

Adaptively Constrained Stochastic Model Predictive Control for Closed-Loop Constraint Satisfaction

Frauke Oldewurtel*, David Sturzenegger†, Peyman Mohajerin Esfahani†
Göran Andersson*, Manfred Morari†, John Lygeros†

Abstract—Stochastic Model Predictive Control (SMPC) for discrete-time linear systems subject to additive disturbances with chance constraints on the states and hard constraints on the inputs is considered. Current chance constrained MPC methods—based on analytic reformulations or on sampling approaches—tend to be conservative partly because they fail to exploit the predefined violation level in closed-loop. For many practical applications, this conservatism can lead to a loss in performance.

We propose an adaptive SMPC scheme that starts with a standard conservative chance constrained formulation and then on-line adapts the formulation of constraints based on the experienced violation frequency. Using martingale theory we establish guarantees of convergence to the desired level of constraint violation in closed-loop for a special class of linear systems. Comments are given on how to extend this to a broader class of (non-)linear systems. The developed methodology is demonstrated with an illustrative example.

Keywords: Stochastic model predictive control; Chance constraints, Adaptive control; Closed-loop violation.

I. INTRODUCTION

Robust Model Predictive Control (Robust MPC) has been successfully employed for systems with uncertainties in various application areas over the last decades [1], [2]. Its key strength is to guarantee stability and recursive feasibility as long as the disturbance stays within an a priori defined bounded disturbance set. Besides ‘classic’ Robust MPC other advanced formulations have been developed, e.g., *Affine Disturbance Feedback MPC* [3]–[5] or *Tube-based MPC* [2], [6].

For systems where the uncertainty is known to be in a bounded set, e.g., when the uncertainty stems from sensor readings and the manufacturer guarantees the signals to be in a given range, this approach is very powerful. However, for many practical applications it is hard to specify a bounded disturbance set a priori. This is problematic in both possible cases: if the disturbance realizations are larger than assumed a priori, all guarantees on stability and recursive feasibility are lost; on the contrary, if the disturbance set is assumed to

be too large, the controller is very conservative, which can lead to a loss in performance.

For this type of applications, *Stochastic MPC* (SMPC) employing so-called *chance constraints* was proposed [7]–[13]. This method enables the formulation of constraints that have to be fulfilled with a given level of probability and hence directly formulates the tradeoff between performance and (full) constraint satisfaction. Chance constraints and their reformulation to deterministic and tractable optimization problems have been extensively studied in optimization theory (see e.g. [14] and the references therein) and a number of different formulations exist.

A. Chance Constrained Stochastic Model Predictive Control

Chance constraints are in general hard to deal with since they involve the computation of multivariate integrals. For the use in MPC problems, there are three principle directions of research: (a) Using an analytic reformulation or approximations to obtain a deterministic optimization problem [7], [11]; (b) Imposing an artificial bound on the uncertainty, solving the resulting robust optimization problem and using Chernoff bounds on the tails to get a probabilistic guarantee [10], [15]; (c) Using *samples* of the uncertainty, also called *scenarios* or *particles*, and using those to formulate an optimization problem that appropriately replaces the chance constrained problem [13], [16], [17], [19]–[22].

The main limitations of (a) are the restriction to special distributions, which might not hold for practical applications, as well as the restriction to convex problems. The main limitation of (b) is its conservatism, which often leads to a loss in performance in practical applications. The restriction to special distributions is overcome by the sampling-based approach of [16], [17], which is independent of the distribution and just requires the uncertainties to be i.i.d. as well as convexity in the decision variables (for fixed uncertain variables). This approach also leads to a conservative formulation unless the sampling and discarding approach of [18] is applied, which might be computationally demanding. The sampling-based approach by [21] does also not require an assumption on the distribution and even works for non-convex problems. It might however require a prohibitive number of samples.

Still, a remaining drawback of all above approaches is that they only give guarantees on the violation level of the solution of the open-loop optimization problem, not for the

*Frauke Oldewurtel and Göran Andersson are with the Power Systems Laboratory, Department of Electrical Engineering, Swiss Federal Institute of Technology in Zurich (ETHZ), Switzerland. {oldewurtel|andersson}@eeh.ee.ethz.ch

†David Sturzenegger, Peyman Mohajerin Esfahani, Manfred Morari, and John Lygeros are with the Automatic Control Laboratory, Department of Electrical Engineering, Swiss Federal Institute of Technology in Zurich (ETHZ), Switzerland.

closed-loop receding horizon problem. This failure to exploit the allowed violation level in closed-loop might lead to a loss in performance in many practical applications.

The probability of constraint violation can be interpreted as a frequency, i.e., “The system is allowed to violate X times out of Y trials”. From this perspective it makes sense to measure the violations during operation of the MPC controller and adapt the formulation along the path so as to meet the desired constraint violation in closed-loop. Hence we propose in this work an *Adaptive Stochastic MPC* formulation.

B. Main Idea and Outline

The work is organized as follows. Section II provides some standard definitions and facts of stochastic systems used in this paper for improving readability. In Section III the problem formulation is stated. Section IV presents the Adaptive Stochastic MPC formulation. The convergence results are presented in Section V as well as comments on the rate of convergence and a modification for active constraints. An illustrative example is given in Section VI. Section VII provides a discussion and conclusions.

C. Notation

The real number set is denoted by \mathbb{R} , the non-negative real numbers including zero is denoted by \mathbb{R}_+ , the set of non-negative integers by \mathbb{N} ($\mathbb{N}_+ := \mathbb{N} \setminus \{0\}$), the set of consecutive non-negative integers $\{j, \dots, k\}$ by \mathbb{N}_j^k . Denote by $\mathbf{0}$ the zero matrix with dimension deemed obvious by context. For matrices A and B of equal dimension inequalities $A \{<, \leq, >, \geq\} B$ hold component-wise. The expectation of a random variable with respect to the probability measure \mathbb{P} is denoted by $\mathbb{E}[\cdot]$. In order to distinguish between the *prediction* of a state and the *actual* state, denote with $x_{t+k|t} \in \mathbb{R}^{n_x}$ the prediction for the actual state $x_{t+k} \in \mathbb{R}^{n_x}$ at time t , where $k \in \mathbb{N}_1^N$ and the prediction horizon is N . Furthermore, vectors (and matrices) that contain all predictions of a variable along the prediction horizon are denoted with bold letters, i.e. $\mathbf{x}_{t+1} := \{x_{t+1|t}, x_{t+2|t}, \dots, x_{t+N|t}\}$ denotes an ordered collection of the vectors $x_{t+k|t}$, which can also be considered to be a large stacked-up vector $\mathbf{x}_{t+1} \in \mathbb{R}^{Nn_x}$ when used in algebraic equations. Also note that at time t variables with a time index larger than t are (unknown) random variables (since information is only revealed up to time t and the future is uncertain). For $a, b \in \mathbb{R}$ let $a \wedge b := \min\{a, b\}$.

