

On convexity of stochastic optimization problems with constraints

Mayank Agarwal, Eugenio Cinquemani, Debasish Chatterjee and John Lygeros

Abstract— We investigate constrained optimal control problems for linear stochastic dynamical systems evolving in discrete time. We consider minimization of an expected value cost over a finite horizon. Hard constraints are introduced first, and then reformulated in terms of probabilistic constraints. It is shown that, for a suitable parametrization of the control policy, a wide class of the resulting optimization problems are either convex or amenable to convex relaxations.

I. INTRODUCTION

This work stems from the attempt to address optimal infinite-horizon constrained control of discrete-time stochastic processes by a model predictive control strategy [9], [19], [18], [1], [2], [4], [16]. We focus on linear dynamical systems driven by stochastic noise and a control input, and consider the problem of finding a control policy that minimizes an expected cost function while simultaneously fulfilling constraints on the control input and on the state evolution. In general, no control policy exists that guarantees satisfaction of deterministic (hard) constraints over the whole infinite horizon. One way to cope with this issue is to relax the constraints in terms of probabilistic (soft) constraints [16], [17]. This amounts to requiring that constraints will not be violated with sufficiently large probability or, alternatively, that an expected reward for the fulfillment of the constraints is kept sufficiently large.

Two considerations lead to the reformulation of an infinite horizon problem in terms of subproblems of finite horizon length. First, given any bounded set (e.g. a safe set), the state of a linear stochastic dynamical system is guaranteed to exit the set at some time in the future with probability one whatever the control policy. Therefore, soft constraints may turn the original (infeasible) hard-constrained optimization problem into a feasible problem only if the horizon length is finite. Second, even if the constraints are reformulated so that an admissible infinite-horizon policy exists, the computation of such a policy is generally intractable. The aim of this note is to show that, for certain parameterizations of the policy space [3], [11] and the constraints, the resulting finite horizon optimization problem is tractable.

An approach to infinite horizon constrained control problems that has proved successful in many applications is model predictive control [14]. In model predictive control, at every time t , a finite-horizon approximation of the infinite-horizon problem is solved but only the first control of the

resulting policy is implemented. At the next time $t + 1$, a new problem is considered, the control policy is updated, and the process is repeated in a receding horizon fashion. Under time-invariance assumptions, the finite-horizon optimal control problem is the same at all times, giving rise to a stationary optimal control policy that can be computed offline.

Motivated by the previous considerations, here we study the convexity of certain stochastic finite-horizon control problems with soft constraints. Convexity is central for the fast computation of the solution by way of numerical procedures [5], [7]. One may argue that non-convex problems can be tackled by convex approximations (see e.g. [15]) or by Monte Carlo procedures. However, for many of the classes of problems considered here, tight convex approximations are usually difficult to derive. Moreover, randomized solutions are typically time-consuming and can only provide probabilistic guarantees. In particular, this is critical in the case where the system dynamics or the problem constraints are time-varying, since in that case optimization must be performed in real-time.

In Section II we state the optimization problem of our concern. Hard constraints are considered at this stage. A suitable parametrization of the control policies is introduced in Section III motivated by [3], [11]. Two probabilistic reformulations of the original constraints and conditions for the convexity of the space of admissible control policies are discussed in Section IV; Section V is dedicated to integrated chance constraints. We present a simple example in Section VI illustrating our results.

II. PROBLEM SETTING

Let $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. Consider the following dynamical model:

$$x(t+1) = Ax(t) + Bu(t) + w(t), \quad t \in \mathbb{N}_0, \quad (1)$$

where $x(t) \in \mathbb{R}^n$ is the state, $u(t) \in \mathbb{R}^m$ is the control input, $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $w(t) \in \mathbb{R}^n$ is a (zero mean, Gaussian) white noise input with covariance $\Sigma = \mathbb{E}[w(t)w(t)^T] > 0$. We assume that $x(0) = x_0$ is given and that, at any time t , $x(t)$ is observed exactly.

Fix a horizon $N \in \mathbb{N}$. The evolution of the system from $t = 0$ through $t = N$ can be described in compact form as follows:

$$\bar{x} = \bar{A}x_0 + \bar{B}\bar{u} + \bar{D}\bar{w}, \quad (2)$$

Mayank Agarwal is with the Department of Electrical Engineering, Stanford University, California, USA. mayankag@stanford.edu

E. Cinquemani, D. Chatterjee and J. Lygeros are with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland. cinquemani, chatterjee, lygeros@control.ee.ethz.ch

where

$$\bar{x} := \begin{bmatrix} x(0) \\ x(1) \\ \vdots \\ x(N) \end{bmatrix}, \quad \bar{u} := \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(N-1) \end{bmatrix}, \quad \bar{w} := \begin{bmatrix} w(0) \\ w(1) \\ \vdots \\ w(N-1) \end{bmatrix},$$

$$\bar{A} := \begin{bmatrix} I_n \\ A \\ \vdots \\ A^N \end{bmatrix}, \quad \bar{B} := \begin{bmatrix} 0_{n \times m} & \cdots & \cdots & 0_{n \times m} \\ B & \ddots & & \vdots \\ AB & B & \ddots & \vdots \\ \vdots & & \ddots & 0_{n \times m} \\ A^{N-1}B & \cdots & AB & B \end{bmatrix},$$

$$\bar{D} := \begin{bmatrix} 0_n & \cdots & \cdots & 0_n \\ I_n & \ddots & & \vdots \\ A & I_n & \ddots & \vdots \\ \vdots & & \ddots & 0_n \\ A^{N-1} & \cdots & A & I_n \end{bmatrix}.$$

