2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

On Feedback Design and Risk Allocation in Chance Constrained Control

Michael P. Vitus
Department of Aeronautics and Astronautics
Stanford University
Email: vitus@stanford.edu

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

***Abstract*— This paper considers the problem of planning for linear, Gaussian systems, and extends existing chance constrained optimal control solutions. Due to the imperfect knowledge of the system state caused by process uncertainty and sensor noise, the system constraints cannot be guaranteed to be satisfied and consequently must be considered probabilistically. Therefore they are formulated as convex constraints on a Gaussian random variable, with the violation probability of all the constraints guaranteed to be below a threshold. Previous work considered optimizing the feedback controller to shape the uncertainty of the system to facilitate the satisfaction of the stochastic constraints. The joint constraints were bounded using an ellipsoidal relaxation technique which assigns uniform risk to each constraint. However this results in a large amount of conservatism degrading the performance of the overall system. Instead of using the ellipsoidal relaxation technique, this work bounds the joint constraints using Boole's inequality which results in a tighter approximation. The conservatism is further reduced by optimizing the risk assigned to each constraint along with the feedback controller. A two-stage optimization algorithm is proposed that alternates between optimizing the feedback controller and the risk allocation until convergence. This solution methodology is shown to reduce the conservatism in previous approaches and improve the performance of the overall system.**

I. Introduction

The primary motivating application for this work is in stochastic motion planning for robotic systems. Robotic and autonomous systems are becoming increasingly prevalent in everyday life; there are robots which clean floors, telepresence robots that can buy and deliver breakfast, personal robots for cleaning a room or emptying a dishwasher, robotic surgeons, driver assistance systems for automatic parking and adaptive cruise control, and even fully autonomous cars. A critical challenge for planning in these systems is the presence of uncertainty. Other planning applications that share this challenge include air traffic control around weather, chemical process control, energy efficient control of buildings and electric cars, and financial engineering. For all of these examples, there is a pressing need to develop robust and safe algorithms in the presence of uncertainty.

Uncertainty in stochastic systems arises from three different sources: (i) process uncertainty, (ii) sensing noise 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 problem must be solved in the space of probability distributions of the system state, defined as the belief space. For a stochastic system, however, planning in the belief space does not guarantee success because there is always a small probability that a large disturbance will occur. Therefore, there is a trade-off between the conservativeness of the plan and the performance of the system.

One way of planning in the belief space is using model predictive control (MPC). In robust MPC [1], the worst-case objective function, which corresponds to the worst possible realization of the uncertainty, is minimized. Tube-based MPC [2] is another approach that tightens the constraints to guarantee the system evolves in a "tube" of trajectories around the predicted nominal trajectory. Through this tightening the system is guaranteed to satisfy the original constraints under any realization of the uncertainty.

Planning under uncertainty can alternatively be handled by *chance constrained programming* introduced by Charnes and Cooper [3]. This formulation allows constraints with non-deterministic constraint parameters, named chance constraints, while only guaranteeing constraint satisfaction up to a specified limit. A thorough account of existing literature employing this problem formulation is given in [4].

The work by van Hessem et al. [5] used chance constrained programming to model process control as a stochastic control problem. They optimized over the feedback control laws and open-loop inputs while ensuring that the chance constraints were satisfied. They used an ellipsoidal relaxation technique to convert the stochastic problem into a deterministic one but this leads to a conservative solution. Blackmore [6] extended van Hessem's work to handle non-convex environments but still used the same conservative approximation.

Blackmore et al. used the chance constrained programming framework to solve the problem of motion planning in the presence of uncertainty. In [7], they extended their previous work to handle non-Gaussian belief distributions by approximating them using a finite number of particles. This transforms the original stochastic control problem into a deterministic one that can be efficiently solved. This sampling approach, however, becomes intractable as the number of samples needed to fully represent the true belief state increases. The work by Blackmore et al. [8] uses the work presented in [9] to approximate the chance constraints using Boole's inequality which typically leads to a very small amount of over-conservativeness. They also used the idea of risk allocation introduced by [9] to distribute the risk of violating each chance constraint while still guaranteeing the specified level of safety. By using the risk allocation technique instead of assuming a constant amount of risk for each constraint, the performance of the overall system can be significantly increased.