II. PRELIMINARIES

As basic datum, we take a filtered space $(\Omega, \mathfrak{F}, \{\mathfrak{F}_t\}, \mathbb{P})$, where $(\Omega, \mathfrak{F}, \mathbb{P})$ is a probability triple with sample space Ω , σ -algebra \mathfrak{F} and probability measure \mathbb{P} on (Ω, \mathfrak{F}) . $\{\mathfrak{F}_t : t \geq 0\}$ is a filtration, i.e., an increasing family of sub σ -algebras of \mathfrak{F} :

$$\mathfrak{F}_0 \subseteq \mathfrak{F}_1 \subseteq \dots \subseteq \mathfrak{F}.$$

We define

$$\mathfrak{F}_\infty := \sigma \left(\bigcup_n \mathfrak{F}_n \right) \subseteq \mathfrak{F}.$$

In this work, *stochastic processes* in discrete time are considered, i.e. infinite sequences of random variables $Z_0, Z_1, \dots, Z_n, \dots$, where the index denotes the time instant.

Definition 1 (Adapted process [25]): A process $Z = (Z_t : t \geq 0)$ is called adapted to the filtration $\{\mathfrak{F}_t\}$ if for each t , Z_t is \mathfrak{F}_t -measurable.

Definition 2 (Almost surely): A statement S is said to be true almost surely (a.s.) or with probability 1 if $\mathbb{P}[S \text{ is true}] = 1$.

Definition 3 (Martingale, super/sub-martingale [25]): A process Z is called a martingale (relative to $(\{\mathfrak{F}_t\}, \mathbb{P})$) if

- (a) Z is an adapted process,
- (b) $\mathbb{E}[|Z_t|] < \infty, \forall t$,
- (c) $\mathbb{E}[Z_{t+1} | \mathfrak{F}_t] = Z_t, \text{ a.s. } (t \geq 0)$.

A super-(sub-)martingale is defined with (c) replaced by

$$\mathbb{E}[Z_{t+1} | \mathfrak{F}_t] \leq (\geq) Z_t, \text{ a.s. } (t \geq 0).$$

Lemma 1 (Measurable function [25]): Sums and products of measurable functions are measurable.

Proposition 1 (Properties of conditional expectation [25]): All X satisfy $\mathbb{E}[|X|] < \infty$. \mathfrak{G} and \mathfrak{H} denote sub- σ -algebras of \mathfrak{F} .

- (a) (**Linearity**) $\mathbb{E}[a_1 X_1 + a_2 X_2 | \mathfrak{G}] = a_1 \mathbb{E}[X_1 | \mathfrak{G}] + a_2 \mathbb{E}[X_2 | \mathfrak{G}], \text{ a.s.}$
- (b) If X is \mathfrak{G} measurable, then $\mathbb{E}[X | \mathfrak{G}] = X, \text{ a.s.}$

III. PROBLEM FORMULATION

Consider a discrete linear time-invariant system with time step t

$$x_{t+1} = Ax_t + Bu_t + w_{t+1}, \quad t \in \mathbb{N}, \quad (1)$$

where $x_t \in \mathbb{R}^{n_x}$, $u_t \in \mathbb{R}^{n_u}$, and $w_t \in \mathbb{R}^{n_w}$.

Remark 1: Note the deviation from usual notation in the index of w . The reason is the assumption in MPC schemes that the uncertainty is revealed only after the input at time t is applied resulting in a new state at $t+1$. Hence, at time t , w is unknown. It will only reveal itself at $t+1$. We therefore use the notation in (1) to comply with the notation used for stochastic systems and to highlight that at time t , x_{t+1} is a random variable which linearly depends on the random variable w_{t+1} , whereas both x_t and u_t are known.

We seek to solve an MPC problem with planning horizon N .

Assumption 1 (System): (A, B) is stabilizable and at each sample instant a measurement of the state is available.

Definition 4: Let $w_t \in \mathbb{W} \subset \mathbb{R}^{n_w}$, and $\mathbb{W}^{\mathbb{N}} = \Omega$. Furthermore, $\mathfrak{W} = \sigma(\{w : w_t \in \mathbb{W}\} : t \in \mathbb{N})$ and $\mathfrak{F}_t = \sigma(\{w_s : w_s \in \mathbb{W}\} : s \leq t)$.

Assumption 2 (Disturbances): (a) The disturbances w are independent and identically distributed (i.i.d.).

(b) The disturbances w are unbiased.

The system is subject to hard convex constraints on the inputs and probabilistic constraints on the states, which are convex in the decision variables. Note that the convexity is required for the efficient solution of the MPC problem. In fact any type of constraint that can be handled by the applied SMPC algorithm can also be handled by the adaptive SMPC algorithm. Polytopic constraints are considered in the following for simplicity of presentation.

We hence formulate the following constraints on the predicted inputs and predicted states for setting up the MPC problem:

$$S\mathbf{u}_{t+k|t} \leq \mathbf{s} \quad \forall k \in \mathbb{N}_0^{N-1}, \quad (2)$$

where $S \in \mathbb{R}^{q \times n_u}$ and $\mathbf{s} \in \mathbb{R}^r$ and

$$\mathbb{P} [G\mathbf{x}_{t+k|t} \leq \mathbf{g}] \geq 1 - \alpha \quad \forall k \in \mathbb{N}_1^N, \quad (3)$$

where $G \in \mathbb{R}^{r \times n_x}$ and $\mathbf{g} \in \mathbb{R}^r$.

Defining

$$\begin{aligned} \mathbf{x}_{t+1} &:= \{x_{t+1|t}, x_{t+2|t}, \dots, x_{t+N|t}\} \in \mathbb{R}^{Nn_x} \\ \mathbf{u}_t &:= \{u_{t|t}, u_{t+1|t}, \dots, u_{t+N-1|t}\} \in \mathbb{R}^{Nn_u} \\ \mathbf{w}_{t+1} &:= \{w_{t+1|t}, w_{t+2|t}, \dots, w_{t+N|t}\} \in \mathbb{R}^{Nn_w} \end{aligned}$$

the dynamics along the prediction horizon can be written more compactly as

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t + \mathbf{E}\mathbf{w}_{t+1}, \quad (4)$$

with \mathbf{x}_{t+1} , \mathbf{u}_t , and \mathbf{w}_{t+1} denoting the states, inputs, and disturbances, along the prediction horizon, respectively, and appropriate matrices \mathbf{A} , \mathbf{B} , and \mathbf{E} .

