Convex Chance Constrained Predictive Control without Sampling

Lars Blackmore*

***Abstract*— In this paper we consider finite-horizon predictive control of dynamic systems subject to stochastic uncertainty; such uncertainty arises due to exogenous disturbances, modeling errors, and sensor noise. Stochastic robustness is typically defined using chance constraints, which require that the probability of state constraints being violated is below a prescribed value.**

Prior work showed that in the case of linear system dynamics, Gaussian noise and convex state constraints, optimal chance-constrained predictive control results in a convex optimization problem. Solving this problem in practice, however, requires the evaluation of multivariate Gaussian densities through sampling, which is time-consuming and inaccurate.

We propose a new approach to chance-constrained predictive control that does not require the evaluation of multivariate densities. We use a new bounding approach to ensure that chance constraints are satisfied, while showing empirically that the conservatism introduced is small. This is in contrast to prior bounding approaches that are extremely conservative. Furthermore we show that the resulting optimization is convex, and hence amenable to online control design.

I. INTRODUCTION

Robust predictive (or *finite horizon*) control of systems subject to stochastic uncertainty has received a great deal of attention in recent years[1], [2], [3], [4], [5], [6], [7], [8], [9]. Stochastic models can be used to characterize, for example, exogenous disturbances, modeling error and sensor noise. In many cases, stochastic uncertainty models are more realistic than set-bounded models, for example in the case of wind disturbances. Previous work has posed problems such as Unmanned Air Vehicle (UAV) path planning [7], chemical reactor control [2] and network traffic control [5] as robust predictive control under stochastic uncertainty.

For a predictive controller to be *robust*, it must take into account uncertainty so that state constraints, such as obstacle avoidance constraints, are not violated. With most stochastic uncertainty models however, it is not possible to guarantee that state constraints are satisfied, since there is always some small probability of an arbitrarily large disturbance occurring. Previous work therefore described robustness in terms of *chance constraints*, which require that the probability of state constraint violation is below a prescribed value. By setting this value appropriately, the operator can trade conservatism against performance; a control strategy that is less risky will typically take more fuel or time (and vice versa).

*L. Blackmore is with the NASA Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109 USA e-mail: lars@jpl.nasa.gov.

A number of approaches to chance-constrained predictive control have been proposed in recent years. In the case of Gaussian uncertainty distributions, [1] considered chance constraints on individual scalar values, while [2], [3] considered chance constraints on joint random variables. This extension to joint random variables is essential if we wish to constrain the probability of failure over the entire planning horizon. [7] considered control in nonconvex feasible regions, while [8] extended this line of research to arbitrary probability distributions and hybrid discrete-continuous systems. The problem of explicitly optimizing feedback design as well as feedforward controls was considered by [4], [6]. Other related work includes that of [9], which provides early results on chance-constrained approaches in receding horizon.

By assuming Gaussian distributions for the uncertain variables, and convexity of the feasible region, the work of [2], [3] uses the result of [10] to show that the optimization resulting from the chance-constrained predictive control problem is convex, and can therefore be solved effectively using standard nonlinear solvers. This approach is limited, however, by the need to evaluate the multivariate Gaussian integrals in the constraint functions. These integrals are approximated through sampling, which is time-consuming and leads to approximation error. In this paper we present a new approach that solves the chance-constrained predictive control problem without the need for sampling. The key idea is to bound the joint probability of multivariate constraint violation conservatively using Boole's inequality. This leads to constraints involving the sum of many univariate probabilities, which can be evaluated efficiently. We show that the resulting optimization is convex, and can therefore be solved efficiently using nonlinear solvers. Critically, we show with empirical examples that the conservatism introduced by the bounding approach is small. We show also that this is in contrast with prior bounding approaches that are conservative by very many orders of magnitude.

II. PROBLEM STATEMENT

In this paper, we are concerned with the following discrete-time Linear Time Invariant (LTI) plant:

$$\mathbf{x}_{k+1} = A\mathbf{x}_k + B^w\mathbf{w}_k + B\mathbf{u}_k, \tag{1}$$

where $\mathbf{x} \in \Re^{n_x}$ is the system state, $\mathbf{u} \in \Re^{n_u}$ are the system inputs, and $\mathbf{w} \in \Re^{n_w}$ is a noise vector. The noise vector can model disturbances, uncertainty in the system model, and sensor noise. We assume that $\mathbf{w}$ is a Gaussian noise process and that the initial state $\mathbf{x}_0$ is a Gaussian random variable; these two are uncorrelated. We use $\mathbf{x}_k$ to denote the value of

$\mathbf{x}$ at time step k, and $\mathbf{x}'$ to denote the transpose of $\mathbf{x}$. We use $P(A)$ to denote the probability of event A and $p(\mathbf{x})$ to denote the probability distribution function of random variable $\mathbf{x}$. We use $\bar{\mathbf{x}}$ to denote the mean of the random variable $\mathbf{x}$, and use $S_{\mathbf{x}}$ to denote its covariance. Note that the plant definition (1) can model an LTI plant with a fixed-gain linear feedback; we will use this in Section VI.