Note that \bar{w} is a zero-mean Nn -dimensional Gaussian random vector with covariance matrix $\bar{\Sigma} := \text{diag}(\Sigma, \dots, \Sigma)$. We define the cost function

$$V(\bar{x}, \bar{u}) = [\bar{x}^T \quad \bar{u}^T] M \begin{bmatrix} \bar{x} \\ \bar{u} \end{bmatrix}, \quad (3)$$

where $M \in \mathbb{R}^{((N+1)n+Nm) \times ((N+1)n+Nm)} \geq 0$. Let $\eta : \mathbb{R}^{(N+1)n+Nm} \rightarrow \mathbb{R}^r$, $r \in \mathbb{N}$, be a measurable function. We are interested in constrained optimization problems of the following kind:

$$\inf_{\bar{u} \in \mathcal{U}} \mathbb{E}[V(\bar{x}, \bar{u})] \quad (4)$$

subject to (2) and $\eta(\bar{x}, \bar{u}) \leq 0$

where the expectation $\mathbb{E}[\cdot]$ is defined in terms of the underlying probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, and \mathcal{U} is a class of causal, deterministic control policies. The inequality in (4) is interpreted componentwise.

In general, there exists no policy that ensures that the constraints $\eta(\bar{x}, \bar{u}) \leq 0$ are satisfied for all outcomes of the stochastic input \bar{w} . A natural way to deal with this is to reformulate the hard constraints in terms of probabilistic (soft) constraints. In Sections IV and V we shall discuss some probabilistic interpretations of the constraints that lead to well-posed convex optimization problems.

We first note that standard types of constraints on the state and the input are handled easily within this framework. For instance, sequential ellipsoidal constraints of the type

$$\begin{bmatrix} x^T(t) & u^T(t) \end{bmatrix} S(t) \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \leq 1, \quad t = 0, 1, \dots, N-1,$$

$$x^T(N)S(N)x(N) \leq 1,$$

with $0 < S(t) \in \mathbb{R}^{(n+m) \times (n+m)}$ for $t = 0, 1, \dots, N-1$ and $0 < S(N) \in \mathbb{R}^{n \times n}$, are captured by the definition

$$\eta(\bar{x}, \bar{u}) = [\eta_0(\bar{x}, \bar{u}) \quad \eta_1(\bar{x}, \bar{u}) \quad \cdots \quad \eta_N(\bar{x})]^T$$

where, for $t = 0, 1, \dots, N$,

$$\eta_t(\bar{x}, \bar{u}) = [\bar{x}^T \quad \bar{u}^T] \Xi_t \begin{bmatrix} \bar{x} \\ \bar{u} \end{bmatrix} - 1,$$

and each matrix Ξ_t is immediately constructed in terms of $S(t)$. Note however that in addition our formulation also allows for cross-constraints between states and inputs at different times. Polytopic constraints are discussed in Section IV. We mention that $x(0)$ in the definition of \bar{x} ensures that constraints and costs depending on the future inputs and $x(0)$ are expressible compactly in terms of \bar{x} .

III. FEEDBACK FROM THE NOISE INPUT

By the hypothesis that the state is observed without error, one may reconstruct the noise sequence from the sequence of observed states and inputs by the formula

$$w(t) = x(t+1) - Ax(t) - Bu(t), \quad t \in \mathbb{N}_0. \quad (5)$$

In the light of this, and following [3], [11], we consider policies of the form:

$$u(t) = \sum_{i=0}^{t-1} G_{t,i} w(i) + d_t, \quad (6)$$

where the feedback gains $G_{t,i} \in \mathbb{R}^{m \times n}$ and the affine terms $d_t \in \mathbb{R}^m$ must be chosen based on the control objective. By definition the policy is causal since the value of u at time t depends on the values of w up to time $t-1$. Using (5) we see that $u(t)$ is a function of the observed states up to time t . It was shown in [11] that there exists a one-to-one (nonlinear) mapping between control policies in the form (6) and the class of affine state feedback policies. That is, provided one is interested in affine state feedback policies, parametrization (5) constitutes no loss of generality. Of course, this choice is generally suboptimal, but it will ensure the tractability of a large class of optimal control problems. In compact notation, the control sequence up to time $N-1$ is given by

$$\bar{u} = \bar{G}\bar{w} + \bar{d}, \quad (7)$$

where

$$\bar{G} := \begin{bmatrix} 0_{m \times n} & & & & \\ G_{1,0} & 0_{m \times n} & & & \\ \vdots & \ddots & & \ddots & \\ G_{N-1,0} & \cdots & G_{N-1,N-2} & 0_{m \times n} & \end{bmatrix}$$

and $\bar{d} := [d_0^T, d_1^T, \dots, d_{N-1}^T]^T$. The resulting closed-loop system dynamics can be written compactly as the equality constraint

$$\bar{x} = \bar{A}x_0 + \bar{B}(\bar{G}\bar{w} + \bar{d}) + \bar{D}\bar{w}. \quad (8)$$

In the sequel we shall use the notation $(\bar{x}_\theta, \bar{u}_\theta)$ to make the dependence of (\bar{x}, \bar{u}) on the optimization parameters $\theta := (\bar{G}, \bar{d})$ explicit. It is not difficult to verify that with the control policy (6) the cost function in (4) is convex in the optimization parameters θ . Therefore we shall concentrate on convexity of the constraints from now on.

