$$\Psi_\xi(\nu, S) := \min\left\{1, \left(\frac{2eS}{\xi}\right)^{\xi} 2^{1-\nu S/2}\right\}, \qquad (10)$$

where $\xi \in \mathbb{Z}[1, \infty)$ is an upper bound on the so-called VC-dimension [Alamo et al. 2009, Definition 6]. Then the following theorem states the minimum sample size S_1 such that the resulting control law is probabilistically feasible.

Theorem 2. Assume that ξ is an upper bound for the VC-dimension associated to (4). Let $u(t) = \kappa(x(t)) := u_0^\star$ be the control law given by any feasible solution of RMPC in (4). If S_1 satisfies

$$\int_0^1 \Psi_\xi(\nu, S_1) d\nu \leq \epsilon,$$

then the resulting control law κ is probabilistically feasible. □

Typically, RMPC-SLT comes with a much larger sample size than Algorithm 2 (RNMPC). This not only results in a more conservative control law, but also makes it computationally demanding to even find a local minimum, see also the case study in Section 5.

4. DISCUSSION

4.1 Choice of M in RNMPC

From Algorithm 2 it can be seen that any choice of $M \geq m + 1$ results in a probabilistically feasible controller κ, provided S_1 is chosen according to (9). Roughly speaking, M determines the complexity of the polyhedron $\mathbb{U}_M$ in (7), which is defined by the M optimizers $u_1^\star, \ldots, u_M^\star$. Although a large M results in a large admissible search space $\mathbb{U}_M$ for RNMPC in (8), it should be kept in mind that the sample size S_1 also increases with M because of (9). The dependence of S_1 on M is illustrated in Figure 1 for the case $\zeta = m = 2$, and it can be actually shown that S_1 depends on M as $\sim \ln(M)$ [Grammatico et al. 2014a]. For M fixed, the issue of choosing the costs $c_1, \ldots, c_M$ arises. Roughly speaking, they should be chosen such that $\mathbb{U}_M$ is as big as possible. One way of doing that is by letting the vectors c_i point into directions uniformly distributed on a hypersphere in dimension m. Since the dependence on M is logarithmic, it can be chosen relatively large without having an adverse impact on S_1.

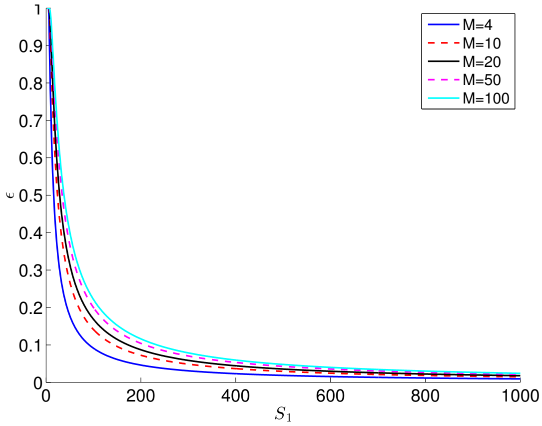

Fig. 1. Bound on the closed-loop violation probability for RNMPC with $\zeta = m = 2$ as a function of S_1.

4.2 Comparison between the three RMPC Algorithms

RLMPC is a standard approach for many multi-stage optimal control problems with linear dynamics. This is because the resulting optimization problem is convex and the sample size in Algorithm 1 is much lower compared to RNMPC and RMPC-SLT, see [Zhang et al. 2013b] for an example where RLMPC is applied. In cases where the RLMPC algorithm does not perform as desired, one is left with the choice between the RNMPC algorithm and RMPC-SLT. The latter is only applicable to problems for which a finite upper bound on the VC-dimension can be found. In contrast, our methodology RNMPC is applicable to all RMPC problems of the form (5).

Let now ξ be such an upper bound for the VC-dimension. Figure 2 depicts the required sample sizes for some realistic values of ξ. The corresponding sample sizes are much higher than those of RNMPC, rendering the (non-convex) RMPC problem in (4) computationally demanding to solve. Morever, it is in general difficult to estimate ξ, making the RMPC-SLT approach not applicable in general.