In finite-horizon predictive control, we plan over a finite horizon of time instances from $k = 0$ to $k = T$. For notational convenience we 'lift' the variables of interest over the time horizon using the following definitions:

$$\mathbb{X} = \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_T \end{bmatrix} \quad \mathbb{U} = \begin{bmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_T \end{bmatrix} \quad \mathbb{W} = \begin{bmatrix} \mathbf{w}_0 \\ \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_T \end{bmatrix}. \tag{2}$$

The lifted system dynamics are given by:

$$\mathbb{X} = G_{xx}\mathbf{x}_0 + G_{xu}\mathbb{U} + G_{xw}\mathbb{W}, \tag{3}$$

where the matrices G_{xx}, G_{xu}, G_{xw} are calculated through repeated multiplication of the system matrices in (1), see for example [4]. The chance-constrained predictive control problem can now be stated.

Problem 1 (Chance-constrained control problem).

$$\begin{array}{ll} \textit{Minimize} & f(\bar{\mathbb{X}}, \mathbb{U}) \\ \textit{Subject to:} & \mathbb{U} \in F_U \\ & P(\mathbb{X} \notin F_X) \leq \delta \\ & \mathbb{X} = G_{xx}\mathbf{x}_0 + G_{xu}\mathbb{U} + G_{xw}\mathbb{W}. \end{array}$$

In other words, we must choose the control inputs to minimize cost, while ensuring that the system state leaves the feasible region with probability at most δ. We assume that the cost function $f(\mathbb{X}, \mathbb{U})$ is convex in $\mathbb{X}$ and $\mathbb{U}$, the control constraint set F_U is convex, and the state constraint set F_X is a convex polytope.

III. Existing Results

While notationally simple, Problem 1 is made challenging by the chance constraint $P(\mathbb{X} \notin F_X) \leq \delta$. There are three key challenges resulting from this constraint. First, we must determine the distribution of $\mathbb{X}$ as a function of the control inputs $\mathbb{U}$. Second, we must perform a multidimensional integral over this distribution. Finally, we must optimize with constraints on this integral.

The first of these challenges is removed by the assumptions of linear system dynamics and Gaussian noise. In this case, the system state $\mathbb{X}$ is a Gaussian random variable with mean and covariance given explicitly by:

$$\begin{aligned} \bar{\mathbb{X}} &= G_{xx}\bar{\mathbf{x}}_0 + G_{xu}\mathbb{U} + G_{xw}\bar{\mathbb{W}} \\ S_{\mathbb{X}} &= G_{xx}S_{\mathbf{x}_0}G'_{xx} + G_{xw}S_{\mathbb{W}}G'_{xw}. \end{aligned} \tag{4}$$

Problem 1 can now be restated as follows:

Problem 2 (Linear-Gaussian control problem).

$$\begin{array}{ll} \textit{Minimize} & f(\bar{\mathbb{X}}, \mathbb{U}) \\ \textit{Subject to:} & \mathbb{U} \in F_U \\ & \int_{\mathbf{z} \notin F_X} \mathcal{N}(\bar{\mathbb{X}}, S_{\mathbb{X}}) d\mathbf{z} \leq \delta \\ & \bar{\mathbb{X}} = G_{xx}\bar{\mathbf{x}}_0 + G_{xu}\mathbb{U} + G_{xw}\bar{\mathbb{W}}, \end{array}$$

where $\mathcal{N}(\cdot)$ is the multivariate normal distribution:

$$\mathcal{N}(\bar{\mathbb{X}}, S_X) = \frac{1}{(2\pi)^{n_x/2}|S_X|^{1/2}} e^{-\frac{1}{2}(\mathbf{z}-\bar{\mathbf{x}})'(S_x^{-1})(\mathbf{z}-\bar{\mathbf{x}})}. \tag{5}$$

Note that, since S_X is not a function of the control inputs $\mathbb{U}$, it can be precomputed.

Prior worked used this result, together with the convexity result of [10], to show that Problem 2 is convex[2], [3]. Convexity of an optimization problem means that a local optimum is also a global optimum (first-order necessary conditions for global optimality are also sufficient), and that standard nonlinear solvers can find such optima efficiently. Hence [2] showed that the chance constrained control problem can be solved, in principle, using nonlinear solvers. Practical implementation of this method, however, requires evaluation of the multidimensional integral:

$$I(\mathbb{U}) = \int_{\mathbf{z} \notin F_X} \mathcal{N}(\bar{\mathbb{X}}, S_{\mathbb{X}}) d\mathbf{z}. \tag{6}$$

This integral cannot be evaluated in closed form. As a result [2] use a sampling approach to approximate the value and its derivatives. In a control problem with $n_x = 4$ and $T = 20$, the value (6) is an integral in 84 dimensions, hence achieving a good approximation requires a very large number of samples. Performing this sampling procedure at each iteration of the optimization is time-consuming and hence limits the applicability of the approach to real-time control problems. Furthermore, since the sample-approximated constraint function used in the optimization is now a random variable, the theoretical guarantees of convexity no longer apply. In the next section we present a new approach that does not require this sampling procedure.

IV. New Approach

The new approach can be summarized as follows. First we pose an alternative form of Problem 2 that does not require evaluation of multivariate densities, which we call the *conservative problem*. Then we show that a feasible solution to the conservative problem is a feasible solution to Problem 2. Next we show that the conservative problem is convex. Finally, in Section VI we show empirically that the conservatism introduced is small.