This work is supported in part by NASA (#NNA06CN22A) and ONR under MURI grant N00014-09-1-1051.

This work extends [5] which studied the problem of feedback design for stochastic systems formulated as a chance constrained program. The problem consists of two parts: optimizing over the feedback gains to shape the covariance of the system and satisfying the joint chance constraints. In previous work, an ellipsoidal relaxation technique was used to assign a static, uniform allocation of risk to each constraint. This transformed the problem into a deterministic program for optimizing the feedback parameters. This work investigates optimizing over both the feedback controller and the risk allocation of the joint chance constraints. Instead of using the ellipsoidal relaxation technique, the joint chance constraints are bounded using Boole's inequality, as in [8], which results in a tighter approximation. To reduce the conservatism even further, the risk allocation for each constraint is optimized as well as the feedback controller. A two-stage optimization algorithm is proposed that alternates between optimizing the feedback controller and the risk allocation until convergence. The algorithm is shown to significantly outperform previous approaches.

The paper proceeds as follows. Section II describes the stochastic problem formulation. Then, the explicit controller used in the closed-loop system is defined and a technique to convert the problem into a convex optimization program is presented in Section III. In Section IV, several techniques for evaluating the chance constraints are derived. The final solution methodology is presented in Section V, and an example is presented in Section VI which characterizes the performance of the algorithm.

II. Problem Formulation

Consider the following linear stochastic system defined by,

$$x_{k+1} = Ax_k + Bu_k + w_k, \forall k \in [0, N-1], \tag{1}$$

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

$$y_k = Cx_k + v_k, \forall k \in [0, N-1], \tag{2}$$

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, control inputs, measurements, and noise parameters for all time-steps are concatenated to form, $\mathbb{X} = \begin{bmatrix} x_0^{\mathrm{T}} & \dots & x_N^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$, $\mathbb{U} = \begin{bmatrix} u_0^{\mathrm{T}} & \dots & u_{N-1}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$, $\mathbb{Y} = \begin{bmatrix} y_0^{\mathrm{T}} & \dots & y_{N-1}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$, $\mathbb{W} = \begin{bmatrix} w_0^{\mathrm{T}} & \dots & w_{N-1}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$, and $\mathbb{V} = \begin{bmatrix} v_0^{\mathrm{T}} & \dots & v_{N-1}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$. Using this compact notation, the system equations can be written as,

$$\mathbb{X} = T_0 x_0 + H\mathbb{U} + G\mathbb{W} \tag{3}$$

$$\mathbb{Y} = C\mathbb{X} + \mathbb{V} \tag{4}$$

where in an abuse of notation $C = [\mathrm{diag}\,(C, \dots, C)\, 0] \in \mathbb{R}^{pN \times n(N+1)}$ with $\mathrm{diag}(\cdot)$ forming a block diagonal matrix from its arguments,

$$T_0 = \begin{bmatrix} I \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \; H = \begin{bmatrix} 0 & & & & \\ B & 0 & & & \\ AB & B & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ A^{N-1}B & A^{N-2}B & \dots & \dots & B \end{bmatrix},$$
$$G = \begin{bmatrix} 0 & & & & \\ I & 0 & & & \\ A & I & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ A^{N-1} & A^{N-2} & \dots & \dots & I \end{bmatrix}. \tag{5}$$

The system state is restricted to be in a feasible region denoted by F_X. For simplicity, 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 that the feasible region is convex, it can be defined by a conjunction of N_{F_X} linear inequality constraints,

$$F_X \triangleq \bigcap_{i=1}^{N_{F_X}} \left\{ \mathbb{X} : a_i^{\mathrm{T}} \mathbb{X} \leq b_i \right\} \tag{6}$$

where $a_i \in \mathbb{R}^{nN}$ and $b_i \in \mathbb{R}$. In this work, the environment parameters, a_i and b_i, are assumed to be deterministic. The expected value of the control inputs are also constrained to be in a feasible region F_U; the control inputs could also be constrained probabilistically.