The aim is to minimize a quadratic cost function $J : \mathbb{R}^{n_x} \times \mathbb{R}^{Nn_u} \times \mathbb{R}^{Nn_w} \rightarrow \mathbb{R}_{+0}$

$$J(x_t, \mathbf{u}_t, \mathbf{w}_{t+1}) := \mathbf{u}_t^T \mathbf{R} \mathbf{u}_t + \mathbf{x}_{t+1}^T \mathbf{Q} \mathbf{x}_{t+1}, \quad (5)$$

where $\mathbf{R} = \mathbf{R}^T \succ 0$ and $\mathbf{Q} = \mathbf{Q}^T \succeq 0$.

IV. ADAPTIVE STOCHASTIC MPC

In this section, the adaptive SMPC procedure is described. The predicted states can be written explicitly as

$$x_{t+k|t} = A^k x_t + \sum_{i=0}^{k-1} A^{k-1-i} B u_{t+i|t} + \sum_{i=0}^{k-1} A^{k-1-i} w_{t+1+i|t}.$$

The aim is to tighten the *nominal constraint* (i.e. the constraint resulting from setting $w_{t+1+k|t} = 0$ for all $k \in \mathbb{N}_1^N$) by some amount $\tilde{h}_{t,k} \in \mathbb{R}^r$ that needs to be determined and to include the following constraint in the MPC problem

$$GA^k x_t + \sum_{i=0}^{k-1} GA^{k-1-i} B u_{t+i|t} \leq \mathbf{g} - \tilde{h}_{t,k} \quad \forall k \in \mathbb{N}_1^N.$$

Loosely speaking, a choice for $\tilde{h}_{t,k}$ should be made such that the constraint is robust against most uncertainties w_{t+1} ,

but so that the chance constraint in (3) is met in closed-loop without unnecessary conservativeness.

In most existing approaches, $\tilde{h}_{t,k} \in \mathbb{R}^r$ is determined such that the chance constraint is fulfilled for the predicted states $x_{t+k|t}$ (i.e. for the optimization problem) leading to a constraint violation which is much smaller than required in closed-loop.

The proposed adaptive SMPC scheme starts by determining an initial (conservative) constraint tightening, which is then successively adapted based on the experienced violations. There are many possibilities for determining an initial constraint tightening. Here the so-called *scenario approach* is used. Let $\xi^{(1)}, \dots, \xi^{(\lambda)}$ denote an available collection of λ i.i.d. random samples from the disturbances along the horizon N , i.e., $\xi^{(i)} = \{w_{t+1}^{(i)}, w_{t+2}^{(i)}, \dots, w_{t+N}^{(i)}\} \in \mathbb{W}^N$ for $i = 1, \dots, \lambda$. We then employ the procedure in [22] and compute

$$\tilde{h}_{t,k} := \max_{j \in \{1, \dots, \lambda\}} \sum_{i=0}^{k-1} GA^{(k-1-i)} w_{t+1+i}^{(j)} \quad \forall k \in \mathbb{N}_1^N, \quad (6)$$

where the maximization is applied to each element individually. Following [23], one should generate

$$\lambda \geq \frac{1}{\alpha} \left(d + \ln \frac{1}{\epsilon} + \sqrt{2d \ln \frac{1}{\epsilon}} \right) \quad (7)$$

samples, where d is the number of decision variables of the respective optimization problem (in this case $d = Nn_u$) and $\epsilon \in (0, 1)$ is a confidence level. The solution of the corresponding optimization problem will then guarantee the chance constraint in (3) with confidence of at least $(1 - \epsilon)$.

Remark 2: Due to Assumption 2 it is likely that all entries of $\tilde{h}_{t,k}$ are positive (which corresponds to a tightening of the nominal constraint). If this is not the case, a new set of samples is drawn.

In [19], [22] this computation is done once off-line and then used throughout the receding horizon implementation of the optimization problem solution. By contrast, in the method proposed here the off-line computed bound is only used as a starting point for the constraint tightening at the initial step. Then on-line the constraint tightening is updated based on constraint violations experienced along the path so that the empirical constraint violation probability along the trajectory converges to the desired value α .

Let $\tilde{\mathbf{h}}_t := \{\mathbf{0}, \tilde{h}_{t,1}, \tilde{h}_{t,2}, \dots, \tilde{h}_{t,N}\}$. Then the above specifications are assembled to the following finite horizon optimal control problem with planning horizon N and the time-dependent constraint tightening $\tilde{\mathbf{h}}_t$.

Problem 1 (ASMPc Problem): Determine

$$\mathbf{u}_t(x_t) = \arg \min_{u_{t|t} \dots u_{t+N-1|t}} J(x_t, \mathbf{u}_t, \mathbf{0})$$

subject to

$$\mathbf{S}\mathbf{u}_t \leq \mathbf{s}$$

$$\mathbf{G}(\mathbf{A}x_t + \mathbf{B}\mathbf{u}_t) \leq \mathbf{g} - \tilde{\mathbf{h}}_t,$$

where $\mathbf{S} \in \mathbb{R}^{Nq \times Nn_u}$ and $\mathbf{s} \in \mathbb{R}^{Nq}$ are describing the input constraints and $\mathbf{G} \in \mathbb{R}^{Nr \times Nn_x}$ and $\mathbf{g} \in \mathbb{R}^{Nr}$ are describing the state constraints along the prediction horizon.

Remark 3: Note that we choose to optimize over the nominal cost $J(x_t, \mathbf{u}_t, \mathbf{0})$, which is the cost in (5) with $\mathbf{w} = \mathbf{0}$.

In order to determine the iterative update of $\tilde{h}_{t,k}$ the violation at time step t is determined as

$$V_t(w_t) := \begin{cases} 1 & \text{if } G(Ax_{t-1} + Bu_{t-1|t-1}^* + w_t) > g \\ 0 & \text{if } G(Ax_{t-1} + Bu_{t-1|t-1}^* + w_t) \leq g. \end{cases}$$

Then the empirical probability of violation is given as

$$Y_t(w) := \frac{\sum_{i=1}^t V_i(w_i)}{t}. \quad (8)$$

Remark 4: V is an \mathfrak{F}_t -adapted process, i.e. at time t the random variable V_t is realized and hence known. Being a function of V , Y is also an \mathfrak{F}_t -adapted process.