A. The Conservative Problem

The convex polytopic feasible region F_X can be defined by a conjunction of N linear inequality constraints:

$$F_X \triangleq \bigcap_{i=1}^{N} \{\mathbb{X} : \mathbf{a}'_i\mathbb{X} \leq b_i\}. \tag{7}$$

Now consider the following problem:

Problem 3 (Conservative Problem).

$$\begin{array}{ll} \textit{Minimize} & f(\bar{\mathbb{X}}, \mathbb{U}) \\ \textit{Subject to:} & \mathbb{U} \in F_U \\ & P(\mathbf{a}_i'\mathbb{X} > b_i) \leq \epsilon_i \ \ \forall i \\ & \sum_i^N \epsilon_i \leq \delta \\ & \mathbb{X} = G_{xx}\mathbf{x}_0 + G_{xu}\mathbb{U} + G_{xw}\mathbb{W}. \end{array}$$

Lemma 1. *A feasible solution to Problem 3 (the conservative problem) is a feasible solution to Problem 1 (the chance-constrained problem).*

Proof: From (7):

$$\mathbf{x} \notin F \Longleftrightarrow \mathbf{x} \notin \bigcup_{i=1}^{N} \{\mathbf{x} : \mathbf{a}_i'\mathbf{x} > b_i\}. \tag{8}$$

Boole's inequality gives us the following bound for a countable set of events $A_1, \ldots, A_N$:

$$P\left[\bigcup_{i=1}^{N} A_i\right] \leq \sum_i P(A_i). \tag{9}$$

Setting A_i as the event $\{\mathbf{a}_i'\mathbf{x} > b_i\}$ gives us:

$$P(\mathbf{x} \notin F) \leq \sum_i P(\mathbf{a}_i'\mathbf{x} > b_i). \tag{10}$$

For a feasible solution to Problem 3 we know that $\sum_i P(\mathbf{a}_i'\mathbf{x} > b_i) \leq \delta$, and hence $P(\mathbf{x} \notin F) \leq \delta$. The constraints in Problem 1 are therefore satisfied by a feasible solution to Problem 3. □

Problem 3 is therefore a conservative approximation of Problem 1. In this paper we propose to solve this approximation instead of solving Problem 1. The intuition is that, in solving Problem 3, we explicitly optimize the probability of each individual constraint being violated, denoted ϵ_i. This idea of *risk allocation* was previously proposed by [11], however in that work the optimization of the ϵ_i was carried out in a separate optimization step. Other related work[7] used Boole's inequality to generate conservative solutions, but assumed that the ϵ_i were equal and fixed *a priori*, which leads to unnecessary conservatism. In the following sections we show that Problem 3 is convex and does not require the evaluation of multidimensional integrals.

B. Constraint Evaluation

The key difference between Problem 3 and Problem 1 is that Problem 3 no longer involves multivariate integrals. Instead, it has N constraints on *univariate* integrals. To see this, define $y_i \triangleq \mathbf{a}_i'\mathbf{x}$ and note that y_i is a univariate Gaussian random variable with mean and variance given by:

$$\begin{aligned} \bar{y}_i &= \mathbf{a}_i'G_{xx}\bar{\mathbf{x}}_0 + \mathbf{a}_i'G_{xu}\mathbb{U} + \mathbf{a}_i'G_{xw}\bar{\mathbb{W}} \\ S_{y_i} &= \mathbf{a}_i'G_{xx}S_{\mathbf{x}_0}G_{xx}'\mathbf{a}_i + \mathbf{a}_i'G_{xw}S_{\mathbb{W}}G_{xw}'\mathbf{a}_i. \end{aligned} \tag{11}$$

We can now write the probability of each individual constraint being violated as follows:

$$P(\mathbf{a}_i'\mathbb{X} > b_i) = P(y_i > b_i) = \frac{1}{\sqrt{2\pi S_{y_i}}} \int_{b_i}^{\infty} e^{-\frac{(y_i - \bar{y}_i)^2}{2S_{y_i}}} dy_i. \tag{12}$$

Because this is a singlevariate integral, we can express this in terms of the standard singlevariate Gaussian cdf:

$$P(\mathbf{a}_i'\mathbb{X} > b_i) = \frac{1}{\sqrt{2\pi}} \int_{\frac{b_i - \bar{y}_i}{\sqrt{S_{y_i}}}}^{\infty} e^{\frac{-z^2}{2}} dz = 1 - \mathrm{cdf}\left(\frac{b_i - \bar{y}_i}{\sqrt{S_{y_i}}}\right), \tag{13}$$

where cdf is the standard Gaussian cumulative distribution function:

$$\mathrm{cdf}(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{z^2}{2}} dz. \tag{14}$$

In order to evaluate each constraint, we therefore need to evaluate $\mathrm{cdf}(\cdot)$ only once. While $\mathrm{cdf}(\cdot)$ cannot be evaluated in closed form, it can be evaluated quickly and accurately using a series expansion or a one-dimensional lookup. In order to evaluate every constraint requires N such lookups, where N is the number of state constraints. This is significantly less computationally intensive than drawing the very large number of samples necessary to approximate the multidimensional integral (6) to the same precision. Hence constraints in the conservative control problem (Problem 3) can be evaluated far more efficiently than in the chance constrained control problem (Problem 1).