Finally, the stochastic control problem can be stated as a chance constrained optimization problem in Program P2.1.

Chance Constrained Problem

$$\begin{aligned} &\text{minimize} \quad \mathbf{E}\left(f(\mathbb{X}, \mathbb{U})\right) \\ &\text{subject to} \\ &\qquad \mathbb{X} = T_0 x_0 + H\mathbb{U} + G\mathbb{W} \\ &\qquad \mathbb{Y} = C\mathbb{X} + \mathbb{V} \\ &\qquad \mathbb{W} \sim \mathcal{N}(0, \Sigma_w) \\ &\qquad \mathbb{V} \sim \mathcal{N}(0, \Sigma_v) \\ &\qquad \mathrm{E}\left(\mathbb{U}\right) \in F_U \\ &\qquad \mathrm{P}(\mathbb{X} \notin F_X) \leq \delta \end{aligned} \tag{P2.1}$$

In the current problem formulation there are two complications that prevent solving the optimization program P2.1 directly: (i) optimizing the controller that is used in the closed-loop system, and (ii) evaluating and satisfying the chance constraints $\mathrm{P}(\mathbb{X} \notin F_X) \leq \delta$. The next section addresses this first issue by developing the controller that is used to provide recourse into the system to reduce the uncertainty of the system state.

III. Closed-Loop System

A feedback controller can be used to shape the uncertainty of the system state, facilitating the satisfaction of the chance constraints and improving the objective function cost. As has been considered in the past [5], [12], only the set of affine causal output feedback controllers will be considered, i.e.,

$$u_k = \bar{u}_k + \sum_{\tau=0}^{k} F_{k,\tau} y_\tau, \forall k = 0, \ldots, N-1. \tag{7}$$

The control law could have also been defined in terms of a reference output, y_τ^r, i.e. $u_k = \bar{u}_k + \sum_{\tau=0}^{k} F_{k,\tau}\left(y_\tau - y_\tau^r\right)$. The compact form of Eqn. (7) is given by,

$$\mathbb{U} = \bar{\mathbb{U}} + F\mathbb{Y} \tag{8}$$

where

$$F = \begin{bmatrix} F_{0,0} & 0 & \ldots & 0 \\ F_{1,0} & F_{1,1} & \ddots & \vdots \\ \vdots & & \ddots & 0 \\ F_{N-1,0} & F_{N-1,1} & \ldots & F_{N-1,N-1} \end{bmatrix}, \tag{9}$$

and $\bar{\mathbb{U}} = \left[\bar{u}_0^{\mathrm{T}}, \ldots, \bar{u}_{N-1}^{\mathrm{T}}\right]^{\mathrm{T}}$.

Now the characteristics of the closed-loop state and input can be determined. Substituting for $\mathbb{X}$ from Eqn. (4) yields,

$$\mathbb{Y} = C\left(T_0 x_0 + H\mathbb{U} + G\mathbb{W}\right) + \mathbb{V}. \tag{10}$$

Substituting for $\mathbb{U}$ from Eqn. (8) results in,

$$\mathbb{Y} = C\left(T_0 x_0 + H\left(\bar{\mathbb{U}} + F\mathbb{Y}\right) + G\mathbb{W}\right) + \mathbb{V}. \tag{11}$$

After simplifying and combining common terms,

$$\left(I - CHF\right)\mathbb{Y} = C\left(T_0 x_0 + H\bar{\mathbb{U}} + G\mathbb{W}\right) + \mathbb{V}. \tag{12}$$

The term CHF is lower triangular, therefore $I - CHF$ is invertible. Solving for $\mathbb{Y}$ yields,

$$\mathbb{Y} = \left(I - CHF\right)^{-1}\left(C\left(T_0 x_0 + H\bar{\mathbb{U}} + G\mathbb{W}\right) + \mathbb{V}\right). \tag{13}$$

Substituting this into Eqn. (8) and simplifying yields,

$$\begin{aligned} \mathbb{U} = \; & F\left(I - CHF\right)^{-1} CT_0 x_0 + \\ & \left(I + F\left(I - CHF\right)^{-1} CH\right)\bar{\mathbb{U}} + \\ & F\left(I - CHF\right)^{-1} CG\mathbb{W} + \\ & F\left(I - CHF\right)^{-1}\mathbb{V}. \end{aligned} \tag{14}$$