The empirical probability of violation Y is desired to converge to the pre-defined violation level α . In order to achieve this, we aim to iteratively update $\tilde{h}_{t,k}$. An intuitive way is given by the following updating rule

$$\tilde{h}_{t,k} := \tilde{h}_{t-1,k} - \frac{\alpha - Y_t}{\gamma} \tilde{h}_{t-1,k}, \quad (9)$$

where γ is a design parameter that influences the rate of convergence (see Section V-B). The updating rule in (9) simply takes the $\tilde{h}_{t-1,k}$ determined at the previous time step and subtracts $\tilde{h}_{t-1,k}$ itself scaled by a factor which is proportional to the deviation of Y_t from α . Observe that if α is larger (smaller) than Y_t , $\tilde{h}_{t,k}$ is shrinking (growing), i.e., the constraints in Problem (1) are relaxed (tightened) and hence the probability of violations is increased (reduced).

A block diagram showing the adaptive SMPC procedure is given in Figure 1 and the adaptive MPC algorithm is summarized as follows.

Algorithm 1: Adaptive Stochastic MPC

Initialization

- 1: $t = 0$
 - 2: Draw λ samples according to Assumption 2 and (7)
 - 3: Compute $\tilde{h}_{0,k}$ for all $k \in \mathbb{N}_0^k$ according to (6)
-

On-line

- 1: Solve ASMP Problem
 - 2: Apply u_t^*
 - 3: $t = t + 1$
 - 4: Measure x_t and observe realizations of V_t and Y_t
 - 5: Compute new constraint tightening $\tilde{h}_{t,k} := \tilde{h}_{t-1,k} - \frac{\alpha - Y_t}{\gamma} \tilde{h}_{t-1,k}$
 - 6: Repeat
-

V. CONVERGENCE RESULTS

A. Supermartingale Property

Defining

$$Z_t := |\alpha - Y_t|, \quad (10)$$

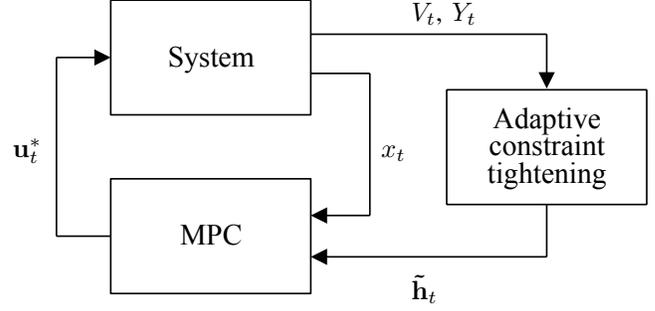


Fig. 1. Block diagram of adaptive SMPC algorithm.

the aim of this section is to argue that the empirical probability of violation converges to the desired probability of violation for $t \rightarrow \infty$. Comments on the rate of convergence are given in Section V-B.

Assumption 3: (a) The input constraints are such that the system can always cancel out the disturbance from the previous time step, if needed, and additionally provide enough control input to steer the system to the desired state. (b) The system has reachability index 1, i.e., it is one-step controllable.

Remark 5: Assumption 3(a) is essential for the proposed strategy. Although it seems restrictive on the first sight, for many practical applications it might actually hold, see discussion. Assumption 3(b) basically means that there are at least as many inputs as states and B is full row rank. This assumption could be dropped leading to a more complicated formulation, see discussion.

Let $\mathcal{K}(\alpha, t) := (\alpha - \frac{1}{2(t+1)}, \alpha + \frac{1}{2(t+1)})$ and denote with τ the first time Y_t hits \mathcal{K} , i.e.,

$$\tau_{\mathcal{K}}(w) := \inf\{t \in \mathbb{N} \mid Y_t(w) \in \mathcal{K}(\alpha, t)\}. \quad (11)$$

Furthermore, let $\hat{Z}_t := Z_{t \wedge \tau}$.

Theorem 1: \hat{Z} is a supermartingale.

Proof: The proof establishes the fulfillment of the three conditions given in Definition 3.

(a) Let $\mathfrak{F}_t = \sigma(\mathbb{W}_0, \mathbb{W}_1, \dots, \mathbb{W}_t)$. Since $Z_t = f_t(\mathbb{W}_0, \mathbb{W}_1, \dots, \mathbb{W}_t)$, where f_t is a measurable function on \mathbb{R}^{t+1} , with Lemma 1 it follows that Z is adapted and hence \hat{Z} is adapted.

(b) We need to show that $\mathbb{E}[\hat{Z}_{t+1}] < \infty, \forall t$. Since V_t is either 0 or 1, we have $Y_t \in [0, 1] \forall t$. Furthermore, $\alpha \in [0, 1]$ and hence $Z_t = |\alpha - Y_t| \in [0, 1] \forall t$, which establishes the assertion.

(c) We need to show that $\mathbb{E}[\hat{Z}_{t+1} \mid \mathfrak{F}_t] \leq \hat{Z}_t$, a.s. For this, it suffices to show that $\mathbb{E}[Z_{t+1} \mid \mathfrak{F}_t] \leq Z_t$ for the stochastic interval $0 \leq t \leq \tau$, a.s.

It is easy to show that Y_{t+1} can be rewritten based on the definition in (8) as

$$Y_{t+1} = Y_t \frac{t}{t+1} + \frac{V_{t+1}}{t+1}. \quad (12)$$

Since V_{t+1} is either 0 or 1, let $p_t := \mathbb{P}[V_{t+1} = 1 \mid \mathfrak{F}_t]$ and hence $(1 - p_t) = \mathbb{P}[V_{t+1} = 0 \mid \mathfrak{F}_t]$. It suffices to show that the following difference Δ_t is ≤ 0 for any $Y_t \notin \mathcal{K}(\alpha, t)$.

$$\begin{aligned} \Delta_t(Y_t, p_t) &:= \mathbb{E}[Z_{t+1} \mid \mathfrak{F}_t] - Z_t \\ &= \mathbb{E}\left[\left|\alpha - Y_t \frac{t}{t+1} - \frac{V_{t+1}}{t+1}\right| \mid \mathfrak{F}_t\right] - |\alpha - Y_t| \\ &\stackrel{\text{Prop. 1(a),(b)}}{=} p_t \left(\left|\alpha - Y_t \frac{t}{t+1} - \frac{1}{t+1}\right|\right) \\ &\quad + (1 - p_t) \left(\left|\alpha - Y_t \frac{t}{t+1}\right|\right) - |\alpha - Y_t| \\ &= p_t \left(\left|\alpha - Y_t \frac{t}{t+1} - \frac{1}{t+1}\right| - \left|\alpha - Y_t \frac{t}{t+1}\right|\right) \\ &\quad + |\alpha - Y_t \frac{t}{t+1}| - |\alpha - Y_t|. \end{aligned} \quad (13)$$

Observe that the only part one can influence is the value of p_t . Note that due to Assumption 3 at each step p_t can be chosen freely in the interval $[0, 1]$.