C. Gradient Evaluation

In this section we show that derivatives of the constraints in Problem 3 can also be computed efficiently, without the need for sampling. Specifically we want to compute the gradient of $P(y_i > b_i)$ with respect to the control input sequence $\mathbb{U}$. The chain rule gives:

$$\nabla_{\mathbb{U}} P(y_i > b_i) = \frac{\partial P(y_i > b_i)}{\partial \bar{y}_i} \nabla_{\mathbb{U}} \bar{y}_i. \tag{15}$$

The Leibniz integral rule gives:

$$\begin{aligned} \frac{\partial P(y_i > b_i)}{\partial \bar{y}_i} &= \frac{\partial}{\partial \bar{y}_i} \frac{1}{\sqrt{2\pi}} \int_{\frac{b_i - \bar{y}_i}{\sqrt{S_{y_i}}}}^{\infty} e^{-\frac{z^2}{2}} dz \\ &= \frac{1}{\sqrt{2\pi S_{y_i}}} e^{-\frac{(b_i - \bar{y}_i)^2}{2S_{y_i}}}, \end{aligned} \tag{16}$$

and from (11) we have:

$$\nabla_{\mathbb{U}} \bar{y}_i = G_{xu}'\mathbf{a}_i. \tag{17}$$

Hence the gradient of each constraint in Problem 3 is:

$$\nabla_{\mathbb{U}} P(\mathbf{a}_i'\mathbb{X} > b_i) = \nabla_{\mathbb{U}} P(y_i > b_i) = \frac{G_{xu}'\mathbf{a}_i}{\sqrt{2\pi S_{y_i}}} e^{-\frac{(b_i - \bar{y}_i)^2}{2S_{y_i}}}. \tag{18}$$

Note that this can be evaluated exactly without the need for sampling *or* table lookups. This is possible because Problem 3 involves only singlevariate constraints.

D. Convexity

We now prove that Problem 3 is a convex optimization problem. First we restate Problem 3 using (13) to express the probabilities in integral form:

Problem 4 (Conservative Problem - Integral Form).

$$\begin{array}{ll} \textit{Minimize} & f(\bar{\mathbb{X}}, \mathbb{U}) \\ \textit{Subject to:} & \mathbb{U} \in F_U \\ & 1 - \textit{cdf}\Big(\frac{b_i - \bar{y}_i}{\sqrt{2 S_{y_i}}}\Big) \leq \epsilon_i \ \ \forall i \\ & \bar{y}_i = \mathbf{a}_i' G_{xx} \bar{\mathbf{x}}_0 + \mathbf{a}_i' G_{xu} \mathbb{U} + \mathbf{a}_i' G_{xw} \bar{\mathbb{W}} \ \ \forall i \\ & S_{y_i} = \mathbf{a}_i' G_{xx} S_{\mathbf{x}_0} G_{xx}' \mathbf{a}_i + \mathbf{a}_i' G_{xw} S_{\mathbb{W}} G_{xw}' \mathbf{a}_i \ \ \forall i \\ & \sum_i^N \epsilon_i \leq \delta. \end{array}$$

Lemma 2. *$\mathit{cdf}(x)$ is a concave function of x in the range $x \in [0, \infty]$. Hence for $\lambda \in [0,1]$, if:*

$$x^{(*)} = \lambda x^{(1)} + (1-\lambda) x^{(2)} \tag{19}$$

then:

$$\mathit{cdf}(x^{(*)}) \geq \lambda \mathit{cdf}(x^{(1)}) + (1-\lambda)\mathit{cdf}(x^{(2)}). \tag{20}$$

Proof: Following the proof in [12], concavity comes from the fact that $\mathrm{cdf}(\cdot)$ is the integral of a function that is monotonically decreasing in the range $[0, \infty]$. □

Lemma 3. *For $b_i \geq \bar{y}_i$ the constraint:*

$$1 - \mathit{cdf}\left(\frac{b_i - \bar{y}_i}{\sqrt{2 S_{y_i}}}\right) \leq \epsilon_i \tag{21}$$

is convex in $(\bar{y}_i, \epsilon_i)$.

Proof: Consider two solutions $(\bar{y}_i^{(1)}, \epsilon_i^{(1)})$ and $(\bar{y}_i^{(2)}, \epsilon_i^{(2)})$ that satisfy (21). In order to show convexity, we must show that (21) is also satisfied by $(\bar{y}_i^{(*)}, \epsilon_i^{(*)})$, where $\lambda \in [0,1]$:

$$(\bar{y}_i^{(*)}, \epsilon_i^{(*)}) \triangleq \Big(\lambda \bar{y}_i^{(1)} + (1-\lambda)\bar{y}_i^{(2)}, \lambda \epsilon_i^{(1)} + (1-\lambda)\epsilon_i^{(2)}\Big). \tag{22}$$