Substituting this into Eqn. (3) results in,

$$\begin{aligned} \mathbb{X} = \; & \left(I + HF\left(I - CHF\right)^{-1} C\right) T_0 x_0 + \\ & H\left(I + F\left(I - CHF\right)^{-1} CH\right)\bar{\mathbb{U}} + \\ & \left(I + HF\left(I - CHF\right)^{-1} C\right) G\mathbb{W} + \\ & HF\left(I - CHF\right)^{-1}\mathbb{V}. \end{aligned} \tag{15}$$

In this formulation, the controller gains, F, are being optimized in order to shape the covariance to satisfy the constraints and improve the objective. In general, the program is not convex in the original design variables F and $\bar{\mathbb{U}}$. However, [5], [12] showed by a change of variables the problem can be cast as a convex optimization problem. This is accomplished by using the Youla parametrization as follows,

$$Q = F\left(I - CHF\right)^{-1}, \tag{16}$$

and the original gain matrix F can be efficiently solved for by

$$F = \left(I + QCH\right)^{-1} Q, \tag{17}$$

resulting in a block lower triangular matrix. Also, define

$$r = \left(I + QCH\right)\bar{\mathbb{U}} \tag{18}$$

such that

$$\bar{\mathbb{U}} = \left(I + QCH\right)^{-1} r = \left(I + FCH\right) r. \tag{19}$$

The variables Q and r can now be used in place of F and $\bar{\mathbb{U}}$ to generate a convex optimization program. Using this change of variables, the control input and state simplify to

$$\mathbb{U} = QCT_0 x_0 + r + QCG\mathbb{W} + Q\mathbb{V}, \tag{20}$$

$$\mathbb{X} = \left(I + HQC\right) T_0 x_0 + Hr + \left(I + HQC\right) G\mathbb{W} + HQ\mathbb{V}, \tag{21}$$

which are affine expressions in Q and r. Similarly, the uncertainty of the state and control input can be expressed in terms of Q and r,

$$\Sigma_{\mathbb{X}} = LT_0\Sigma_0 T_0^{\mathrm{T}} L^{\mathrm{T}} + LG\Sigma_{\mathbb{W}} G^{\mathrm{T}} L^{\mathrm{T}} + HQ\Sigma_{\mathbb{V}} Q^{\mathrm{T}} H^{\mathrm{T}}, \tag{22}$$

$$\Sigma_{\mathbb{U}} = QCT_0\Sigma_0\left(QCT_0\right)^{\mathrm{T}} + QCG\Sigma_{\mathbb{W}}\left(QCG\right)^{\mathrm{T}} + Q\Sigma_{\mathbb{V}} Q^{\mathrm{T}}, \tag{23}$$

where $L = I + HQC$.

Now that the feedback controller has been defined, the next complication in solving the optimization program is evaluating and satisfying the state chance constraints. While they may first appear to be easy to evaluate, they require the integration of a multivariate Gaussian density which does not have an analytic solution. Fortunately, there are many different ways to simplify the problem, which are presented in the next section.

IV. Chance Constraints

Given the complexity in integrating the multivariate Gaussian density, a simplification will be needed to reduce it to a tractable problem. Previously when optimizing over the controller feedback gains ([5], [6]), an ellipsoidal relaxation technique was used to simplify the constraints, but this typically leads to a very conservative approximation of the probability of failure which degrades the overall objective cost. Another approach to simplify the chance constraints is to use Boole's inequality which provides a tighter approximation of the probability of failure, but it complicates the solution of the optimization program. These two approaches are described in more detail below.

A. Ellipsoidal Relaxation

The ellipsoidal relaxation method has been previously developed by [5] to simplify the joint chance constraints and is presented for completeness.