We proceed with two steps: first, the optimal policy p_t^* is determined based on analysis of the term in the round bracket of (13) and second, it is analyzed whether the last two terms can render $\Delta_t > 0$ when applying the optimal strategy p_t^* .

Step 1: If

$$\left|\alpha - Y_t \frac{t}{t+1} - \frac{1}{t+1}\right| \leq \left|\alpha - Y_t \frac{t}{t+1}\right|, \quad (14)$$

the optimal strategy is $p_t^*(Y_t) = 1$. Let $\beta := \alpha - Y_t \frac{t}{t+1}$. Then (14) yields

$$\left|\beta - \frac{1}{t+1}\right| \leq |\beta|. \quad (15)$$

For this to hold we need $\beta \geq \frac{1}{2(t+1)}$.

$$\begin{aligned} \Leftrightarrow \alpha - Y_t \frac{t}{t+1} &\geq \frac{1}{2(t+1)} \\ \Leftrightarrow \alpha - Y_t + \frac{2\alpha - 1}{2t} &\geq 0. \end{aligned}$$

The opposite case of (14) is symmetrical, we therefore have as optimal policy

$$p_t^*(Y_t) = \begin{cases} 1 & \text{if } \alpha - Y_t + \frac{2\alpha - 1}{2t} \geq 0 \\ 0 & \text{if } \alpha - Y_t + \frac{2\alpha - 1}{2t} < 0. \end{cases} \quad (16)$$

Step 2: By virtue of the analysis in Step 1, in the ‘critical’ region where $\Delta_t^* := \Delta_t(Y_t, p_t^*)$ may be positive it holds that

$$\begin{aligned} \left|\alpha - Y_t \frac{t}{t+1}\right| &> |\alpha - Y_t| \\ \Leftrightarrow \alpha - Y_t &> -\frac{Y_t}{2(t+1)} \\ \Leftrightarrow \alpha - Y_t \frac{t}{t+1} &> \frac{Y_t}{2(t+1)}. \end{aligned}$$

According to the optimal policy in (16), there are two cases.

Case I: $\alpha - Y_t \frac{t}{t+1} > \frac{1}{2(t+1)} > \frac{Y_t}{2(t+1)}$ with $p_t^*(Y_t) = 1$. Using this in (13) yields

$$\begin{aligned} \Delta_t^*(Y_t) &= \left|\alpha - Y_t \frac{t}{t+1} - \frac{1}{t+1}\right| - |\alpha - Y_t| \\ &= \left|\alpha - Y_t + \frac{1}{t+1}(Y_t - 1)\right| - |\alpha - Y_t|. \end{aligned}$$

The ‘critical’ region in Case I is then given by

$$\begin{aligned} \Delta_t^*(Y_t) &> 0 \\ \Leftrightarrow \frac{-Y_t}{2(t+1)} &< \alpha - Y_t < \frac{1}{2(t+1)} - \frac{Y_t}{2(t+1)}. \end{aligned}$$

Case II: This case is symmetrical to Case I. We have $\frac{1}{2(t+1)} > \alpha - Y_t \frac{t}{t+1} > \frac{Y_t}{2(t+1)}$ with $p_t^* = 0$. Carrying out the same analysis as in Case I yields the following condition for the ‘critical’ region of Case II

$$\frac{-Y_t}{2(t+1)} < \alpha - Y_t < \frac{1}{2(t+1)}.$$

Using the largest upper and smallest lower bound of both Cases I and II yields the overall ‘critical’ region

$$\mathcal{K}(\alpha, t) = \left\{y \in \mathbb{R} \mid \alpha - \frac{1}{2(t+1)} < y < \alpha + \frac{1}{2(t+1)}\right\}$$

and the desired result. \blacksquare

This leads us to the following time-dependent optimal updating rule (observe the similarity with the intuitive updating rule in (9))

$$\tilde{h}_{t,k} := \tilde{h}_{t-1,k} - \frac{\alpha - Y_t + \frac{2\alpha - 1}{2t}}{\gamma} \tilde{h}_{t-1,k}. \quad (17)$$

Due to the monotonicity $h \uparrow \Rightarrow p \downarrow$ and $h \downarrow \Rightarrow p \uparrow$, this updating rule satisfies the behavior already described for the intuitive updating rule in (9).

Corollary 1: At each time step t , there is at most one possible Y_t such that $\Delta_t^*(\alpha, Y_t) > 0$.

Proof: Observe that the width of \mathcal{K} is $1/(t+1)$, which is smaller than $1/t$. Furthermore, observe that due to the definition of Y_t , at each time t , there is a finite number of possible values of Y_t with the minimum distance of $1/t$ between any of them. Hence, at each time step t , there is at most one possible Y_t such that $\Delta_t^*(\alpha, Y_t) > 0$. \blacksquare

\hat{Z} is a supermartingale with arbitrary starting point. We can therefore view Z as a concatenation of supermartingales. Furthermore, since both \mathcal{K} and Δ_t^* are inversely proportional with respect to time, both the probability of hitting the ‘critical’ region as well as the distortion are shrinking over time. Intuitively arguing, these observations suggest that $\lim_{t \uparrow \infty} Z_t = 0$ a.s., but we do not provide a rigorous proof here.

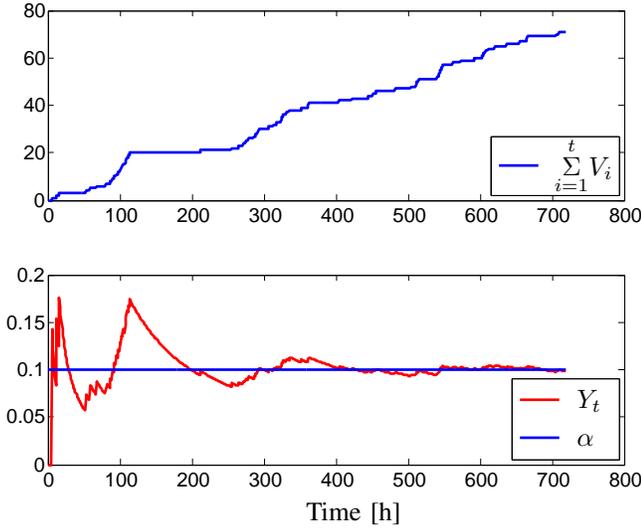


Fig. 2. Heuristic updating rule (9) and time invariant constraints. Top: Cumulative number of violations $\sum_{i=1}^t V_i$ over time. Bottom: Empirical probability of violation Y_t and desired violation level α over time.