IV. CHANCE CONSTRAINTS

A. Polytopic Constraint Functions

Polytopic constraint functions

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = T^x \bar{x}_\theta + T^u \bar{u}_\theta - y, \quad (9)$$

where $T^x \in \mathbb{R}^{r \times (N+1)n}$, $T^u \in \mathbb{R}^{r \times Nm}$, and $y \in \mathbb{R}^r$, describe one of the most common types of constraints. It is clear from (7) and (8) that $\eta(\bar{x}_\theta, \bar{u}_\theta)$ is affine in θ . For a given $\alpha \in]0, 1[$, we relax the hard constraint $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ by

$$\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha; \quad (10)$$

the smaller α , the better the approximation at the expense of a more constrained optimization problem. Note that

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = T^x \bar{x} + T^u \bar{u} - y = h_\theta + P_\theta \bar{w}, \quad (11)$$

where $h_\theta = (T^x \bar{A}x_0 - y) + (T^x \bar{B} + T^u) \bar{d}$, and $P_\theta = T^x \bar{D} + (T^x \bar{B} + T^u) \bar{G}$ are affine functions of θ , since $\bar{x}_\theta = \bar{A}x_0 + \bar{B}\bar{u}_\theta + \bar{D}\bar{w}$ and $\bar{u}_\theta = \bar{C}\bar{w} + \bar{d}$.

We now describe two approaches to approximate problems involving such chance constraints by problems that are convex in the optimization parameters.

1) *Approximation via constraint separation:* The constraint (10) requires us to satisfy $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ with probability of at least $1 - \alpha$. Here $\eta \in \mathbb{R}^r$ and the inequality $\eta \leq 0$ is to be interpreted componentwise. The underlying idea of the current method is to satisfy the inequality $\mathbb{P}(\eta_i \leq 0) \geq 1 - \alpha_i$ for each of the r components of η , where $\sum_{i=1}^r \alpha_i = \alpha$. This choice guarantees that $\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha$ is satisfied.

Proposition 1: Let T_i^x and T_i^u denote the i -th rows of the matrices T^x and T^u , respectively, and $\alpha_i \in]0, 1[$. We define $\eta_i(\bar{x}_\theta, \bar{u}_\theta) = T_i^x \bar{x} + T_i^u \bar{u} - y_i$. Then the constraint

$$\mathbb{P}(\eta_i(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha_i$$

is a second-order cone constraint.

Proof: From the definition of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ in (11) we can write $\eta_i = h_{i,\theta} + P_{i,\theta}^T \bar{w}$, where $h_{i,\theta} = (T_i^x \bar{A}x_0 - y_i) + (T_i^x \bar{B} + T_i^u) \bar{d}$, and $P_{i,\theta} = (T_i^x \bar{D} + (T_i^x \bar{B} + T_i^u) \bar{G})^T$. Note that the variables $h_{i,\theta}, P_{i,\theta}$ are affine in the original parameters $\theta = (G, \bar{d})$. Since the vector \bar{w} is Gaussian with distribution $\mathcal{N}(0, \bar{\Sigma})$, we conclude that η_i is Gaussian with distribution $\mathcal{N}(h_{i,\theta}, P_{i,\theta}^T \bar{\Sigma} P_{i,\theta})$. It follows from definitions that for a real-valued Gaussian random variable X with distribution $\mathcal{N}(\mu, \sigma^2)$, $\mathbb{P}(X \leq 0) \geq 1 - \alpha \iff \mu + k\sigma \leq 0$, where $k = \sqrt{2} \operatorname{erf}^{-1}(1 - 2\alpha)$. Here $\operatorname{erf}^{-1}(\cdot)$ denotes the inverse function of the standard error function $\operatorname{erf}(x) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$. In view of this fact, the constraint $\mathbb{P}(\eta_i(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha_i$ is equivalent to $h_{i,\theta} + k_i \left\| \Sigma^{\frac{1}{2}} P_{i,\theta} \right\| \leq 0$, where k_i is a fixed number depending upon the choice of α_i . The above constraint is a second-order cone constraint in θ , which proves the claim. ■

This approach can be too conservative since the probability of a union of events is approximated by the sum of the probabilities of the individual events. The second approach described below calculates a conservative approximation of the union at once.

2) *Approximation via confidence ellipsoids:* The constraint function is given by (11). Thus, $\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) = \mathbb{P}(P_\theta \bar{w} + h_\theta \leq 0)$, and the constraint $\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha$ is equivalent to $\mathbb{P}(P_\theta \bar{w} + h_\theta \leq 0) \geq 1 - \alpha$.