Let $z \in \mathbb{R}^{n_z}$ and $z \sim \mathcal{N}(\bar{z}, \Sigma_z)$. The chance constraint $\mathrm{P}\left(h_i^{\mathrm{T}} z > g_i,\ \forall i\right) \leq \delta$ is equivalent to $\mathrm{P}\left(h_i^{\mathrm{T}} z \leq g_i,\ \forall i\right) \geq 1-\delta$, which is also equivalent to,

$$\alpha \int_{\mathcal{P}_z} \exp\left(-\frac{1}{2}(z-\bar{z})^{\mathrm{T}}\Sigma_z(z-\bar{z})\right) dz \geq 1-\delta, \quad (24)$$

where $\alpha = \dfrac{1}{\sqrt{(2\pi)^n \det(\Sigma)}}$ and $\mathcal{P}_z$ is the feasible region for z. There is no straightforward way of handling this integral constraint, therefore it was suggested to tighten the constraint as follows. If the constraint

$$\bar{z} + \mathcal{E}_r \subset \mathcal{P}_z \quad (25)$$

is ensured to be satisfied for the ellipsoid,

$$\mathcal{E}_r = \left\{\xi : \xi^{\mathrm{T}}\Sigma_z^{-1}\xi \leq r^2\right\}, \quad (26)$$

with an appropriately chosen r, then the original constraint $\mathrm{P}\left(h_i^{\mathrm{T}} z \leq g_i,\ \forall i\right) \geq 1-\delta$ is implied by

$$\alpha \int_{\bar{z}+\mathcal{E}_r} \exp\left(-\frac{1}{2}(z-\bar{z})^{\mathrm{T}}\Sigma_z(z-\bar{z})\right) dz \geq 1-\delta. \quad (27)$$

After simplifying Eqn. (27) to a one-dimensional integral, r is chosen such that

$$\frac{1}{2^{\frac{n_z}{2}}\Gamma\left(\frac{n_z}{2}\right)} \int_0^{r^2} \chi^{n_z/2-1}\exp(-\frac{\chi}{2}) d\chi = 1-\delta. \quad (28)$$

Now the original constraint $\mathrm{P}\left(h_i^{\mathrm{T}} z \leq g_i,\ \forall i\right) \geq 1-\delta$ can be replaced by requiring $\bar{z} + \mathcal{E}_r \subset F_z$ which is equivalent to requiring

$$h_i^{\mathrm{T}}\bar{z} + r\sqrt{h_i^{\mathrm{T}}\Sigma_z h_i} \leq g_i, \forall i. \quad (29)$$

B. Boole's Inequality

The chance constraint $\mathrm{P}(\mathbb{X} \notin F_X)$ can be converted into univariate integrals by using Boole's inequality to conservatively bound the probability of violation. Consequently, from Eqn. (6) and Boole's inequality the probability of the state not being contained inside the feasible region is bounded by,

$$\begin{aligned} \mathrm{P}(\mathbb{X} \notin F_X) &= \mathrm{P}\left(\mathbb{X} \in \bigcup_{i=1}^{N_{F_X}} \left\{\mathbb{X} : a_i^{\mathrm{T}}\mathbb{X} > b_i\right\}\right) \\ &\leq \textstyle\sum_{i=1}^{N_{F_X}} \mathrm{P}(a_i^{\mathrm{T}}\mathbb{X} > b_i). \end{aligned} \quad (30)$$

Now that the multivariate constraints have been converted to univariate constraints in Eqn. (30), they can be efficiently evaluated through,

$$\begin{aligned} \mathrm{P}(a_i^{\mathrm{T}}\mathbb{X} > b_i) &= \mathrm{P}(y_i > b_i) \\ &= 1-\Phi(\frac{b_i - a_i^{\mathrm{T}}\bar{\mathbb{X}}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}} a_i}}). \end{aligned} \quad (31)$$

The function $\Phi(\cdot)$ is the Gaussian cumulative distribution function which does not have an analytic solution, but it can be efficiently evaluated using a series approximation or a pre-computed lookup table. If the constraints

$$\begin{gathered} 1-\Phi\left(\frac{b_i - a_i^{\mathrm{T}}\mathbb{X}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}} a_i}}\right) \leq \delta_i \\ \sum \delta_i \leq \delta \end{gathered} \quad (32)$$

are satisfied then the original chance constraint will also be satisfied.