From Lemma 2 we have:

$$\begin{aligned} & 1 - \mathrm{cdf}\left(\frac{b_i - \bar{y}_i^{(*)}}{\sqrt{2 S_{y_i}}}\right) \\ & \leq 1 - \lambda \mathrm{cdf}\left(\frac{b_i - \bar{y}_i^{(1)}}{\sqrt{2 S_{y_i}}}\right) - (1-\lambda)\mathrm{cdf}\left(\frac{b_i - \bar{y}_i^{(2)}}{\sqrt{2 S_{y_i}}}\right) \\ & \leq 1 - \lambda(1 - \epsilon_i^{(1)}) - (1-\lambda)(1-\epsilon_i^{(2)}) = \epsilon_i^{(*)}, \end{aligned} \tag{23}$$

where the final inequality comes from knowing that $(\bar{y}_i^{(1)}, \epsilon_i^{(1)})$ and $(\bar{y}_i^{(2)}, \epsilon_i^{(2)})$ satisfy (21). Hence (21) is satisfied by $(\bar{y}_i^{(*)}, \epsilon_i^{(*)})$, which proves the convexity of (21). □

Lemma 4. *For $\delta \leq 0.5$, Problem 4 is convex, and so is Problem 3 (the conservative problem).*

Proof: To show convexity of Problem 4 it suffices to show that all constraints are convex, since we have assumed that the cost function $f(\cdot)$ is convex. The control constraint $\mathbb{U} \in F_{\mathbb{U}}$ is convex since we have assumed the feasible control set $F_{\mathbb{U}}$ to be convex. All other constraints are linear, and hence convex, except for (21). For $\delta \leq 0.5$, we know that a feasible solution has $\epsilon_i \leq 0.5$ for all i since $\sum_i^N \epsilon_i \leq \delta$. By the definition of $\mathrm{cdf}(\cdot)$, for (21) to be satisfied with $\epsilon_i \leq 0.5$, we must have $b_i \geq \bar{y}_i$, in which case (21) is convex by Lemma 3. Hence Problem 4 is convex. Since Problem 4 is equivalent to Problem 3, Problem 3 is also convex. □

E. Summary

We have proposed a new conservative approximation of the chance-constrained control problem. A feasible solution to this approximation is a feasible solution to the original problem. The conservative problem is convex, meaning that existing nonlinear solvers can be used to find the globally optimal solution in practice. Furthermore, the constraint values and derivatives needed to perform this optimization can be computed without the need for sampling. While the approach is conservative, we show empirically in Section VI that the conservatism introduced is small.

V. Comparison with Set Conversion Techniques

Alternative conservative approximations of the chance constrained problem (Problem 1) have been proposed previously, for example [4]. One particularly computationally tractable approach is to convert all stochastic distributions into sets. That is, we define, before optimization begins, a set $G(\mathbb{U})$ such that the following condition holds:

$$P\big(\mathbb{X} \notin G(\mathbb{U})\big) \leq \delta. \tag{24}$$

We can then use algorithms for robust control under *set-bounded* uncertainty to ensure that $G(\mathbb{U}) \subseteq F_x$, for example [13]. This ensures that the required chance constraints are satisfied:

$$\big\{G(\mathbb{U}) \subseteq F_X\big\} \wedge \big\{P\big(\mathbb{X} \notin G(\mathbb{U})\big) \leq \delta\big\} \Rightarrow P\big(\mathbb{X} \notin F_X\big) \leq \delta. \tag{25}$$

With Gaussian uncertainty, we typically choose the set $G(\mathbb{U})$ to be ellipsoidal with principal axes aligned with the covariance of $\mathbb{X}$. This leads to a set-bounded problem where we must optimize the location of the center of the ellipsoid subject to the ellipsoid lying within F_X. The choice of ellipsoidal $G(\mathbb{U})$ leads to tractable determination of the smallest ellipsoid size satisfying (24) as well as tractable methods for ensuring $G(\mathbb{U}) \subseteq F_x$; see for example [4], for details. We now discuss how this set conversion approach relates to our new convex optimization approach. In Section VI we provide an empirical comparison.

Figure 1 shows, schematically, a two-dimensional Gaussian distribution and an ellipsoid containing exactly 99% of the probability density. In most chance-constrained problems of interest, the chance constraint is tight in the optimal solution. We can measure the conservatism of a particular approach by the difference between the value $P(\mathbb{X} \notin F_X)$ from δ in the returned solution. With set conversion techniques, the only way that $P(\mathbb{X} \notin F_X)$ can be close to δ is if the feasible region approximates the set $G(\mathbb{U})$, as shown in Figure 1. This is, however, extremely unlikely in the general case. In most optimization problems of interest, the feasible region will be significantly larger than $G(\mathbb{U})$ and of a different geometry, as in Figure 2. Observe that, using a set conversion approach in this case, $P(\mathbb{X} \notin F_X)$ is far below the constraint δ, indicating a great deal of conservatism. As the dimensionality of the distribution increases, the level of conservatism increases dramatically.

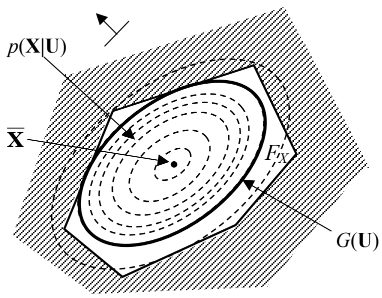

Fig. 1. Chance constrained optimization problem approximated using ellipsoidal set conversion. Shown is a two-dimensional Gaussian distribution for $\mathbb{X}$, represented using contours of the pdf (dashed). The thick ellipse is the set $G(\mathbb{U})$ containing 99% of the probability mass. In this case, the feasible region F_X is almost identical in geometry to $G(\mathbb{U})$. The cost is defined so that the mean $\bar{\mathbb{X}}$ moves as far as possible in the direction of the arrow. With this particular geometry $P(\mathbb{X} \notin F_X)$ can be close to the constraint δ because the integral of the pdf over F_X is close to the integral over $G(\mathbb{U})$.

