A Hybrid Method for Chance Constrained Control in Uncertain Environments

Michael P. Vitus
Electrical Engineering and Computer Sciences
University of California at Berkeley
Email: vitus@eecs.berkeley.edu

Claire J. Tomlin
Electrical Engineering and Computer Sciences
University of California at Berkeley
Email: tomlin@eecs.berkeley.edu

Abstract— **This paper introduces a novel hybrid method for solving a stochastic control problem for linear, Gaussian systems through uncertain environments. Due to the imperfect knowledge of the system state caused by motion, sensor and environment uncertainty, the system constraints cannot be guaranteed to be satisfied and consequently must be considered probabilistically. Due to the environmental uncertainty, the constraints are sums of products of random variables which do not have a closed-form analytical expression. Previous approaches have either approximated the distribution leading to a nonconvex optimization program, or used sampling alone to represent the uncertainty which requires a large number of samples to accurately represent the distribution. To address these limitations, a novel hybrid method is proposed that uses both analytical functions and sampling to represent the uncertainty. It is shown that under certain conditions, the resulting optimization program is convex. Also, this method drastically reduces the computational complexity over previous methods, which is demonstrated through an example.**

I. Introduction

For many systems, it is desired to determine a control strategy to complete an objective while satisfying a set of constraints, however, in many circumstances there is noise in the system that can cause the constraints to be violated. Furthermore, the constraint parameters could also be uncertain which significantly increases the complexity in determining this control strategy. Therefore, the problem of interest in this paper is how to efficiently solve the stochastic control problem with uncertain constraint parameters.

There are many applications that could benefit from this capability including robot-assisted surgery, energy efficient control of buildings, chemical process control, financial engineering, autonomous control of vehicles, robotic assistance for elderly and disabled people, routing aircraft around weather, and home automation for tasks.

The uncertainty in the previous applications arise from three different sources: (i) process uncertainty, (ii) sensing uncertainty and (iii) environment uncertainty. The presence of these uncertainties means that the exact system state is never truly known. Consequently, in order to maximize the probability of success, the control problem must be performed in the space of probability distributions of the system, defined as the belief space. For a stochastic system, however, planning in the belief space is not enough to guarantee success because there is always a small probability that a large disturbance will be experienced causing the system to violate the constraints. Therefore, a trade-off must be made between the conservativeness of the solution and the performance of the system.

Charnes, Miller and Wagner [1] introduced the problem of *chance constrained programming* which only guarantees constraint satisfaction up to a specified probability limit. A thorough account of existing literature employing this problem formulation is given in [2].

A group of researchers used the chance constrained programming formulation to model the planning problem as an optimal control problem. The work in [3] utilized Boole's inequality to bound the chance constraints, typically resulting in a very small amount of over-conservativeness. They also used the idea of risk allocation introduced by [4] to distribute the risk of violating each chance constraint while still guaranteeing safety. The work in [5] optimized over the feedback control laws and open-loop inputs while ensuring that the chance constraints on the overall system were satisfied. He used an ellipsoidal relaxation technique to convert the stochastic problem into a deterministic one, but this leads to a conservative solution. Vitus and Tomlin [6] reduced the conservatism by using Boole's inequality and proposed an efficient iterative two stage optimization scheme.

Incorporating environment uncertainty into the motion planning problem has also received some attention. Vitus and Tomlin [7] studied the problem of belief space planning for linear, Gaussian systems in uncertain environments. They formulated the problem as a chance constrained optimization problem, and showed the probability of colliding with polygonal obstacles can be accurately approximated by a Gaussian distribution; however, this leads to a nonconvex optimization program. Du Toit and Burdick [8] investigated obstacle avoidance in dynamic, uncertain environments, and proposed an approximation method for calculating the probability of collision for a spherical robot and obstacle distributed via a Gaussian distribution. Missiuro and Roy [9] handled uncertain environments by modifying the sampler used in a probabilistic roadmap. However, while the algorithm accounts for the environment uncertainty, motion noise or sensing noise is not accounted for.

This work extends previous chance constrained programming formulations to solve the stochastic control problem in uncertain environments. Given the uncertainty in the constraint parameters, the probabilistic constraints on the system state are shown to be distributed via the sums of products of random variables. In general the constraint expression does not have a closed-form analytical expression and requires the evaluation of multivariate integrals. To reduce the

This work has been supported in part by NSF under CPS:ActionWebs (CNS-931843), by ONR under the HUNT (N0014-08-0696) and SMARTS (N00014-09-1-1051) MURIs and by grant N00014-12-1-0609, by AFOSR under the CHASE MURI (FA9550-10-1-0567).

computational complexity over previous approaches, a novel hybrid method is proposed that uses both analytical functions and sampling to represent the uncertainty in the system and environment. It is shown that under certain conditions, the resulting optimization program is convex. By using this dual representation, the computational complexity was drastically reduced by 30-110 times over previous methods. Consequently, this method has the ability to enable real-time stochastic control for the motivating applications.