Since \bar{w} is Gaussian with distribution $\mathcal{N}(0, \bar{\Sigma})$, the random variable $\eta(\bar{x}_\theta, \bar{u}_\theta) = P_\theta \bar{w} + h_\theta$ is also Gaussian with distribution $\mathcal{N}(h_\theta, \bar{\Sigma}_\theta)$, with $\bar{\Sigma}_\theta := P_\theta \bar{\Sigma} P_\theta^T$. Thus the constraint (10) with η as above restricts the choice of P_θ and h_θ to be such that, with a probability of $1 - \alpha$ or more, the realization of random vector $\eta(\bar{x}_\theta, \bar{u}_\theta)$ should lie in the negative orthant i.e. $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$. In general, it is difficult to describe this constraint explicitly since it involves the integration of the pdf of the normal distribution over the negative orthant, for which no known closed-form expression exists. Our approach to deal with this problem is to ensure that the $100(1 - \alpha)\%$ confidence ellipsoid of the random vector $\eta(\bar{x}_\theta, u_\theta)$ is contained in the negative orthant. This relaxation will imply that the probability of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ being in the negative orthant is strictly greater than $1 - \alpha$.

For an r -dimensional random vector X with distribution $\mathcal{N}(h, \Sigma')$, we define the ellipsoids

$$\mathcal{E}(X, \beta) = \{x \mid (x - h)^T \Sigma'^{-1} (x - h) \leq \beta\}, \quad (12)$$

where $\beta \geq 0$ is a parameter specifying the size of the ellipsoid $\mathcal{E}(X, \beta)$. Since X is $\mathcal{N}(h, \Sigma')$, we know that the random variable $(X - h)^T \Sigma'^{-1} (X - h)$ has a χ_r^2 distribution. Thus, for a given value of α , one can choose $\beta(\alpha)$ from the distribution χ_r^2 such that the corresponding ellipsoid $\mathcal{E}(X, \beta(\alpha))$ is a $100(1 - \alpha)\%$ confidence ellipsoid. We have the following

Proposition 2: Let $\alpha \in]0, 1[$ be given, and suppose η defined in (11) takes values in \mathbb{R}^r , $r \in \mathbb{N}$. Then the constraint $\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha$ can be conservatively approximated by $\mathcal{E}(\eta(\bar{x}_\theta, \bar{u}_\theta), \beta(\alpha)) \subset]-\infty, 0]^r$, where \mathcal{E} is as defined in (12) and $\beta(\alpha)$ is chosen as above. Moreover, the latter constraint can be modelled as a sequence of r second-order cone constraints.

Proof: Let g denote the probability distribution function of a χ_r^2 random variable. Since $\alpha \in]0, 1[$ is given, we select parameter $\beta(\alpha) = g^{-1}(1 - \alpha)$ such that ellipsoid $\mathcal{E}(\eta(\bar{x}_\theta, \bar{u}_\theta), \beta(\alpha))$ is the $100(1 - \alpha)\%$ ellipsoid. The first claim follows immediately from the definition of \mathcal{E} . To verify the second claim, note that the ellipsoid

$$\mathcal{E}(\eta(\bar{x}_\theta, \bar{u}_\theta), \beta(\alpha)) = \{x \mid (x - h_\theta)^T \bar{\Sigma}_\theta^{-1} (x - h_\theta) \leq \beta\}$$

can alternatively be represented as

$$\mathcal{E}(\eta(\bar{x}_\theta, \bar{u}_\theta), \beta(\alpha)) = \{y \in \mathbb{R}^r \mid y = M_\theta u + h_\theta, \|u\| \leq 1\}, \quad (13)$$

where $M_\theta = \sqrt{\beta} \bar{\Sigma}_\theta^{\frac{1}{2}}$. Since $y \leq 0 \iff e_i^T y \leq 0$, $i = 1, \dots, r$, where e_i denote the standard basis vectors, we may rewrite (13) as $e_i^T (M_\theta u + h_\theta) \leq 0 \quad \forall \|u\| \leq 1$, or equivalently $\sup_{\|u\| \leq 1} e_i^T (M_\theta u + h_\theta) \leq 0$. The supremum is attained for $u = M_\theta^T e_i / \|M_\theta^T e_i\|$; therefore, the above is equivalent to $e_i^T h_\theta + \|M_\theta^T e_i\| \leq 0$. Since $M_\theta = \sqrt{\beta} \bar{\Sigma}_\theta^{\frac{1}{2}}$, we have $\|M_\theta^T e_i\| = \sqrt{\beta e_i^T \bar{\Sigma}_\theta e_i} = \sqrt{\beta (\bar{\Sigma}_\theta)_{ii}}$. Finally,

since $\bar{\Sigma}_\theta = P_\theta \bar{\Sigma} P_\theta^T$, we have $(\bar{\Sigma}_\theta)_{ii} = p_{i,\theta}^T \bar{\Sigma} p_{i,\theta}$, and thus $\sqrt{(\bar{\Sigma}_\theta)_{ii}} = \left\| \bar{\Sigma}^{\frac{1}{2}} p_{i,\theta} \right\|$, where $p_{i,\theta}^T$ denotes the i^{th} row of P_θ . Thus, the constraint $\mathcal{E}(\eta(\bar{x}_\theta, \bar{u}_\theta), \beta(\alpha)) \subset]-\infty, 0]^r$ reduces to

$$e_i^T h_\theta + \sqrt{\beta(\alpha)} \left\| \bar{\Sigma}^{\frac{1}{2}} p_{i,\theta} \right\| \leq 0 \quad \forall i = 1, \dots, r.$$