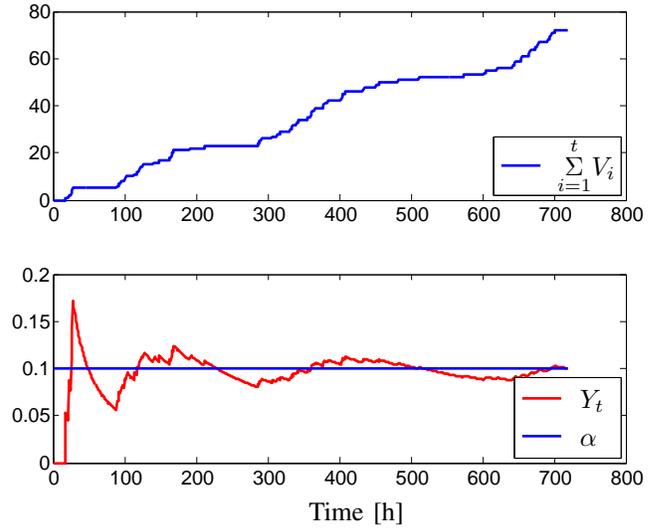


Fig. 4. Optimal updating rule (17) and time varying constraints. Top: Cumulative number of violations $\sum_{i=1}^t V_i$ over time. Bottom: Empirical probability of violation Y_t and desired violation level α over time.

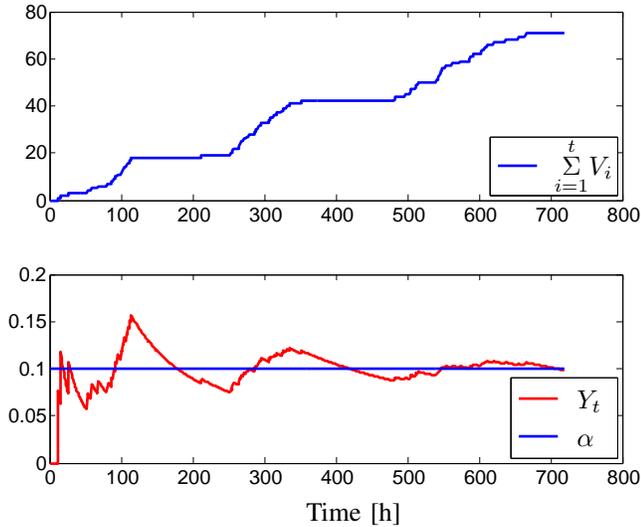


Fig. 3. Optimal updating rule (17) and time invariant constraints. Top: Cumulative number of violations $\sum_{i=1}^t V_i$ over time. Bottom: Empirical probability of violation Y_t and desired violation level α over time.

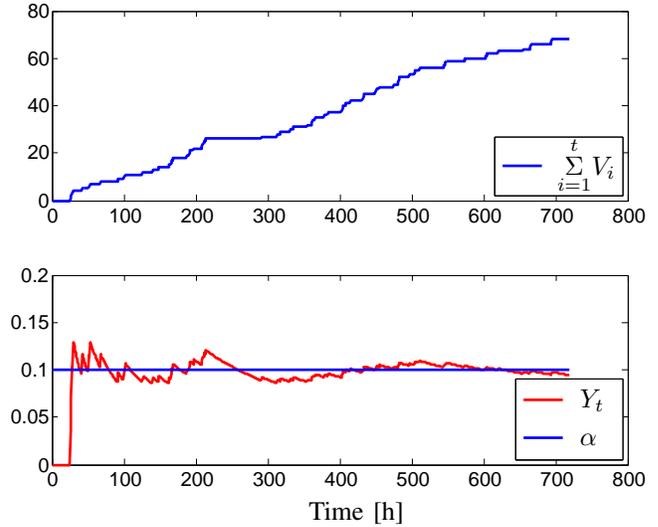


Fig. 5. Modified updating rule (19) and time varying constraints. Top: Cumulative number of violations $\sum_{i=1}^t V_i$ over time. Bottom: Empirical probability of violation Y_t and desired violation level α over time.

B. Rate of Convergence

For both the intuitive updating rule in (9) and the optimal updating rule in (17), γ determines the rate of convergence. γ should in all cases be > 0 in order to tighten if $Y_t > \alpha$ and relax if $Y_t < \alpha$. Given $\gamma > 0$, the larger γ , the slower the convergence and the smaller γ , the faster the convergence.

For the updating rule $\tilde{h}_{t,k} = \mu_t \tilde{h}_{t-1,k}$, we require $\mu_t > 0$. This is to ensure that the relaxation is not more than up to the nominal constraint. This requirement yields $\gamma > \alpha$ for the heuristic updating rule in (9) and $\gamma > 2\alpha$ for the optimal updating rule in (17). If the initial tightening is found with the scenario approach, it might furthermore make sense to bound the tightening with the updating rule by not allowing a tighter bound than given by the initial tightening, since it is known that this tightening (conservatively) satisfies the chance constraint. The initial tightening and the nominal constraint provide an upper and lower bound of the tightening.

C. Modification of Updating Rule for Active Constraints

In this section a modification of the proposed approach is introduced with two goals: (a) The first value of \tilde{h} is determined here with the scenario approach and all subsequent updates only scale this first version. However, it appears to be desirable to also be able to change the vector orientation as some dimension might turn out to be more conservative than another; (b) In many cases, it might happen that a particular constraint cannot be violated with the required probability level at a particular time instant since the state is too far away from the nominal constraint. This is, e.g., regularly the case in building climate control, where the temperature rises during the day. So only in the mornings and evenings it is critical to maintain the comfort constraint. Hence, the \tilde{h} should only be updated for the constraints which are actually prone to violations, i.e., at times when the measurement of violations actually provides some feedback of the tightening.