The paper proceeds as follows. Section II describes the probabilistic problem formulation. Then, two previous methods for handling the chance constraints are presented Section III. In Section IV, the hybrid method is presented and an example is presented in Section V which characterizes the performance of the algorithm.

II. Problem Formulation

Consider the following linear stochastic system defined by,

$$x_{k+1} = Ax_k + Bu_k + w_k, \quad k = 0, \dots N-1, \tag{1}$$

where $x_k \in \mathbb{R}^n$ is the system state, $w_k \in \mathbb{R}^n$ is the process noise and N is the time horizon. The initial state, x_0, is assumed to be a Gaussian random variable with mean $\bar{x}_0$ and covariance Σ_0 i.e., $x_0 \sim \mathcal{N}(\bar{x}_0, \Sigma_0)$. At each time step, a noisy measurement of the state is taken, defined by

$$y_k = Cx_k + v_k, \quad k = 1, \dots N, \tag{2}$$

where $y_k \in \mathbb{R}^p$ and $v_k \in \mathbb{R}^p$ are the measurement output and noise of the sensor at time k, respectively. The process and measurement noise have zero mean Gaussian distributions, $w_k \sim \mathcal{N}(0, \Sigma_w)$ and $v_k \sim \mathcal{N}(0, \Sigma_v)$. The process noise, measurement noise and initial state are assumed to be mutually independent. For notational convenience, the state and control inputs for all time-steps are concatenated to form, $\mathbb{X} = \begin{bmatrix} x_1^{\mathrm{T}} & \dots & x_N^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$ and $\mathbb{U} = \begin{bmatrix} u_0^{\mathrm{T}} & \dots & u_{N-1}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$, and $\bar{\mathbb{X}}$, $\bar{\mathbb{U}}$ refer to the mean of the state and control inputs.

The control inputs are required to be in a convex region denoted by F_U and the system state is restricted to be in a feasible region denoted by F_X. To simplify the presentation of the material, the feasible region F_X is assumed to be convex. Nonconvex regions can still be handled, however, either by (i) performing branch and bound on the set of conjunction and disjunction linear state constraints directly [10], or by (ii) decomposing the space into convex regions and using branch and bound to determine when to enter/exit each convex subregion [11]. Given this assumption, the feasible region can be defined by a conjunction of N_{F_X} linear inequalities,

$$F_X \triangleq \bigcap_{i=1}^{N_{F_X}} \left\{ \mathbb{X} : h_i^{\mathrm{T}} \mathbb{X} \le b_i \right\} \tag{3}$$

where $h_i \in \mathbb{R}^{nN}$ and $b_i \in \mathbb{R}$. In this work, the environment is uncertain but the parameters of the probability distribution describing h_i and b_i are assumed to be known.

In this formulation, the design variables are the control inputs for the system. They can be specified directly as the open-loop values or generated through a linear trajectory controller with the desired trajectory as the design variable. In either case, the distribution of the closed-loop system can be calculated *a priori* and is given by a Gaussian distribution,

$$\mathbb{X} \sim \mathcal{N}\left(\bar{\mathbb{X}}, \Sigma_{\mathbb{X}}\right). \tag{4}$$

The general belief space planning problem is posed as the optimization program (5). The optimization variables are the control inputs for the system. The objective function, $f(\cdot)$, is assumed to be a convex function in $\mathbb{X}$ and $\mathbb{U}$.

$$\begin{array}{ll} \text{minimize} & \mathbf{E}\left[f(\mathbb{X}, \mathbb{U})\right] \\ \text{subject to} & \\ & x_{k+1} = Ax_k + Bu_k + w_k, \quad k = 0, \dots N-1 \\ & y_k = Cx_k + v_k, \quad k = 1, \dots N \\ & w_k \sim \mathcal{N}(0, \Sigma_w), \quad k = 0, \dots N-1 \\ & v_k \sim \mathcal{N}(0, \Sigma_v), \quad k = 1, \dots N \\ & \bar{\mathbb{U}} \in F_U \\ & \mathrm{P}(\mathbb{X} \notin F_X) \le \delta \end{array} \tag{5}$$

The difficulty in solving the optimization program (5) is in evaluating and satisfying the joint chance constraints: $\mathrm{P}(\mathbb{X} \notin F_X) \le \delta$. The complexity arises from evaluating the multivariate integrals over the uncertain state and environment to calculate the desired probability of failure. In particular, by allowing H and b to be uncertain, the distribution of $H^{\mathrm{T}}\mathbb{X} - b$ becomes a sum of multiple products of random variables which in general results in a nonconvex optimization program. In addition, the properties of the probability distribution for this constraint are not easy to calculate analytically, increasing the complexity of the problem. The following two sections describe three approaches to deal with such uncertain constraint parameters.