In the current problem formulation both δ_i and $\Sigma_{\mathbb{X}}$ are variables and as of now there isn't a direct way of handling them. However, if either variable is fixed then the optimization problem can be solved efficiently. This property is exploited in the proposed solution methodology in Section V.

For a fixed covariance, $\Sigma_{\mathbb{X}}$, and $\delta \leq 0.5$ the constraints in Eqn. (32) are convex since the function $\Phi(x)$ is concave in the range $x \in [0, \infty)$ [8], [13].

If the risk allocation, δ_i, is fixed such that $\sum \delta_i \leq \delta$ then the constraints in Eqn. (32) can be converted into equivalent second order cone constraints as follows. The constraints require

$$1-\Phi\left(\frac{b_i - a_i^{\mathrm{T}}\mathbb{X}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}} a_i}}\right) \leq \delta_i, \quad (33)$$

which can be simplified to

$$a_i^{\mathrm{T}}\mathbb{X} + \Phi^{-1}(1-\delta_i)\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}} a_i} \leq b_i \quad (34)$$

where Φ^{-1} is the inverse of the Gaussian cumulative distribution function. Finally, the standard deviation of the constraint uncertainty can be converted into a second order cone constraint,

$$\begin{aligned} &\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}} a_i} = \\ &\| \left[(I+HQC)\, T_0\Sigma_0^{\frac{1}{2}}, (I+HQC)G\Sigma_{\mathbb{W}}^{\frac{1}{2}}, HQ\Sigma_{\mathbb{V}}^{\frac{1}{2}}\right] a_i \|. \end{aligned} \quad (35)$$

C. Comparison

The previous two techniques are equally valid for simplifying and bounding the joint chance constraints, but the ellipsoidal relaxation technique is very conservative even for a small number of states. For example, for two states and two constraints the ellipsoidal relaxation method's conservativeness in the violation probability is 72.9% whereas using Boole's inequality results in only 0.7% conservativeness. Figure 1 shows a comparison of the two constraint relaxation techniques in which the original constraints are the red solid lines, the blue dash line is the ellipsoidal relaxation technique, and the green dash-dot line is for Boole's inequality with a uniform risk allocation. The feasible region for each technique is below the corresponding lines. The ellipsoidal relaxation technique clearly requires a larger backoff from the original constraints resulting in a larger over-approximation.

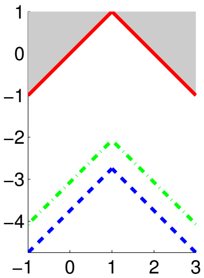

Fig. 1. A comparison of the two joint chance constraint relaxation techniques. The original constraints are the red solid lines, the ellipsoidal relaxation technique is shown as the blue dash line, and the Boole's inequality method is shown as the green dash-dot line with a uniform risk allocation. The feasible region for the state is below the corresponding lines.

D. Optimization Problem

Using either method for bounding the chance constraints, the Program P2.1 can be simplified to P4.1.

Chance Constrained Problem

$$
\begin{aligned}
&\underset{Q,r}{\text{minimize}} \quad \mathbf{E}\left(f(\mathbb{X},\mathbb{U})\right)\\
&\text{subject to}\\
&\quad \mathbb{X} = (I + HQC)\,T_0 x_0 + Hr + (I+HQC)\,G\mathbb{W} + HQ\mathbb{V}\\
&\quad \mathbb{U} = QCT_0x_0 + r + QCG\mathbb{W} + Q\mathbb{V}\\
&\quad \mathbb{W} \sim \mathcal{N}(0,\Sigma_w)\\
&\quad \mathbb{V} \sim \mathcal{N}(0,\Sigma_v)\\
&\quad \mathrm{E}(\mathbb{U}) \in F_U\\
&\quad L = I + HQC\\
&\quad \nu_i = \left\| \left[LT_0\Sigma_0^{\frac{1}{2}}, LG\Sigma_{\mathbb{W}}^{\frac{1}{2}}, HQ\Sigma_{\mathbb{V}}^{\frac{1}{2}}\right] a_i \right\|\\
&\quad a_i^{\mathrm{T}}\mathbb{X} + \beta_i\nu_i \le b_i,\ \forall i\\
&\quad Q \text{ lower triangular}
\end{aligned}
\tag{P4.1}
$$