This is an intersection of second order cone constraints since $\beta(\alpha), \bar{\Sigma}$ are known and the variables (P_θ, h_θ) are affine in the original parameters $\theta = (G, \bar{d})$. ■

3) Quality of approximation via confidence ellipsoids:

Simulation results indicate that the number $\beta(\alpha)$ increases quite rapidly with the dimension of the vector $\eta(\bar{x}_\theta, \bar{u}_\theta)$. An explanation of this phenomenon is provided by the following fact, that is better known by the name of (classical) ‘‘concentration of measure’’ inequalities; proofs may be found in, e.g., [6].

Proposition 3: Let $\Gamma_{h,\Sigma'}$ be the r -dimensional Gaussian measure with mean h and (nonsingular) covariance Σ' , i.e.,

$$\begin{aligned} \Gamma_{h,\Sigma'}(d\xi) &= \frac{1}{(2\pi)^{r/2} \sqrt{\det \Sigma'}} \exp\left(-\frac{1}{2} \langle \xi - h, \Sigma'^{-1}(\xi - h) \rangle\right) d\xi. \end{aligned}$$

Then for $\varepsilon \in]0, 1[$,

- (i) $\Gamma_{h,\Sigma'}\left(\left\{\xi \mid \|\xi - h\|_{\Sigma'^{-1}} > \sqrt{\frac{r}{1-\varepsilon}}\right\}\right) \leq e^{-\frac{r\varepsilon^2}{4}}$;
- (ii) $\Gamma_{h,\Sigma'}\left(\left\{\xi \mid \|\xi - h\|_{\Sigma'^{-1}} < \sqrt{r(1-\varepsilon)}\right\}\right) \leq e^{-\frac{r\varepsilon^2}{4}}$.

The above proposition states that as the dimension r of the Gaussian measure increases, the mass concentrates in an ellipsoidal shell of ‘mean-size’ \sqrt{r} . It readily follows that since $\eta(\bar{x}_\theta, \bar{u}_\theta)$ is a r -dimensional Gaussian random vector, its mass concentrates around a shell of size \sqrt{r} . Note that the bounds corresponding to (i) and (ii) of Proposition 3 in the case of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ are independent of the optimization parameters θ ; of course the relative sizes of the confidence ellipsoids change with θ (because the mean and the covariance of $\eta(\bar{x}_\theta, \bar{u}_\theta)$ depend on θ), but Proposition 3 shows that the size of the confidence ellipsoids grows quite rapidly with the dimension of the noise or the length of the optimization horizon. Further analysis and implications of this phenomenon will be reported elsewhere.

Remark 4: In the above we have considered the variable $\eta(\bar{x}_\theta, \bar{u}_\theta)$ to be a χ_r^2 random vector, where r is the dimension of η . This is true if $r \leq Nn$, but otherwise r should be replaced by $\min\{r, Nn\}$ because there are Nn independent Gaussian random variables in the optimization problem.

B. Ellipsoidal Constraint Functions

Consider the constraint function

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = \left(\begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix} - \delta \right)^T \Xi \left(\begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix} - \delta \right) - 1,$$

where $\Xi \geq 0$ and δ are given. Then the constraint $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ restricts the vector $\begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix}$ to an ellipsoid with

center δ and shape determined by Ξ . In this section we relax the hard constraint $\eta \leq 0$ in (4) by the chance constraint

$$\mathbb{P}(\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0) \geq 1 - \alpha \quad (14)$$

for some prespecified value of $\alpha \in]0, 1[$, and give an approximation of this chance constraint that is a semi-definite program in the optimization parameters. To this end, let

$$\begin{aligned} y_\theta &= \begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix} - \delta = \begin{bmatrix} \bar{A}x_0 + \bar{B}(\bar{G}\bar{w} + \bar{d}) + \bar{D}\bar{w} \\ \bar{G}\bar{w} + \bar{d} \end{bmatrix} - \delta \\ &= h'_\theta + P'_\theta \bar{w}, \end{aligned}$$

where $h'_\theta = \begin{bmatrix} \bar{A}x_0 + \bar{B}\bar{d} \\ \bar{d} \end{bmatrix} - \delta$ and $P'_\theta = \begin{bmatrix} \bar{B}\bar{G} + \bar{D} \\ \bar{G} \end{bmatrix}$. The idea is to ensure that the $100(1-\alpha)\%$ confidence ellipsoid of \bar{w} is such that (14) holds, and ensure that this reformulated problem is convex.