III. Chance Constraints

The primary focus of most previous chance constrained problem formulations was handling uncertainty in the system. In this work the constraint parameters are also uncertain which drastically increases the complexity. In the first method, the constraints are approximated by a Gaussian distribution, but this leads to a nonconvex optimization problem which does not guarantee the optimal solution. The second approach upper bounds the original chance constraint by a convex function and uses sampling to calculate the value of the bound. While this approach is useful for handling arbitrary probability distributions, it usually requires a large number of samples resulting in a large computational complexity. These two methods will be further developed in the following subsections.

A. Gaussian Approximation using Boole's Inequality

In [7], we proposed to use Boole's inequality to bound the multivariate chance constraint by the summation of a set of univariate constraints. Boole's inequality states that for a countable set of events $E_1, E_2, \dots$, the probability that at least one event happens is no larger than the sum of the individual probabilities $\mathrm{P}\left(\bigcup_i E_i\right) \le \sum_i \mathrm{P}\left(E_i\right)$. Consequently, from Eqn. (3) and Boole's inequality the probability of the state not being contained inside the feasible region is bounded by,

$$\begin{aligned} \mathrm{P}(\mathbb{X} \notin F_X) &= \mathrm{P}\left(\mathbb{X} \in \bigcup_{i=1}^{N_{F_X}} \left\{\mathbb{X} : h_i^{\mathrm{T}}\mathbb{X} > b_i\right\}\right) \\ &\le \textstyle\sum_{i=1}^{N_{F_X}} \mathrm{P}(h_i^{\mathrm{T}}\mathbb{X} > b_i). \end{aligned} \tag{6}$$

By allowing h_i and b_i to be uncertain, the distribution of $\sum_{j=1}^{nN} h_{ij}\mathbb{X}_j - b_i$ is now a sum of multiple products of random variables, which does not have an analytical expression. Fortunately, the work of [7] showed the true distribution can be accurately approximated by a Gaussian distribution to allow the efficient evaluation of the constraints,

$$\begin{aligned} \mathrm{P}(h_i^{\mathrm{T}}\mathbb{X} - b_i > 0) &= \frac{1}{\sqrt{2\pi}} \int_{\frac{\bar{b}_i - \bar{h}_i^{\mathrm{T}}\bar{\mathbb{X}}}{\sigma_i}}^{\infty} \exp\left(-\frac{z^2}{2}\right) dz \\ &= 1 - \Phi\left(\frac{\bar{b}_i - \bar{h}_i^{\mathrm{T}}\bar{\mathbb{X}}}{\sigma_i}\right) \end{aligned} \tag{7}$$

where $\Phi(\cdot)$ is the Gaussian cumulative distribution function and $\sigma_i^2 = \mathrm{var}(h_i^{\mathrm{T}}\mathbb{X} - b_i)$. Using a risk allocation technique for each of the univariate constraints ($\mathrm{P}(h_i^{\mathrm{T}}\mathbb{X} > b_i)$) yields the final set of constraints for this method:

$$\begin{gathered} \frac{1}{N_s} \sum_{j=1}^{N_s} \left(1 - \Phi\left(\frac{\bar{b}_i - \bar{h}_i^{\mathrm{T}}\bar{\mathbb{X}}}{\sigma_i}\right)\right) \leq \epsilon_i, \forall i \\ \sum \epsilon_i \leq \delta. \end{gathered} \tag{8}$$

Unfortunately, for this approach the resulting optimization program is nonconvex [7] and therefore only a locally optimal solution can be guaranteed. Also, this method has a large computational complexity which will prevent applying this technique in real-time control applications.

B. Convex Bounding Method

Another approach to handling the uncertainty in the constraint parameters is the convex bounding method developed by [12], [13]. Since the probability distribution of the chance constraints may not be a convex function, it is difficult to include them in the optimization program. This method finds a suitable conservative, convex approximation for the probability distribution of the chance constraints which results in a convex optimization program.

Consider a single individual chance constraint of the form

$$\mathrm{P}\left(h_i^{\mathrm{T}}\mathbb{X} - b_i > 0\right) \leq \delta_i. \tag{9}$$

The probability in Eqn. (9) can be calculated via

$$\mathrm{P}\left(h_i^{\mathrm{T}}\mathbb{X} - b_i > 0\right) = \mathrm{E}\left[\mathbf{1}\left(h_i^{\mathrm{T}}\mathbb{X} - b_i\right)\right] \tag{10}$$

where $\mathbf{1}(\cdot)$ is the indicator function defined as

$$\mathbf{1}(z) = \begin{cases} 1, & \text{if } z > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{11}$$

Since the indicator function $\mathbf{1}(z)$ is a nonconvex function, this greatly complicates the evaluation of the chance constraints in the optimization problem (5). However, the intuition behind this method is that by bounding the indicator function by a convex function the optimization program simplifies to a convex program.