5. STOCHASTIC CONTROL OF THE DISCRETE-TIME UNICYCLE

In this section, we consider a case study of the system

$$\dot{x} = \cos(\theta)v, \quad \dot{y} = \sin(\theta)v, \quad \dot{\theta} = w, \qquad (11)$$

which is the continuous model of a unicycle [M'Closkey and Murray 1994]. The states (x, y) represent the Cartesian position of the center of mass of the car, whereas θ

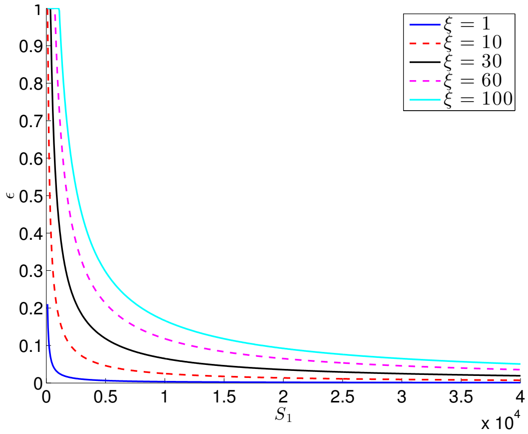

Fig. 2. Upper bound on the closed-loop violation probability for RNMPC-LST as a function of S_1.

is the angle of the car with respect to the x axis. The controlled inputs $v \in \mathbb{R}$ and $w \in \mathbb{R}$ are the forward and angular velocities, respectively.

The unicycle in (11) is not straightforward to stabilize to the origin because it is a nonholonomic system and violates Brockett's necessary condition for stabilization via continuous feedback [Brockett 1983]. Since MPC can automatically generate a discontinuous control law, it is a natural candidate for nonholonomic systems. We next study a tracking problem where we investigate the influence of S_2 on the empirical closed-loop cost and violation probability.

5.1 Tracking a Circle

It is intuitively clear that the exact choice of S_2 for all three algorithms will influence the closed-loop cost, which is investigated next. We thus define $(x_1,\ x_2,\ x_3) := (x,\ y,\ \theta)$, $(u_1,\ u_2) := (v,\ w)$, let $\delta = (\delta_1,\ \delta_2) \in \mathbb{R}^2$, and discretize (11) with Euler step $\tau = 0.1$. The system dynamics then read as

$$\begin{aligned} x_1^+ &= x_1 + \tau\left(\cos(x_3)u_1 + \delta_1\right) \\ x_2^+ &= x_2 + \tau\left(\sin(x_3)u_1 + \delta_2\right) \\ x_3^+ &= x_3 + \tau u_2, \end{aligned}$$

where (δ_1, δ_2) are uniformly distributed on $[-0.4, 0.4]^2$. The input constraints are $|u_1| \leq 2,\quad |u_2| \leq \pi/\tau$. We consider the constraint set

$$\mathbb{X} := \{(x_1,\ x_2,\ x_3) \in \mathbb{R}^3 \mid \|(x_1,\ x_2)\|_2 \leq 10\}.$$

Let $\{x^\star(t) := (10\cos\theta(t),\ 10\sin\theta(t),\ \theta(t) + \pi/2)\}_{t\geq 0}$ denote the sequence of states to be tracked, for a sequence $\{\theta(t) \in [0, 2\pi]\}_{t\geq 0}$, which generates a circle for an unperturbed unicycle.