Proposition 5: Let $S'_\theta = \sqrt{\beta(\alpha)} \Xi^{\frac{1}{2}} P'_\theta \bar{\Sigma}^{\frac{1}{2}}$, with $\beta(\alpha)$ computed as in Section IV-A.2, and $\xi_\theta = \Xi^{\frac{1}{2}} h'_\theta$. Then

$$\begin{bmatrix} -\lambda + 1 & 0 & \xi_\theta^T \\ 0 & \lambda I & (S'_\theta)^T \\ \xi_\theta & S'_\theta & I \end{bmatrix} \geq 0, \quad \lambda > 0 \quad (15)$$

is a Linear Matrix Inequality (LMI) in the unknown parameters θ and $\lambda \in \mathbb{R}$. If (θ, λ) is a solution of (15), then θ satisfies (14).

Proof: The inequality (14) may be equivalently represented as

$$\begin{aligned} \mathbb{P}(y_\theta^T \Xi y_\theta - 1 \leq 0) &\geq 1 - \alpha \\ \iff \mathbb{P}(\|\Xi^{\frac{1}{2}} y_\theta\|^2 - 1 \leq 0) &\geq 1 - \alpha \\ \iff \mathbb{P}(\|\Xi^{\frac{1}{2}} (h'_\theta + P'_\theta \bar{w})\|^2 - 1 \leq 0) &\geq 1 - \alpha \\ \iff \mathbb{P}(\|\xi_\theta + S_\theta \bar{w}\|^2 - 1 \leq 0) &\geq 1 - \alpha, \end{aligned} \quad (16)$$

where $S_\theta = \Xi^{\frac{1}{2}} P'_\theta$. Since \bar{w} has distribution $\mathcal{N}(0, \bar{\Sigma})$, one can compute $\beta(\alpha)$ such that $\|\bar{\Sigma}^{-1/2} \bar{w}\|^2 \leq \beta(\alpha)$ specifies the required $100(1-\alpha)\%$ confidence ellipsoid of \bar{w} . Hence, we need to ensure that $\|\bar{\Sigma}^{-1/2} \bar{w}\|^2 \leq \beta(\alpha) \Rightarrow \|\xi_\theta + S_\theta \bar{w}\|^2 \leq 1$. This is equivalent to

$$\begin{aligned} \sup_{\|\bar{\Sigma}^{-1/2} \bar{w}\|^2 \leq \beta(\alpha)} \|\xi_\theta + S_\theta \bar{w}\|^2 &\leq 1 \\ \iff \sup_{\|\bar{v}\|^2 \leq 1} \|\xi_\theta + S'_\theta \bar{v}\|^2 &\leq 1 \text{ where } S'_\theta = \sqrt{\beta(\alpha)} S_\theta \bar{\Sigma}^{1/2}. \end{aligned}$$

We know from [8, p. 653] that $\sup_{\|\bar{v}\| \leq 1} \|\xi_\theta + S'_\theta \bar{v}\|^2 \leq 1$ if and only if there exists $\lambda \geq 0$ such that

$$\begin{bmatrix} -\xi_\theta^T \xi_\theta - \lambda + 1 & \xi_\theta^T S'_\theta \\ (S'_\theta)^T \xi_\theta & \lambda I - (S'_\theta)^T S'_\theta \end{bmatrix} \geq 0.$$

Using Schur complements the last relation is found to be equivalent to (15). This shows that the solutions of (15) are solutions of (14) as well. To verify that (15) is an LMI, note that S'_θ and ξ_θ are affine in the optimization variables $\theta = (G, \bar{d})$. ■

With the constraint (15), problem (4) is a semidefinite program in the optimization parameters θ (and λ). The solutions of this problem conservatively approximate the solutions of problem (4) with constraints (14).

V. INTEGRATED CHANCE CONSTRAINTS

Let, as before, $\theta = (\bar{G}, \bar{d})$ and let \bar{x}_θ and \bar{u}_θ express the dependence of \bar{x} and \bar{u} on θ . Consider a scalar measurable constraint function $\eta : \mathbb{R}^{(N+1)n \times Nm} \rightarrow \mathbb{R}$ (i.e. $r = 1$) and a measurable function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$. For a given $\beta > 0$, consider the constraint

$$\mathbb{E}[\varphi \circ \eta(\bar{x}_\theta, \bar{u}_\theta)] \leq \beta. \quad (17)$$

For the particular choice

$$\varphi(z) = z^+ := \begin{cases} 0, & \text{if } z \leq 0, \\ z, & \text{otherwise,} \end{cases} \quad (18)$$

constraints of the form (17) are known as integrated chance constraints [12], [13]. In fact, one may write (dropping the arguments of η from the notation)

$$\mathbb{E}[\varphi(\eta)] = \mathbb{E} \left[\int_0^{+\infty} 1_{[0, \eta)}(s) ds \right] = \int_0^{+\infty} \mathbb{P}[\eta > s] ds,$$

where $1_S(\cdot)$ is the indicator function of set S and the second equality follows from Tonelli's theorem [10, Theorem 4.4.5]. Therefore, for ϕ as in (18) the constraint (17) is equivalent to

$$\int_0^{+\infty} \mathbb{P}[\eta(\bar{x}_\theta, \bar{u}_\theta) > s] ds \leq \beta, \quad (19)$$

whence the name integrated chance constraint. Note that φ plays the role of a penalty (or barrier) function that penalizes larger violations of the inequality $\eta(\bar{x}_\theta, \bar{u}_\theta) \leq 0$ more than smaller ones, and β is a maximum allowable cost in the sense of (19). Of course, different choices of φ may be motivated by different applications, but will not guarantee an equivalence between (17) and (19). It should be noted that in general integrated chance constraints cannot be directly compared with chance constraints; in particular, they cannot be viewed as relaxations of chance constraints.