Suppose such a nonnegative, nondecreasing, convex function $\psi : \mathbb{R} \to \mathbb{R}$ can be found such that for any $\alpha > 0$, $\psi(z/\alpha) \geq \mathbf{1}(z)$ for all z then

$$\mathrm{E}\left[\psi\left(\left(h_i^{\mathrm{T}}\mathbb{X} - b_i\right)/\alpha\right)\right] \geq \mathrm{P}\left(h_i^{\mathrm{T}}\mathbb{X} - b_i > 0\right). \tag{12}$$

Consequently, if the following convex constraint is satisfied, then the original chance constraint in Eqn. (9) is guaranteed to hold:

$$\mathrm{E}\left[\psi\left(\left(h_i^{\mathrm{T}}\mathbb{X} - b_i\right)/\alpha\right)\right] \leq \delta_i. \tag{13}$$

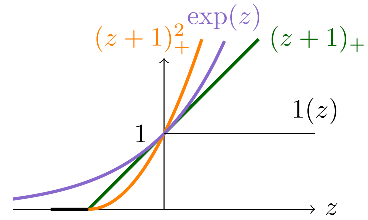

Fig. 1. A comparison of several convex bounds that can be used to approximate the indicator function $\mathbf{1}(z)$. The functions $(z+1)_+$, $\exp(z)$, and $(z+1)_+^2$ (where $+$ denotes $\max\{z+1, 0\}$) are represented as the green, purple and orange lines respectively.

Note, the constraint in Eqn. (13) holds for any α, and the conservativeness can be reduced by including α as an optimization variable. To handle joint chance constraints, the maximum over all constraint violations can be used in the convex bounding function, i.e.,

$$\mathrm{E}\left[\psi\left(\max\left(H^{\mathrm{T}}\mathbb{X} - b\right)/\alpha\right)\right] \leq \delta. \tag{14}$$

Now that the form of the convex constraint used to bound the original chance constraint has been formulated, the next step is to determine what form of function to use for ψ. The restrictions on the function, as stated previously, are that it needs to be a convex function and $\psi(z) \geq \mathbf{1}(z)$ for all z. Several examples of possible functions are shown in Figure 1.

In this work, the functional form for $\psi(z)$ used is $\psi(z) = (z+1)_+$ (where the subscript $+$ denotes $\max\{z+1, 0\}$) because it results in the least conservative bound. The final convex constraint that bounds the original chance constraint is then:

$$\mathrm{E}\left[\left(\max\left(H^{\mathrm{T}}\mathbb{X} - b\right) + \alpha\right)_+\right] \leq \alpha\delta. \tag{15}$$

Unfortunately, there is no analytical, closed-form solution to calculate the expectation in Eqn. (15), however, it can be efficiently approximated through sampling. One advantage of this is that sampling can represent arbitrary distributions for the uncertainty in the system, which inherently handles the increased complexity of having uncertain constraint parameters.

To evaluate the expectation, N_s particles are drawn at each time-step from the process noise, measurement noise, initial state, and uncertainty constraint parameters to obtain the sets

$$\begin{gathered} \{w_0^{(1)}, \ldots, w_{N-1}^{(1)}, \ldots, w_0^{(N_s)}, \ldots, w_{N-1}^{(N_s)}\}, \\ \{v_0^{(1)}, \ldots, v_{N-1}^{(1)}, \ldots, v_0^{(N_s)}, \ldots, v_{N-1}^{(N_s)}\}, \\ \{x_0^{(1)}, \ldots, x_0^{(N_s)}\}, \\ \{H^{(1)}, \ldots, H^{(N_s)}\}, \\ \{b^{(1)}, \ldots, b^{(N_s)}\}. \end{gathered} \tag{16}$$

In order to use the convex bounding method, the expectation in Eqn. (15) is not only over the system state but also over the uncertain constraint parameters. To calculate the expectation, each system trajectory particle $\mathbb{X}^{(j)}$ is associated with a set of constraint parameters $H^{(j)}$ and $b^{(j)}$ for all $j = 1, \ldots, N_s$ as follows

$$\begin{aligned} &\mathrm{E}\left[\left(\max\left(H^{\mathrm{T}}\mathbb{X} - b\right) + \alpha\right)_+\right] \approx \\ &\quad \frac{1}{N_s} \sum_{j=1}^{N_s} \left(\max\left({H^{(j)}}^{\mathrm{T}}\mathbb{X}^{(j)} - b^{(j)}\right) + \alpha\right)_+ . \end{aligned} \tag{17}$$

While this approach inherently handles the uncertainty in the system and constraint parameters, for some problems it might result in a large optimization program due to having to use an increased number of samples to accurately represent the underlying distribution. This will be illustrated in the following subsection.