The control objective is to follow $\{x^\star(1), x^\star(2), \ldots\}$ while exiting $\mathbb{X}$ less than 40% of the time. The cost functions are

$$\ell(x,u) = \|Q(x - x^\star)\|_1 + \|Ru\|_1, \quad \ell_f(x) = \|Q_f(x - x^\star)\|_1$$

with $Q = I$, $Q_f = 10Q$, and $R = 0.1I$. We use $J(x(t), \mathbf{u}, \mathbf{0})$ as cost and $x(0) = (9.9,\ 0,\ \pi/2)$ as initial condition. For the predefined bound on the closed-loop violation probability of 40%, we obtain the sample size $S_1 = 29$ from Algorithm 2 (RNMPC) [1]. It can be verified that

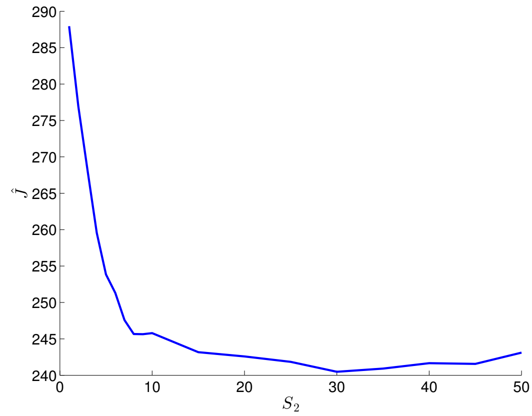

Fig. 3. Empirical closed-loop cost.

the required sample size for RMPC-SLT is 285, which is almost ten times the sample size required by RNMPC. Thus, RMPC-SLT is not considered here.

To assess the influence of S_2 on the closed-loop cost for RNMPC, we simulate the described system for different values of S_2 ranging from 1 to 50, where for each value of S_2 ten runs are conducted. Figure 3 depicts the empirical closed-loop cost as a function of S_2, for fixed $S_1 = 29$. It can be observed that for $S_2 = 1, 2, \ldots, 30$, the cost decreases monotonically. For $S_2 > 30$, the cost starts to increase again, producing overall a plot with a clear global minimum. This can be justified as follows: each time the unicycle leaves $\mathbb{X}$, it has to spend additional time and effort to enter $\mathbb{X}$ again (instead of tracking the circle), thus generating additional cost. Therefore, it is beneficial not to exit $\mathbb{X}$ too often, explaining the decrease in cost for $S_2 = 1, 2, \ldots, 30$. On the other hand, if S_2 is increased further, the solution becomes more robust (and conservative). Therefore, the car tends to always stay away from the boundary of $\mathbb{X}$, which causes the cost to increase.

The above results demonstrate that S_2 indeed influences the closed-loop cost, as one would expect. The amount of impact depends on the individual applications, and should be determined individually for each system. This opens the more general question whether an optimal choice of S_2 exists, which minimizes the closed-loop cost. This issue is subject to future research.

6. CONCLUSION

We have considered the stochastic control problem for non-linear control-affine systems subject to chance constraints. We applied randomization as a tool to effectively approximate the corresponding non-convex chance constrained Finite Horizon Optimal Control Problem. As shown in the numerical example, the proposed Algorithm, referred to as RNMPC (Algorithm 2), requires in certain cases a much smaller sample size than existing methods. When the RNMPC algorithm is implemented in receding horizon, we have shown that the resulting control law is probabilistically feasible with bounds on the average closed-loop constraint violation.

[1] S_1 is obtained by setting $\zeta = 1$ and $\epsilon = 20\%$ in (9). We set $\epsilon = 20\%$ because once the unicycle is outside $\mathbb{X}$, it takes two steps to return into $\mathbb{X}$, effectively doubling ϵ.

REFERENCES

Alamo, T., Tempo, R., and Camacho, E.F. (2009). Randomized strategies for probabilistic solutions of uncertain feasibility and optimazation problems. *IEEE Trans. on Automatic Control*.

Brockett, R. (1983). Asymptotic stability and feedback stabilization. *Differential Geometric Control Theory*, 181–208.

Calafiore, G. and Campi, M.C. (2006). The scenario approach to robust control design. *IEEE Trans. on Automatic Control*, 51(5).

Calafiore, G.C. (2010). Random convex programs. *SIAM Journal on Optimization*, 20(6), 3427–3464.