Now let us consider again the particular case of a scalar polytopic constraint function defined in (11). Consider the function $\phi(z) = z^+$ as above. Since \bar{w} is Gaussian with distribution $\mathcal{N}(0, \bar{\Sigma})$, it is clear that η is also Gaussian, with distribution $\mathcal{N}(h_\theta, \|P_\theta\|_{\bar{\Sigma}}^2)$.

Proposition 6: The scalar constraint $\mathbb{E}[(\eta(\bar{x}_\theta, \bar{u}_\theta))^+] \leq \beta$ with η defined in (11), is convex in θ , and moreover, is equivalent to $\|P_\theta\|_{\bar{\Sigma}} g\left(\frac{h_\theta}{\|P_\theta\|_{\bar{\Sigma}}}\right) \leq \beta$, where $g(\xi) := \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\xi^2}{2}\right) + \frac{\xi}{2} \operatorname{erfc}\left(\frac{-\xi}{\sqrt{2}}\right)$.

Proof: It is not difficult to prove that if $z \sim \mathcal{N}(\mu, \sigma^2)$ then, for $\sigma > 0$, $\mathbb{E}[z^+] \leq \beta \iff \sigma g(\mu/\sigma) \leq \beta$, where

$$g(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) + \frac{x}{2} \operatorname{erfc}\left(\frac{-x}{\sqrt{2}}\right).$$

We claim that $g(x)$ is convex. Indeed, straightforward calculations show that

$$g'(x) = \frac{1}{2} \operatorname{erfc}\left(\frac{-x}{\sqrt{2}}\right) \quad \text{and} \quad g''(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right).$$

Clearly $g'' > 0$, which proves that g is convex. Also recall that for a convex function $g : \mathbb{R}^n \rightarrow \mathbb{R}$, the *perspective* [8, p. 89] of g is the function $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ defined as

$f(x, t) = tg(x/t)$ with domain $\operatorname{dom} f = \{(x, t) \mid x/t \in \operatorname{dom} g, t > 0\}$. It can be shown that perspective operation preserves convexity [8]. Thus, if g is convex, so is its perspective function f . Thus we conclude that $f(\mu, \sigma) = \sigma g(\frac{\mu}{\sigma})$ is convex. Note that f is increasing in σ since

$$\frac{df(\mu, \sigma)}{d\sigma} = g(\mu/\sigma) - \frac{\mu}{\sigma} g'(\mu/\sigma) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) > 0.$$

Using the composition rules of convexity, we know that a convex function, increasing in a convex argument is still convex. The above arguments imply that we have a convex constraint

$$\mathbb{E}[(\eta(\bar{x}_\theta, \bar{u}_\theta))^+] \leq \beta \iff \|P_\theta\|_{\bar{\Sigma}} g\left(\frac{h_\theta}{\|P_\theta\|_{\bar{\Sigma}}}\right) \leq \beta,$$

which proves the claims. \blacksquare

The above proposition extends readily to the case of joint integrated chance constraints, namely, where the constraint function η is vector-valued, and the ICC is interpreted componentwise. In this case, the set of parameters for a joint integrated chance constraint is the intersection of the set of parameters corresponding to each of the (scalar) components, and each of the latter is a convex set.

VI. EXAMPLE

In this section we provide simulation results illustrating some of the results of this paper. Consider the linear dynamical system (1) with

$$A = \begin{bmatrix} 0.93 & 0.85 & 0.67 \\ 0.47 & 0.53 & 0.84 \\ 0.42 & 0.20 & 0.02 \end{bmatrix}, \quad B = \begin{bmatrix} 0.68 & 0.50 \\ 0.38 & 0.71 \\ 0.83 & 0.43 \end{bmatrix}.$$

The noise vector w is assumed to be a 3-dimensional Gaussian vector with mean 0 and covariance matrix I , i.e., $w \sim \mathcal{N}(0, I)$. Let the matrix $M = I$, which means that the cost function $V(\bar{x}, \bar{u})$ in (3) is $\|\bar{x}\|^2 + \|\bar{u}\|^2$.

A. Affine Chance Constraints: Approximation via Constraint Separation

We consider the chance constraint (10) corresponding to the hard constraint given by $\|x\|_\infty \leq 10, \|u\|_\infty \leq 10$. We follow the technique described in IV-A.1. These constraints are equivalent to the polytopic constraint $\eta(\bar{x}, \bar{u}) \leq 0$ where η defined is as defined in (9), and $T^x = [I \ -I \ 0 \ 0]^T$, $T^u = [0 \ 0 \ I \ -I]^T$, $y = [10 \ \dots \ 10]^T$. We fix the initial state $x_0 = [6, 0, -8]^T$, the length of the horizon $N = 9$, and $\alpha = 0.06$. A typical trajectory of the state and the control input are shown in Fig 1(a). The (red) horizontal straight lines indicate the hard bounds on state and control action.