C. Stochastic Motion Planning Example

The following example will motivate the need to develop a new method for handling problems with environmental uncertainty. The system has double integrator dynamics with $\Delta t = 0.1$ seconds and a time-horizon of $N = 20$. The state is composed of the positions followed by the velocities, and the measurement is of the position. where $\Delta t = 0.1$ seconds and a time-horizon of $N = 20$. The noise parameters are $\Sigma_w = \mathrm{diag}(0.0003, 0.0005, 0.0003, 0.0005)$ and $\Sigma_v = \mathrm{diag}(0.001, 0.002)$, where diag places the elements along the diagonal. The allowed probability constraint violation is $\delta = 0.005$. The objective function for this problem is quadratic in the final state as well as the control inputs,

$$f(\bar{\mathbb{X}}, \bar{\mathbb{U}}) = (x_N - x_{\text{ref}})^{\mathrm{T}} Q_{\text{obj}} (x_N - x_{\text{ref}}) + \bar{\mathbb{U}}^{\mathrm{T}} R_{\text{obj}} \bar{\mathbb{U}} \quad (18)$$

with $Q_{\text{obj}} = 50I$, $R_{\text{obj}} = 0.001I$ and $x_{\text{ref}} = \begin{bmatrix} 2 & 1 & 0 & 0 \end{bmatrix}^{\mathrm{T}}$.

For this example, the environment is defined by a set of half plane constraints defined by a series of end points, which are assumed to be uncertain. The uncertainty is modeled by a truncated Gaussian represented as the orange ellipses in Figure 2.

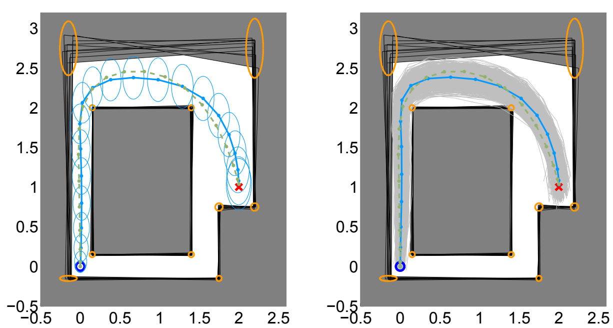

Fig. 2. Results for the Gaussian approximation method and the convex bounding method with ten realizations of the uncertain walls. The white area is the feasible region of the system state and the orange ellipses are the uncertainty of the environment. The blue, solid line is the solution when accounting for the uncertainty of the environment and the green, dotted line plans through the mean environment. The blue ellipses around the path indicate the uncertainty of the system. The start and goal location are marked by an 'o' and 'x', respectively. (a) The Gaussian approximation solution using Boole's inequality. (b) The solution using the convex bounding method with 1500 particles and their trajectories are shown as the gray lines.

The solution for this example is shown in Figure 2(a)-(b) for Boole's method and the convex bounding method, respectively. The blue, solid line is the trajectory of the system when accounting for the uncertainty of the environment, and the green, dotted line is the solution when planning only through the mean environment. The blue ellipses show the 99.7% confidence ellipsoid of the system state at each time-step. For this example, the optimal path is always through the top region, even though the bottom region is shorter. The bottom region is infeasible with respect to the allowed constraint violation because of the large uncertainty of the vertical position of the state.

There are several interesting differences between the solution that accounts for the uncertainty and the solution through the mean environment. The solution from the mean environment initially curves toward the wall with large uncertainty, but when the uncertainty of the environment is incorporated the system deviates away from it. The more noticeable difference between the two solutions in the top region of the environment. The solution which accounts for the uncertainty of the environment stays lower to avoid the highly uncertain top wall.

The Gaussian approximation and convex bounding methods can be compared by their conservativeness and their computational complexity. The estimated true probability of constraint violation using Monte Carlo simulation for the Gaussian approximation method and the convex bounding method is 0.0038 and 0.0047. The Gaussian approximation method is the most conservative due to the approximation from Boole's inequality. In particular, Boole's inequality doesn't account for the dependence between the state at different time-steps violating the constraints. Consider a wall whose uncertainty is purely translational. If at one time-step the system doesn't violate the constraint, then the system will also not violate the constraint for any future motion parallel to the wall. Using Boole's inequality to bound the probability ignores this dependency.

All computations were done using Matlab on a 2.7 GHz Intel Core i7. The Gaussian approximation method was solved using Matlab's `fmincon` with an interior point solver, and the convex bounding method was solved using `CVX` [14]. In terms of computational complexity, the Gaussian approximation solution takes 61.2 seconds and the convex bounding method uses 1500 particles and takes 286.1 seconds. Given the large computational complexity, neither of these methods have the potential of being applied for real-time control for the motivating applications.

The two approaches presented thus far for handling uncertain constraint parameters either approximate the constraints' probability distribution analytically or use sampling. These methods either result in a nonconvex optimization program or had a large computational complexity. To overcome these difficulties, the following section will present a novel hybrid method that results in a convex optimization program under certain conditions and drastically reduces the computational complexity over the prior approaches.