Calafiore, G.C. and Fagiano, L. (2013). Robust model predictive control via scenario optimization. *IEEE Trans. on Automatic Control*, 58.

Campi, M. and Calafiore, G.C. (2009). Notes on the scenario design approach. *IEEE Trans. on Automatic Control*, 54(2), 382–385.

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

Cannon, M., Kouvaritakis, B., and Wu, X. (2009). Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, 45, 167–172.

Chatterjee, D., Hokayem, P., and Lygeros, J. (2011). Stochastic receding horizon control with bounded control inputs: a vector space approach. *IEEE Trans. on Automatic Control*, 56(11), 2704–2710.

Grammatico, S., Zhang, X., Margellos, K., Goulart, P., and Lygeros, J. (2014a). A scenario approach for non-convex control design. *IEEE Trans. on Automatic Control (submitted)* Available online at: http://arxiv.org/abs/1401.2200.

Grammatico, S., Zhang, X., Margellos, K., Goulart, P., and Lygeros, J. (2014b). A scenario approach to non-convex control design: preliminary probabilistic guarantees. In *IEEE American Control Conference*. Portland, Oregon, USA.

Kothare, M.V., Balakrishnan, V., and Morari, M. (1996). Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10), 1361–1379.

Mayne, D.Q., Rawlings, J., Rao, C., and Scokaert, P. (2000). Constrained model predictive control: stability and optimality. *Automatica*, 36, 789–814.

M'Closkey, R.T. and Murray, R.M. (1994). Experiments in exponential stabilization of a mobile robot towing a trailer. In *IEEE American Control Conference, 1994*, volume 1, 988–993.

Primbs, J.A. and Sung, C.H. (2009). Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Trans. on Automatic Control*, 54(2).

Schildbach, G., Fagiano, L., Frei, C., and Morari, M. (2014). The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica (provisionally accepted)*. Available online at: http://arxiv.org/abs/1307.5640.

Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on Stochastic Programing. Modeling and Theory.* SIAM and Mathematical Programming Society.

Zhang, X., Grammatico, S., Schildbach, G., Goulart, P., and Lygeros, J. (2014). On the Sample Size of Randomized MPC for Chance-Constrained Systems with Application to Building Climate Control. In *Proc. of the IEEE European Control Conference*.

Zhang, X., Margellos, K., Goulart, P., and Lygeros, J. (2013a). Stochastic Model Predictive Control Using a Combination of Randomized and Robust Optimization. In *Proc. of the IEEE Conf. on Decision and Control*. Florence, Italy.

Zhang, X., Schildbach, G., Sturzenegger, D., and Morari, M. (2013b). Scenario-based MPC for Energy-Efficient Building Climate Control under Weather and Occupancy Uncertainty. In *Proc. of the IEEE European Control Conference*. Zurich, Switzerland.

Appendix A. PROOFS

Proof of Theorem 1

Step 1: First Stage Violation Probability Fix an initial state x and consider any $\mathbf{u} \in \mathbb{U}^N$. We define the first stage violation probability $V(x, \mathbf{u})$ as

$$V(x,\mathbf{u}) := \mathbb{P}^N[\phi(1; x, \mathbf{u}, \boldsymbol{\delta}) \notin \mathbb{X}].$$

For $S := S_0 + S_1 + S_2$, we prove the following statement.

Lemma 1. Under Assumption 1 it holds for any feasible $\mathbf{u}^\star = \mathbf{u}^\star(\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \boldsymbol{\omega}_2)$ of RNMPC in (8) that

$$\mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star) > \nu\right] \leq \binom{M}{m+1}\Phi(\nu,(m+1)\zeta,S_1).$$