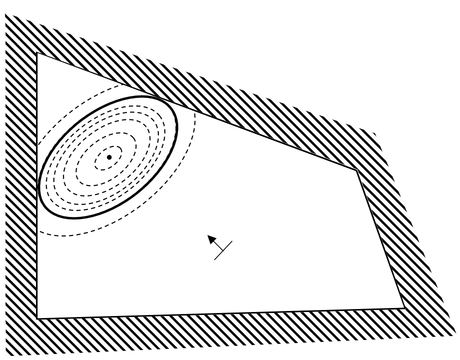

Fig. 2. Chance constrained optimization problem with general feasible region (geometry dissimilar to $G(\mathbb{U})$). In the optimal solution returned by the set conversion method, the set $G(\mathbb{U})$ is constrained to lie within F_X. Observe that large portions of the pdf are outside $G(\mathbb{U})$ but within F_X, meaning that $P(\mathbb{X} \notin F_X)$ is significantly below δ in this solution. Hence the solution is conservative. This conservatism increases as the dimensionality of the distribution increases.

Intuitively, a set bounding approach assumes a 'worst-case' scenario, where all constraints contribute equally to the overall probability of failure. In most finite-horizon control problems, only a small subset of the constraints are active in the optimal solution. We claim, then, that a set conversion approach leads to high conservatism. By contrast, the new approach described in this paper optimizes the violation probabilities assigned to each constraint, denoted ϵ_i. In doing so it can greatly reduce the conservatism of the solution, as illustrated in Figure 3. This 'risk allocation' concept was proposed first by [11].

The new approach does introduce conservatism in the use of Boole's inequality. However we show empirically

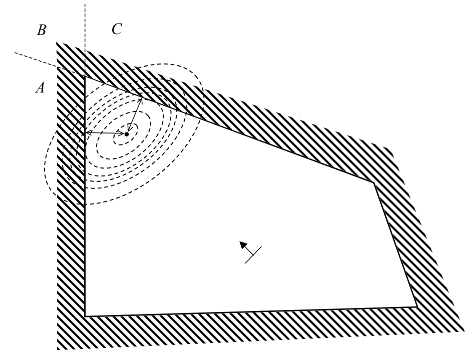

Fig. 3. Chance constrained optimization using new convex optimization approach. The approach optimally allocates the contribution to the failure probability from each of the constraints. In the solution shown, two of the constraints have approximately zero probability of violation, so the algorithm pushes the mean closer to the upper-left corner until the probability of violation of the constraints sums to 0.01. The conservatism introduced by Boole's inequality is the integral of the pdf over region B, which is small in practice.

in Section VI that this conservatism is small, while the set conversion approach is conservative by many orders of magnitude.

VI. Simulation Results

In this section we show simulation results demonstrating the new approach. The system to be controlled has state $\mathbf{x}_k = [x_k' \; y_k']'$ and the system parameters are defined by:

$$A = \begin{bmatrix} 1 & 1 \\ -0.5 & 0 \end{bmatrix} \; B = \begin{bmatrix} 0 \\ 0.0\dot{3} \end{bmatrix} \; B^w = C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \; D = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{26}$$

The noise parameters are:

$$S_{\mathbf{x}_0} = \begin{bmatrix} 0.0001 & 0 \\ 0 & 0.0001 \end{bmatrix} \quad S_{w_k} = \begin{bmatrix} 0.001 & 0 \\ 0 & 0.001 \end{bmatrix} \;\; \forall k \tag{27}$$

$$\bar{\mathbf{x}}_0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \qquad \bar{\mathbf{w}}_k = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \;\; \forall k. \tag{28}$$

The constraints on the state are:

$$-0.25 \leq y_k \leq 0.25 \;\; \forall k. \tag{29}$$

These are encoded using $2(T+1)$ linear inequality constraints. The cost is defined as:

$$f(\bar{\mathbb{X}}, \mathbb{U}) = \sum_{k=0}^{T} (\bar{\mathbf{x}}_k - \mathbf{x}^r)'(\bar{\mathbf{x}}_k - \mathbf{x}^r). \tag{30}$$

In other words, we try to minimize the squared distance of the expected state from some reference state $\mathbf{x}^r$, averaged over the planning horizon. For the convex optimization we used Sequential Quadratic Programming, as implemented in the Matlab `fmincon` function. Optimization was performed on a MacBook Pro with a 2.4GHz processor and 4GB RAM.

Figure 4 shows a single solution to the predictive control problem using the new convex optimization approach. In this case, $\mathbf{x}^r = [1 \; 0]'$, $N = 20$ and $\delta = 0.01$. The new approach optimizes the allocation of risk at each time step, while ensuring that the probability of failure over the entire horizon is less than δ. As shown in Figure 6, the risk allocation values ϵ_i are tiny ($< 10^{-8}$) for all constraints except for 5 of the 42 constraints. This implies that optimizing risk allocation can lead to significant gains over a set conversion approach, which uses an *a priori* fixed backoff from the constraints.