IV. Hybrid Method

Given the limitations of the previous approaches for handling the uncertain constraint parameters, this section presents a novel hybrid approach that uses a combination of analytical functions and sampling to represent the probability distributions. In particular, for some problem formulations, the stochastic variables are naturally separable into two sets: one set that can be accurately represented through an analytic function, and another set whose distribution is best represented through sampling. This section develops this hybrid method, discusses how to handle the joint chance constraints, and shows that the resulting optimization program is convex under certain conditions.

The intuition behind this method comes from examining the previous methods. Analytical approaches that use Boole's inequality have been very successful for solving problems involving only system uncertainty. Their advantage is the ability to efficiently calculate the chance constraint, but adding environmental uncertainty significantly increases the complexity due to the multiplicative constraints. For the convex bounding method, the use of sampling enables handling arbitrary probability distributions but it may require a large number of samples. By combining both approaches, the multiplicative constraints may be eliminated while reducing the required number of samples.

In the following development of the hybrid method, sampling is used to represent the uncertain constraints and analytical functions are used to represent the uncertainty of the system state. A similar approach can be used to employ the hybrid method for other separations of the uncertainty representation.

For the chosen uncertainty representation, only the constraint parameters are sampled, resulting in the following set of particles $\{H^{(1)}, \ldots, H^{(N_s)}\}$ and $\{b^{(1)}, \ldots, b^{(N_s)}\}$. The distribution of the closed-loop state, $\mathbb{X}$, is represented analytically with distribution given by Eqn. (4). Using this hybrid representation of the uncertainty, none of the current methods for enforcing the joint chance constraints, $\mathrm{P}(\mathbb{X} \notin F_X) \leq \delta$, are readily applicable.

Using Boole's inequality, the joint chance constraint can be upper bounded by,

$$\mathrm{P}(\mathbb{X} \notin F_X) \leq \sum_{i=1}^{N_{F_X}} \mathrm{P}(h_i^{\mathrm{T}}\mathbb{X} > b_i), \tag{19}$$

which simplifies the constraints to a set of univariate constraints but it does not have an analytical solution due to the products of random variables. However, by exploiting the dual representation of the uncertainty, the probability can be approximated by the sample average over the environment,

$$\mathrm{P}(h_i^{\mathrm{T}}\mathbb{X} > b_i) \approx \frac{1}{N_s}\sum_{j=1}^{N_s}\left(1 - \mathrm{P}({h_i^{(j)}}^{\mathrm{T}}\mathbb{X} > b_i^{(j)})\right) \tag{20}$$

Using the same risk allocation technique introduced in the Gaussian approximation method yields the following constraints that replace the original joint chance constraint with,

$$\begin{gathered}\frac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \mathrm{P}({h_i^{(j)}}^{\mathrm{T}}\mathbb{X} > b_i^{(j)})\right) \leq \epsilon_i,\ i = 1, \ldots, N_{F_X}, \\ \textstyle\sum_{i=1}^{N_{F_X}} \epsilon_i \leq \delta.\end{gathered} \tag{21}$$

Using the analytical function for the probability distribution of the closed-loop state in Eqn. (4) this simplifies even further

$$\begin{gathered}\frac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \Phi\left(\dfrac{b_i^{(j)} - {h_i^{(j)}}^{\mathrm{T}}\bar{\mathbb{X}}}{\sqrt{{h_i^{(j)}}^{\mathrm{T}}\Sigma_{\mathbb{X}}h_i^{(j)}}}\right)\right) \leq \epsilon_i, \forall i \\ \textstyle\sum_{i=1}^{N_{F_X}} \epsilon_i \leq \delta.\end{gathered} \tag{22}$$

Through the dual representation of the uncertainty the products of random variables have been eliminated, but whether this leads to a computational advantage still needs to be investigated. The following theorem determines when the set of constraints in Eqn. (22) form a set of convex constraints.

Theorem 1: The constraints in Eqn. (22) are convex if $b_i^{(j)} \geq {h_i^{(j)}}^{\mathrm{T}}\bar{\mathbb{X}}$ for all $j = 1, \ldots, N_s$ and $i = 1, \ldots, N_{F_X}$.