To achieve this, we propose to only update active constraints, i.e., we determine all i such that

$$\mathbf{G}^{[i]}(\mathbf{A}x_t + \mathbf{B}u_t^*) = \mathbf{g}^{[i]} - \tilde{\mathbf{h}}_t^{[i]}. \quad (18)$$

Then the update rule is changed to update only the rows i for which the constraint in (18) is satisfied.

$$\tilde{\mathbf{h}}_t^{[i]} := \tilde{\mathbf{h}}_{t-1}^{[i]} - \frac{\alpha - Y_t + \frac{2\alpha-1}{2t}}{\gamma} \tilde{\mathbf{h}}_{t-1}^{[i]}. \quad (19)$$

VI. EXAMPLE

The proposed strategy is tested on a small-scale example based on a simplified version of [26]. It is a building climate control problem, where the goal is to minimize heating energy consumption while satisfying a chance constraint on room temperature. The system matrices of the discrete-time system are given as

$$\begin{aligned} A &= \begin{bmatrix} 0.0016 & 0.8676 & 0.1182 \\ 0.0017 & 0.9935 & 0.0044 \\ 0.0012 & 0.0229 & 0.9725 \end{bmatrix} & B &= \begin{bmatrix} 0.0431 \\ 0.0016 \\ 0.0011 \end{bmatrix} \\ E &= \begin{bmatrix} 0.0126 & 0.0029 \\ 0.0005 & 0.0001 \\ 0.0034 & 0.0001 \end{bmatrix} \end{aligned} \quad (20)$$

and the time step is one hour. The system states $x = [x_{(1)} \ x_{(2)} \ x_{(3)}]^T$ are the room temperature $x_{(1)}$, the temperature in the inner wall $x_{(2)}$, and the temperature in the outer wall $x_{(3)}$, all given in °C. The control input is the heating $u_{(1)}$, given in W/m². The disturbances are the outside temperature $w_{(1)}$ in °C and the irradiation $w_{(2)}$ in W/m². For simulation purposes, i.i.d. disturbance samples are drawn from a component-wise uncorrelated standard normal distribution with each component truncated at 4, i.e., $\mathbb{W} = \{w \in \mathbb{R}^{n_w} \mid \|w\|_\infty \leq 4\}$.

The goal is to minimize the cost function in (5) with $\mathbf{Q} = \mathbf{0}$ while satisfying the following *time-invariant* constraints

$$\begin{aligned} 0 \leq u_{t+k|t} \leq 100 & \quad \forall k \in \mathbb{N}_0^{N-1} \\ \mathbb{P}[[1 \ 0 \ 0]x_{t+k|t} \geq 21] \geq 1 - \alpha & \quad \forall k \in \mathbb{N}_1^N. \end{aligned}$$

As an alternative to the state constraints given above, we investigate time-varying state constraints, so-called *temperature setbacks*, where the room temperature constraints are relaxed during the night. This is given as

$$\begin{cases} \mathbb{P}[[1 \ 0 \ 0]x_{t+k|t} \geq 21] \geq 1 - \alpha & \text{during daytime} \\ \mathbb{P}[[1 \ 0 \ 0]x_{t+k|t} \geq 12] \geq 1 - \alpha & \text{during nighttime} \end{cases}.$$

The simulation starts from steady state. The initial state is $x_0 = [21 \ 21 \ 19.38]^T$. The prediction horizon is $N = 5$ and the desired level of violation $\alpha = 0.1$. Simulation time is 720 hours which equals 30 days.

A. Investigations

Three investigations are carried out. All investigation are using the adaptive SMPC scheme (Algorithm 1), the building example in (20), and $\gamma = 0.7$ for all updating rules. The investigations differ in terms of the used updating rule and the definition of state constraints.

Investigation 1: In Investigation 1 the heuristic update rule (9) is compared with the optimal update rule (17). Time-invariant constraints are used.

Investigation 2: In Investigation 2 the optimal updating rule (17) is used and time invariant constraints are compared with setbacks.

Investigation 3: In Investigation 3 the optimal updating rule (17) is compared with the updating rule (19) that only updates active constraints. Setbacks are used.

B. Results

Note that the results obtained depend on the drawn samples. For comparability all simulations were carried out with the same samples. For all investigated cases the empirical probability of violation converges to the pre-specified violation level α within a couple of days, which is a reasonable time frame for the application at hand.

Investigation 1: When comparing Figures 2 and 3 one can see that the heuristic updating rule (9) produces a larger overshoot of Y_t , especially at the very beginning. With the optimal, time-dependent updating rule this overshoot can be reduced. For large t the difference between the two updating rules goes to zero.

Investigation 2: When comparing Figures 3 and 4 a larger overshoot at the beginning can be seen for the case of setbacks. This is created by the additional relaxation of the constraint during setback times when the constraint is not active and not likely to be violated. Despite the overshoot of more than 50% of the desired value of α at the beginning, the system quickly converges to the desired value of α (within 5-6 days in this example).

Investigation 3: When comparing Figures 4 and 5 we see that the modified updating rule (19) is producing much less overshoot when setbacks are used.

VII. DISCUSSION AND CONCLUSIONS

The key result of the paper is that the empirical probability of violation can be used to adapt the constraint tightening such that it converges in closed-loop to the desired level of constraint violation if the system meets the described assumptions. This is in particular interesting for practical applications, because it enables the exploitation of constraint violations for improving the performance. This paper introduces the basic idea and a convergence result, but there are further points to be addressed.

First, the update of active constraints. This idea was introduced in Section V-C. It can be expected that this becomes more important for more complex systems with many constraints. Additionally it can make sense to define chance constraints not for each time step, but for each individual constraint and then measure the violation of this and do the update according to the particular empirical violation probability. Such formulations are particularly meaningful for some applications, e.g., when the chance constraint is naturally formulated only for one state not for all states of each time step.

Second, the requirement of reachability index 1 (Assumption 3(b)). This requirement basically means that the system needs to have at least as many inputs as states and the matrix B needs to have full row rank. For some practical systems this might hold. If not, this requirement can be dropped by formulating a different convergence result as follows. Since each (controllable) system needs n steps to control the state (with $n \leq n_x$), instead of showing a supermartingale property where the expectation is decreasing at every step, one may require the expectation to decrease every n steps.

Third, enough input power available (Assumption 3(a)). This assumption is restrictive on the first sight, but it holds for many practical applications, simply because the systems are designed such that they can handle the respective disturbances, e.g., in building control systems, systems are designed such that there is enough input power available even in the presence of uncertainties. This question is also strongly linked to the question of recursive feasibility of the MPC problem. In the presented setup, if the SMPC problem with the initial tightening is feasible, then due to Assumption 3 the adaptive SMPC problem remains feasible for all t .