For the sake of comparison, Figure 5 shows a solution to the same problem using the elliptical set conversion approach of [4]. Notice that the state means are very far from the constraints compared to the solution in Figure 4, indicating a great deal of conservatism. This is because the set conversion approach assumes a 'worst-case' allocation of risk to each of the constraints over the time horizon, rather than optimizing the risk allocation. To evaluate the conservatism we performed 10^6 Monte Carlo simulations and determined the empirical probability of constraint violation, which we refer to as the 'true' probability of failure P_{true}. We define the conservatism factor as $(\delta - P_{true})/P_{true}$. For this example, the new convex optimization approach gave a true probability of failure of 0.0086, and hence is conservative by a factor of 0.2. The elliptical set conversion approach [4] has a true probability of failure of less than 10^{-6}, and is hence conservative by a factor of over 10^4.

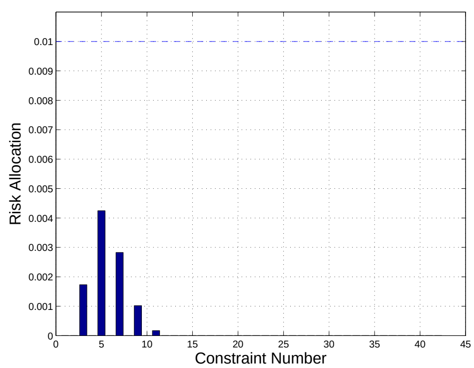

Fig. 6. Risk allocation in optimal solution of Figure 4. The bars show the value ϵ_i, i.e. the risk allocation, for each constraint in the problem. The allocated risks add to the maximum probability of failure for the entire horizon, δ, shown as the dashed line. Only a handful of the ϵ_i are non-negligible.

In order to assess the average performance of the new algorithm, we generated random instances of the control problem by setting $\mathbf{x}^r = [n\ 0]'$ with n uniformly distributed in the range $[0, 1]$. Again we used $N = 20$ and $\delta = 0.01$. The average results for 20 solutions are shown in Table I. Since we are interested in the conservatism of the new approach, we have removed instances where the chance constraints were not tight. In the globally optimal solution, we would expect the true probability of failure to be the same as δ. The results show that the new convex optimization approach is many orders of magnitude less conservative than the elliptical set conversion approach for a small penalty in solution time.

VII. Conclusion

We have proposed a new approach for chance-constrained predictive control that does not require the evaluation of multivariate probability densities. By using a conservative bounding approach we ensure that chance constraints are satisfied, and we have shown analytically that the resulting optimization is convex. This means that existing solvers can find the globally optimal solution efficiently. Empirical results showed that the approach is many orders of magnitude less conservative than existing set conversion techniques for a small penalty in computation time.

Algorithm	Time (s)	P_{true}	Conservatism
Convex Optimization	1.03	0.0079	0.27
Elliptical Set Conversion	0.22	$< 10^{-6}$	$> 10^4$

TABLE I. Optimization time and true probability of failure averaged for 20 randomized problem instances with $\delta = 0.01$. Instances where the chance constraint is not tight have been removed. The convex optimization approach is orders of magnitude less conservative than the set conversion approach for a small penalty in solution time.

VIII. Acknowledgements

The research described in this paper was carried out in part at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The author would also like to acknowledge Masahiro Ono, for discussions that made this paper possible.

References

[1] A. Schwarm and M. Nikolaou, "Chance-constrained model predictive control," *AlChE Journal*, vol. 45, pp. 1743–1752, 1999.

[2] P. Li, M. Wendt, and G. Wozny, "Robust model predictive control under chance constraints," *Computers and Chemical Engineering*, vol. 24, pp. 829–834, 2000.

[3] ——, "A probabilistically constrained model predictive controller," *Automatica*, vol. 38, pp. 1171–1176, 2002.

[4] D. van Hessem and O. H. Bosgra, "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 Control and Decision Conference*, 2003.

[5] Y. Yan and R. R. Bitmead, "Incorporating state estimation into model predictive control and its application to network traffic control," *Automatica*, vol. 41, pp. 595–604, 2005.

[6] I. Batina, "Model predictive control for stochastic systems by randomized algorithms," Ph.D. dissertation, Technische Universiteit Eindhoven, Eindhoven, The Netherlands, 2004.

[7] L. Blackmore, H. X. Li, and B. C. Williams, "A probabilistic approach to optimal robust path planning with obstacles," in *Proceedings of the American Control Conference*, 2006.

[8] L. Blackmore, A. Bektassov, M. Ono, and B. C. Williams, "Robust, optimal predictive control of jump markov linear systems using particles," in *Hybrid Systems: Computation and Control, HSCC*, ser. Lecture Notes in Computer Science, A. Bemporad, A. Bicchi, and G. Buttazzo, Eds. Springer Verlag, 2007, vol. 4416, pp. 104–117.

[9] J. A. Primbs, "Receding horizon control of constrained linear systems with state and control multiplicative noise," in *Proceedings of the American Control Conference*, 2007.

[10] P. Kall and S. W. Wallace, *Stochastic Programming*. Wiley, 1994.