Proof: To show this, consider two feasible solutions $\left(\bar{\mathbb{X}}^{(1)}, \epsilon_i^{(1)}\right)$ and $\left(\bar{\mathbb{X}}^{(2)}, \epsilon_i^{(2)}\right)$. To prove the constraints are convex, it suffices to show that the convex combination of two feasible solutions is feasible. Define the convex combination for $0 \leq \theta \leq 1$ as

$$\left(\bar{\mathbb{X}}^{(*)}, \epsilon_i^{(*)}\right) \triangleq \left(\theta\bar{\mathbb{X}}^{(1)} + (1-\theta)\bar{\mathbb{X}}^{(2)}, \theta\epsilon_i^{(1)} + (1-\theta)\epsilon_i^{(2)}\right).$$

Since $b_i^{(j)} \geq {h_i^{(j)}}^{\mathrm{T}}\bar{\mathbb{X}}$, it restricts the Gaussian cumulative distribution function to its concave subset (i.e. $\Phi(z)$ is concave for $z \in [0, \infty)$). From the concavity, the lefthand side of Eqn. (22) for $\left(\bar{\mathbb{X}}^{(*)}, \epsilon_i^{(*)}\right)$ can be upper bounded by

$$\begin{aligned}&\tfrac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \Phi\left(\nu_{ij}^{(*)}\right)\right) \leq \\ &\quad\tfrac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \theta\Phi\left(\nu_{ij}^{(1)}\right) - (1-\theta)\Phi\left(\nu_{ij}^{(2)}\right)\right),\end{aligned} \tag{23}$$

where $\nu_{ij}^{\triangle} \in \mathbb{R}$ is defined as $\nu_{ij}^{\triangle} = \dfrac{b_i^{(j)} - {h_i^{(j)}}^{\mathrm{T}}\bar{\mathbb{X}}^{\triangle}}{\sqrt{{h_i^{(j)}}^{\mathrm{T}}\Sigma_{\mathbb{X}}h_i^{(j)}}}$ and $\triangle = \{(*), (1), (2)\}$. After simplifying and collecting common terms, the righthand side of Eqn. (23) is equal to

$$1 - \frac{1}{N_s}\sum_{j=1}^{N_s}\theta\Phi\left(\nu_{ij}^{(1)}\right) - \frac{1}{N_s}\sum_{j=1}^{N_s}(1-\theta)\Phi\left(\nu_{ij}^{(2)}\right). \tag{24}$$

From the definition of the two feasible solutions it is known that

$$-\frac{1}{N_s}\sum_{j=1}^{N_s}\Phi\left(\nu_{ij}^{(k)}\right) \leq \epsilon_i^{(k)} - 1 \quad k = 1, 2.$$

Substituting this into Eqns. (23), (24) and simplifying yields,

$$\begin{aligned}\tfrac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \Phi\left(\nu_{ij}^{(*)}\right)\right) &\leq 1 + \theta(\epsilon_i^{(1)} - 1) + \\ &\quad (1-\theta)(\epsilon_i^{(2)} - 1) \\ &= \theta\epsilon_i^{(1)} + (1-\theta)\epsilon_i^{(2)} \\ &= \epsilon_i^{(*)}.\end{aligned}$$

Hence the constraints in Eqn. (22) are convex. ■

The final optimization program for the hybrid approach is given in program (25).

$$\begin{aligned}\text{minimize}\quad & \mathbf{E}\left[f(\mathbb{X}, \mathbb{U})\right] \\ \text{subject to}\quad & \\ & x_{k+1} = Ax_k + Bu_k + w_k, \quad k = 0, \ldots N-1 \\ & y_k = Cx_k + v_k,, \quad k = 1, \ldots N \\ & w_k \sim \mathcal{N}(0, \Sigma_w), \quad k = 0, \ldots N-1 \\ & v_k \sim \mathcal{N}(0, \Sigma_v), \quad k = 1, \ldots N \\ & \bar{\mathbb{U}} \in F_U \\ & z_{ji} = b_i^{(j)} - {h_i^{(j)}}^{\mathrm{T}}\bar{\mathbb{X}},\ \forall i, j \\ & \sigma_{ji}^2 = {h_i^{(j)}}^{\mathrm{T}}\Sigma_{\mathbb{X}}h_i^{(j)},\ \forall i, j \\ & \tfrac{1}{N_s}\textstyle\sum_{j=1}^{N_s}\left(1 - \Phi\left(\frac{z_{ji}}{\sigma_{ji}}\right)\right) \leq \epsilon_i,\ \forall i \\ & \textstyle\sum_{i=1}^{q}\epsilon_i \leq \delta\end{aligned} \tag{25}$$

For a general uncertainty model of the constraints, there is no guarantee that $b_i^{(j)} \geq h_i^{(j)\mathrm{T}} \bar{\mathbb{X}}$. However, if the constraints $b_i^{(j)} - h_i^{(j)\mathrm{T}} \bar{\mathbb{X}} \geq 0 \; \forall i, j$ are added to the optimization program and none of the constraints is active, then the globally optimal solution can be guaranteed.

As will be illustrated in the following examples, the hybrid approach has several benefits over the Gaussian approximation approach and the convex bounding method. Since the hybrid approach is a convex program, it typically can be solved faster than the nonconvex program in the Gaussian approximation method. By using both sampling and analytical functions to represent the uncertainty, fewer particles are needed to fully represent the underlying uncertainty than using sampling alone. As compared to the convex bounding method, this drastically reduces the computational complexity of the problem formulation.