Fourth, the presented method is not limited to linear systems, but can be applied to nonlinear systems as well. Provided Assumption 3 holds, the convergence result also holds in this case.

VIII. ACKNOWLEDGMENTS

The European Commission is gratefully acknowledged for its financial support. The work of F. Oldewurtel is supported by the European Commission under the FP7 project ‘Umbrella’. The work of P. Mohajerin Esfahani is supported by the European Commission under the project ‘MoVeS’.

IX. APPENDIX

A. Analysis of Δ_t depending on p_t and Y_t

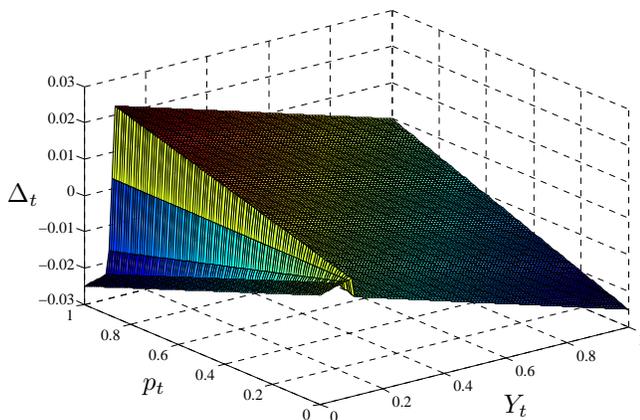


Fig. 6. Plot of Δ_t in (13) for $\alpha = 0.1$ depending on p_t and Y_t . One can see that by choosing p_t , Δ_t can be made negative for almost all values of Y_t . There is only a small ‘critical’ region around α .

REFERENCES

- [1] Bemporad, A. and Morari, M., *Robust model predictive control: A survey*, in: Robustness in identification and control, Lecture Notes in Control and Information Sciences, Garulli, A. and Tesi, A. (eds.), Springer Berlin / Heidelberg, 1999.
- [2] Rawlings, J.B., Mayne, D.Q., *Model predictive control: Theory and design*, Nob Hill Publishing, 2009.
- [3] Ben-Tal, A., Goryashko, A., Guslitzer, E., Nemirovski, A., *Adjustable robust solutions of uncertain linear programs*, Mathematical Programming, vol. 99, no. 2, pp. 351-376, 2004.
- [4] Goulard, P.J., Kerrigan, E.C., Maciejowski, J.M., *Optimization over state feedback policies for robust control with constraints*, Automatica, vol. 42, no. 4, pp. 523-533, 2006.
- [5] Löfberg, J., *Approximation of closed-loop MPC*, Proc. of the 42nd IEEE Conference on Decision and Control, 2003.
- [6] Mayne, D.Q., Seron, M.M., Rakovic, S.V., *Robust model predictive control of constrained linear systems with bounded disturbances*, Automatica, vol. 41, pp. 219-224, 2005.
- [7] Cinquemani, E., Agarwal, M., Chatterjee, D., Lygeros, J., *On convex problems in chance constrained stochastic model predictive control*, <http://arxiv.org/abs/0905.3447>, 2009.
- [8] Cannon, M., Kouvaritakis, B., Wu, X., *Probabilistic constrained MPC for systems with multiplicative and additive stochastic uncertainty*, IFAC World Congress (Seoul, Korea), 2009.
- [9] Hokayem, P., Cinquemani, E., Chatterjee, D., Lygeros, J., *Stochastic MPC with output feedback and bounded control inputs*, <http://arxiv.org/pdf/1001.3015.pdf>, 2010.
- [10] Oldewurtel, F., Jones, C.N., Morari, M., *A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback*, Proc. of the 47th IEEE Conference on Decision and Control, pp. 4731-4736, 2008.
- [11] Van Hessem D.H., Bosgra, O.H., *A full solution to the constrained stochastic closed-Loop MPC problem via state and innovations feedback and its receding horizon implementation*, Proc. of the 42nd IEEE Conference on Decision and Control, pp. 929-934, 2003.
- [12] Primbs, J., Sung, C.H., *Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise*, IEEE Trans. on Automatic Control, vol. 54, no. 2, pp. 221-230, 2009.
- [13] Batina, I., *Model predictive control for stochastic systems by randomized algorithms*, Diss., Technische Universiteit Eindhoven, 2004.
- [14] Kall, P., Mayer, J., *Stochastic linear programming - Models, theory, and computation*, Springer, 2005.
- [15] Bertsimas, D., Sim, M., *Tractable approximations to robust conic optimization problems*, Math. Program., Ser. B 107(1), pp. 5-36, 2006.
- [16] Campi, M.C., Garatti, S., *The exact feasibility of randomized solutions of uncertain convex programs*, SIAM Journal on Optimization, vol. 19, no.3, pp. 1211-1230, 2008.
- [17] Calafiore, G., *Random convex programs*, Siam Journal on Optimization, vol. 20, no. 6, pp. 3427-3464, 2010.
- [18] Campi, M.C., Garatti, S., *A sampling-and-discarding approach to chance-constrained optimization: feasibility and optimality*, Journal of Optimization Theory and Applications 148(2):257-280, 2011.
- [19] Calafiore, G., Campi, M., *The scenario approach to robust control design*, IEEE Trans. on Automatic Control, vol. 51, no. 5, pp. 742-753, 2006.
- [20] Schildbach, G., Fagiano, L., Morari, M., *Randomized solutions to convex programs with multiple chance constraints*, <http://arxiv.org/find/all/1/all:+schildbach/0/1/0/all/0/1>, 2012.
- [21] Blackmore, L., Ono, M., Bektassov, A., William, B.C., *A probabilistic particle-control approximation of chance-constrained stochastic predictive control*, IEEE Trans. on Robotics, vol. 26, no. 3, pp.502-507, 2010.
- [22] Schildbach, G., Calafiore, G., Fagiano, L., Morari, M., *Randomized model predictive control for stochastic linear systems*, Proc. of the American Control Conference, 2011.
- [23] Alamo, T., Tempo, R., Luque A., *On the sample complexity of randomized approaches to the analysis and design under uncertainty*, Proc. of the American Control Conference, 2010.
- [24] Prandini M., Garatti S., Lygeros J., *A randomized approach to stochastic model predictive control* Proc. of the 51st IEEE Conference on Decision and Control, Maui, Hawaii, 2012.
- [25] Williams, D., *Probability with martingales*, Cambridge University Press, 1991.
- [26] Gwerder, M., Tödli, J., *Predictive control for integrated room automation*, Proc. of REHVA Clima conference, 2005.