[11] M. Ono and B. C. Williams, "Efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure," MIT CSAIL, Tech. Rep., Mar. 2008. [Online]. Available: http://dspace.mit.edu/handle/1721.1/40803

[12] ——, "Two-stage optimization approach to robust model predictive control with a joint chance constraint," MIT CSAIL, Tech. Rep., Mar. 2008. [Online]. Available: http://dspace.mit.edu/handle/1721.1/40804

[13] A. Richards, "Robust constrained model predictive control," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2005.

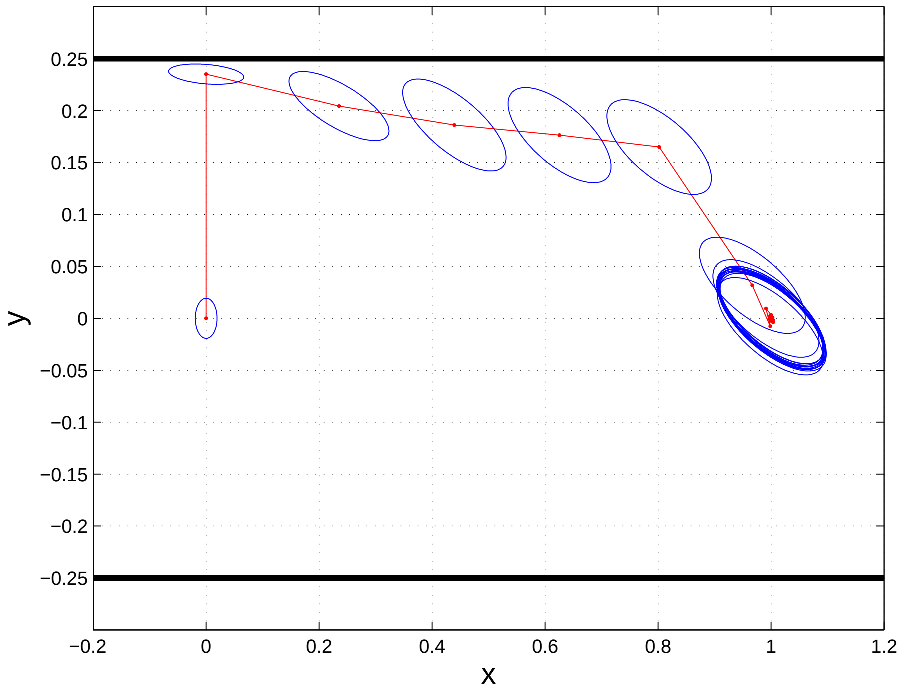

Fig. 4. Single solution using new convex optimization approach for $\mathbf{x}^r = [1\ 0]'$, $N = 20$ and $\delta = 0.01$. The red dots show the state mean $\bar{\mathbf{x}}_k$ for $k = 0, \ldots, N$. The blue ellipses show the covariance (1-sigma) ellipses for $\mathbf{x}_k$. The state constraints are shown as thick black lines. The new approach optimizes the allocation of risk at each time step, while ensuring that the probability of failure over the entire horizon is less than δ.

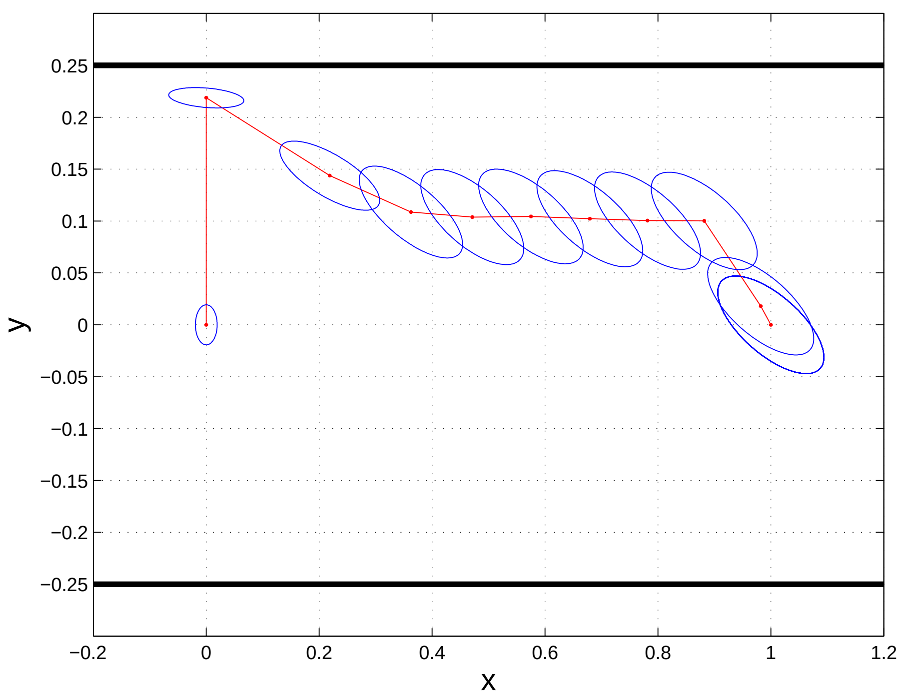

Fig. 5. Single solution for $\mathbf{x}^r = [1\ 0]'$, $N = 20$ and $\delta = 0.01$, using elliptical set conversion approach of [4]. The state means are very far from the constraints compared to the solution in Figure 4, indicating a great deal of conservatism.