V. Results

The hybrid method using $N_s = 50$ samples was applied to the previous example and the results are shown in Figure 3. The estimated true probability of constraint violation using Monte Carlo simulation is 0.0042 which is in between the Gaussian approximation and the convex bounding method. Since the hybrid method uses Boole's inequality to bound the joint chance constraint, it is expected that it is more conservative than the convex bounding method. However, the hybrid method reduces the conservativeness over the Gaussian approximation method by sampling the constraint parameters, which accounts for the interdependence between violating the constraints at different time-steps.

The hybrid method optimization program was solved using an interior point solver with a Newton step. In terms of computational complexity, the hybrid method takes 2.42 seconds to compute the solution, which is a decrease of 30 times over the Gaussian approximation method and 110 times over the convex bounding method. Clearly, using the hybrid approach would enable stochastic control in real-time applications, whereas the other two approaches could only be used for offline calculations.

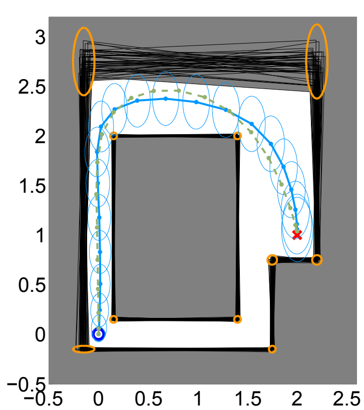

Fig. 3. Results for the hybrid method on the previous example. The solution using the hybrid method with $N_s = 50$ samples.

VI. Conclusion

The stochastic control problem in uncertain environments was formulated as a chance constrained optimization program. A novel hybrid method was proposed that uses a combination of analytical functions and sampling to represent the uncertainty of the system and environment. Under certain restrictions, the method results in a convex optimization program guaranteeing the globally optimal solution, Also, through this dual representation, the computational complexity was drastically reduced by $30 - 110$ times over previous methods which will enable stochastic control in real-time for the motivating applications.

There are several interesting areas of future work that the authors wish to explore. The authors wish to investigate conditions on the uncertainty distributions under which the hybrid method formulation is guaranteed to be convex. Another area of future work is to explore an iterative solution algorithm to remove particles that cause the hybrid method's constraints to be nonconvex while still maintaining a feasible solution. Lastly, the authors wish to apply these algorithms to actual systems to navigate through cluttered environments or to enable energy efficient control of buildings.

References

[1] A. Charnes, W. W. Cooper, and G. H. Symonds, "Cost horizons and certainty equivalents: An approach to stochastic programming of heating oil," *Management Science*, vol. 4, no. 3, pp. 235–263, 1958.

[2] A. Prékopa, *Stochastic Programming*. Dordrecht: Kluwer Academic Publishers, 1995.

[3] L. Blackmore and M. Ono, "Convex chance constrained predictive control without sampling," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Chicago, Illinois), August 2009.

[4] A. Prékopa, "The use of discrete moment bounds in probabilistic constrained stochastic programming models," *Annals of Operations Research*, vol. 85, pp. 21–38, 1999.

[5] D. van Hessem, *Stochastic inequality constrained closed-loop model predictive control with application to chemical process operation*. PhD thesis, Delft University of Technology, 2004.

[6] M. P. Vitus and C. J. Tomlin, "On feedback design and risk allocation in chance constrained control," in *Proceedings of the 50th IEEE Conference on Decision and Control*, (Orlando, Florida), December 2011.

[7] M. P. Vitus and C. J. Tomlin, "Belief space planning for linear, gaussian systems in uncertain environments," in *18th World Congress of the International Federation of Automatic Control*, (Milan, Italy), September 2011.

[8] N. Du Toit and J. Burdick, "Probabilistic collision checking with chance constraints," *IEEE Transactions on Robotics*, vol. 27, pp. 809 –815, August 2011.

[9] P. Missiuro and N. Roy, "Adapting probabilistic roadmaps to handle uncertain maps," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2006)*, (Orlando, FL), pp. 1261–1267, May 2006.

[10] M. Ono, L. Blackmore, and B. C. Williams, "Chance constrained finite horizon optimal control with nonconvex constraints," in *Proceedings of the 2010 American Control Conference*, (Baltimore, Maryland), June 2010.

[11] M. P. Vitus, V. Pradeep, G. M. Hoffmann, S. L. Waslander, and C. J. Tomlin, "Tunnel-MILP: Path planning with sequential convex polytopes," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Honolulu, Hawaii), August 2008.

[12] R. T. Rockafeller and S. P. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–42, 2000.

[13] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal of Optimization*, vol. 17, no. 4, pp. 969–996, 2006.

[14] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21." `http://cvxr.com/cvx`, April 2011.