B. Affine Chance Constraints: Approximation via Confidence Ellipsoids

We address approximating the affine chance constraints described in Section VI-A via confidence ellipsoids as described in Section IV-A.2. We fix the length of the horizon $N = 6$, and $\alpha = 0.06$. Fig 1(b) shows one outcome of the state and control input trajectories. As above, horizontal lines indicate the bounds on state and control.




Fig. 1. A sample outcome of the state trajectory (above) and the corresponding control input (below) for problems VI-A (a), VI-B (b) and VI-C (c).

C. Ellipsoidal Chance Constraints

In this section we address chance constraints with ellipsoidal constraint functions. Let the constraint function be defined as

$$\eta(\bar{x}_\theta, \bar{u}_\theta) = \left(\begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix} - \delta \right)^T Q \left(\begin{bmatrix} \bar{x}_\theta \\ \bar{u}_\theta \end{bmatrix} - \delta \right) - 1.$$

The matrix Q is chosen to be $\frac{1}{5N(n+m)}I$ and $\delta = 0$. This inverse dependence of Q upon $N(m+n)$ is due to the fact that the length of the vector $[\bar{x}_\theta^T, \bar{u}_\theta^T]^T$ is $N(m+n)$, and without scaling the problem becomes infeasible simply because of the increase in the length of the vector. The choice $\delta = 0$ reflects that we want to operate the system near the origin, i.e. both inputs and states should be as close to zero as possible. We fix the initial state to be $x_0 = [2, 0, -2]^T$, the length of the horizon $N = 5$ and $\alpha = 0.10$. One particular realization of the state and control input trajectories is shown in Fig 1(c).

REFERENCES

- [1] J. M. ALDEN AND R. L. SMITH, *Rolling horizon procedures in nonhomogeneous Markov decision processes*, *Operations Research*, 40 (1992), pp. S183–S194.
- [2] I. BATINA, *Model predictive control for stochastic systems by randomized algorithms*, PhD thesis, Technische Universiteit Eindhoven, 2004.
- [3] A. BEN-TAL, A. GORYASHKO, E. GUSLITZER, AND A. NEMIROVSKI, *Adjustable robust solutions of uncertain linear programs*, *Mathematical Programming*, 99 (2004), pp. 351–376.
- [4] D. P. BERTSEKAS, *Dynamic programming and suboptimal control: A survey from ADP to MPC*, *European Journal of Control*, 11 (2005).
- [5] D. BERTSIMAS AND D. B. BROWN, *Constrained stochastic LQC: a tractable approach*, *IEEE Transactions on Automatic Control*, 52 (2007), pp. 1826–1841.
- [6] V. I. BOGACHEV, *Gaussian Measures*, vol. 62 of *Mathematical Surveys and Monographs*, American Mathematical Society, Providence, RI, 1998.
- [7] V. S. BORKAR, *A convex analytic approach to Markov decision processes*, *Probability Theory and Related Fields*, 78 (1988), pp. 583–602.
- [8] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, Cambridge, 2004.
- [9] P. D. COUCHMAN, M. CANNON, AND B. KOUVARITAKIS, *Stochastic MPC with inequality stability constraints*, *Automatica J. IFAC*, 42 (2006), pp. 2169–2174.
- [10] R. M. DUDLEY, *Real Analysis and Probability*, vol. 74 of *Cambridge Studies in Advanced Mathematics*, Cambridge University Press, Cambridge, 2002. Revised reprint of the 1989 original.
- [11] P. J. GOULART, E. C. KERRIGAN, AND J. M. MACIEJOWSKI, *Optimization over state feedback policies for robust control with constraints*, *Automatica*, 42 (2006), pp. 523–533.
- [12] W. K. K. HANEVELD, *On integrated chance constraints*, in *Stochastic programming (Gargnano)*, vol. 76 of *Lecture Notes in Control and Inform. Sci.*, Springer, Berlin, 1983, pp. 194–209.
- [13] W. K. K. HANEVELD AND M. H. VAN DER VLERK, *Integrated chance constraints: reduced forms and an algorithm*, *Computational Management Science*, 3 (2006), pp. 245–269.
- [14] J. MACIEJOWSKI, *Predictive Control with Constraints*, Prentice Hall, 2002.
- [15] A. NEMIROVSKI AND A. SHAPIRO, *Convex approximations of chance constrained programs*, *SIAM Journal on Control and Optimization*, 17 (2006), pp. 969–996.
- [16] J. PRIMBS, *A soft constraint approach to stochastic receding horizon control*, in *Proceedings of the 46th IEEE Conference on Decision and Control*, 2007, pp. 4797 – 4802.
- [17] J. PRIMBS, *Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise*, in *Proceedings of the 26th American Control Conference*, 2007.
- [18] D. H. VAN HESSEM AND O. H. BOSGRA, *Stochastic closed-loop model predictive control of continuous nonlinear chemical processes*, *Journal of Process Control*, 16 (2006), pp. 225–241.
- [19] ———, *A full solution to the constrained stochastic closed-loop mpc problem via state and innovations feedback and its receding horizon implementation*, in *Proceedings of the 42nd IEEE Conference on Decision and Control*, vol. 1, 9-12 Dec. 2003, pp. 929–934.