□

Proof 1. For any $\mathbf{u}^\star$, we first compute its conditional probability $\mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star)|\boldsymbol{\omega}_0,\boldsymbol{\omega}_2\right]$. To do this, consider

$$\tilde{\mathrm{SP}}[\boldsymbol{\omega}_1|\boldsymbol{\omega}_0,\boldsymbol{\omega}_2] : \qquad \text{(A.1)}$$

$$\begin{cases} \min\limits_{\mathbf{u}\in\mathbb{U}^N} & \tilde{J}(x,\mathbf{u}|\boldsymbol{\omega}_0,\boldsymbol{\omega}_2) \\ \text{s.t.} & u_0 \in \mathbb{U}_M(\boldsymbol{\omega}_1), \end{cases}$$

where we define

$$\tilde{J}(x,\mathbf{u}|\boldsymbol{\omega}_0,\boldsymbol{\omega}_2) := \sum_{\boldsymbol{\delta}^{(j)}\in\boldsymbol{\omega}_0} J(x,\mathbf{u},\boldsymbol{\delta}^{(j)}) + \sum_{\boldsymbol{\delta}^{(l)}\in\boldsymbol{\omega}_2}\sum_{i=2}^{N}\chi\left[\phi(i;x,\mathbf{u},\boldsymbol{\delta}^{(l)})\in\mathbb{X}\right],$$

where $\chi[S] = 0$ if the statement S is true, and ∞ otherwise. We notice that for fixed $(\boldsymbol{\omega}_0, \boldsymbol{\omega}_2)$, $\tilde{\mathrm{SP}}[\boldsymbol{\omega}_1|\boldsymbol{\omega}_0,\boldsymbol{\omega}_2]$ in (A.1) can be interpreted as an optimization problem with random *convex* constraint $u_0 \in \mathbb{U}_M(\boldsymbol{\omega}_1)$ and *deterministic* non-convex cost $\tilde{J}$. Hence it follows from [Grammatico et al. 2014a] that

$$\mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star) > \nu \mid \boldsymbol{\omega}_0,\boldsymbol{\omega}_2\right] \leq \binom{M}{m+1}\Phi(\nu,(m+1)\zeta,S_1).$$

Namely, S_1 enters the probabilistic guarantees because the corresponding samples $\boldsymbol{\omega}_1$ are used to constrain the first stage violation probability. The proof of Lemma 1 then follows by integrating over $(\boldsymbol{\omega_0}, \boldsymbol{\omega_2}) \in \Delta^{(S_0+S_2)N}$:

$$\begin{aligned} &\mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star) > \nu\right] \\ &= \int_{\Delta^{(S_0+S_2)N}} \mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star) > \nu \mid \boldsymbol{\omega}_0,\boldsymbol{\omega}_2\right] \mathrm{d}\mathbb{P}^{(S_0+S_2)N} \\ &\leq \int_{\Delta^{(S_0+S_2)N}} \binom{M}{m+1}\Phi(\nu,(m+1)\zeta,S_1)\,\mathrm{d}\mathbb{P}^{(S_0+S_2)N} \\ &= \binom{M}{m+1}\Phi(\nu,(m+1)\zeta,S_1). \end{aligned}$$

■

Step 2: Expected Violation Probability Given Lemma 1, it follows from [Campi and Calafiore 2009] that, for all x and $\mathbf{u}^\star$, the expected first stage violation probability $\mathbb{E}^{SN}[V(x,\mathbf{u}^\star)]$ can be upper bounded by the integral on the left-hand-side of (9), where the term $\min\{1,\cdot\}$ is due to the fact that the probability of any event is no more than 1. The proof can now be concluded by following the steps of [Schildbach et al. 2014, Proof of Theorem 15], which shows that the average closed-loop violation probability is upper bounded by the expected first stage violation probability.

Proof of Theorem 2

The proof follows the lines of that of Theorem 1, with the only difference that for all $\mathbf{u}^\star$

$$\mathbb{P}^{SN}\left[V(x,\mathbf{u}^\star) > \nu \mid \boldsymbol{\omega}_0,\boldsymbol{\omega}_2\right] \leq \left(\frac{2eS_1}{\xi}\right)^{\xi} 2^{1-\nu S_1/2},$$

see e.g. [Alamo et al. 2009] and the references therein. The remainder of the proof is omitted in the interest of space.