In the optimization program P4.1, $\beta_i = r$ for the ellipsoidal relaxation technique, or $\beta_i = \Phi^{-1}(1-\delta_i)$ for Boole's inequality. Given a fixed controller F or fixed risk allocation δ_i, the Program P4.1 is a convex optimization problem which can be efficiently solved. However, optimizing over both the controller parameters and the risk allocation simultaneously isn't a straightforward procedure given the complication of β_i and ν_i being involved in a multiplicative constraint.

V. Solution Methodology: Two-Stage Method

Given the complexity in simultaneously optimizing the risk allocation as well as the controller parameters, an iterative two-stage optimization scheme is utilized. The upper stage allocates the risk associated to each individual constraint while the lower stage solves a second order cone program (SOCP) for the optimal control parameters given the current risk allocation. The difficulty in solving this problem is devising the iterative risk allocation scheme.

A. Upper Stage

There are many heuristic methods that can be devised for the upper stage which allocates the risk for each constraint.

1) Bisection Method: The simplest heuristic is to use a bisection method presented in Algorithm 1 to adjust the uniform allocation of the risk until the actual risk is within a specified threshold of the allowed value. The algorithm starts with two uniform risk allocations that result in a true probability of constraint violation, as calculated in lines 3 and 6 of Algorithm 2, above and below the desired value as described in Algorithm 2. A uniform risk allocation at the midpoint of the current two allocations that bracket the desired violation is then passed to the lower stage to solve for the controller parameters. For this solution, if the true probability of constraint violation, calculated in line 1 of Algorithm 1, is less than the allowed risk then it replaces the current lower bound otherwise it replaces the upper bound. This is continued until the actual risk is within a threshold of the allowed risk.

Algorithm 1 Bisection Method for Risk Allocation

1: $\delta_{i_{\text{true}}} = \sum_i 1 - \Phi\left(\dfrac{b_i - a_i^{\mathrm{T}}\mathbb{X}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}}a_i}}\right)$
2: **if** $|\delta_{i_{\text{true}}} - \delta| \le \epsilon$ **then**
3: Solution found.
4: **end if**
5: **if** $\delta_{i_{\text{true}}} \ge \delta$ **then**
6: $\overline{\delta} = \delta_i$
7: **else**
8: $\underline{\delta} = \delta_i$
9: **end if**
10: $\delta_i = 0.5(\underline{\delta} + \overline{\delta})$

Algorithm 2 Initialization for the Bisection Method

1: $\underline{\delta} = \delta / N_{F_{\mathbb{X}}}$
2: Perform optimization of P4.1 with $\beta_i = \Phi^{-1}(1-\underline{\delta})$
3: $\underline{\delta}_{\text{true}} = \sum_i 1 - \Phi\left(\dfrac{b_i - a_i^{\mathrm{T}}\mathbb{X}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}}a_i}}\right)$
4: $\overline{\delta} = 0.5$
5: Perform optimization of P4.1 with $\beta_i = \Phi^{-1}(1-\overline{\delta})$
6: $\overline{\delta}_{\text{true}} = \sum_i 1 - \Phi\left(\dfrac{b_i - a_i^{\mathrm{T}}\mathbb{X}}{\sqrt{a_i^{\mathrm{T}}\Sigma_{\mathbb{X}}a_i}}\right)$
7: **if** $\overline{\delta}_{\text{true}} < \delta$ **then**
8: Solution found, no need to iterate.
9: **end if**
10: $\delta_i = 0.5(\underline{\delta} + \overline{\delta})$

2) Optimal Risk Given Fixed Controller: As was stated before, if the controller parameters, F or Q, are fixed (and hence the covariance $\Sigma_{\mathbb{X}}$) then the problem simplifies to a convex optimization program. Consequently, this problem can be efficiently solved for the optimal risk allocation, $\delta_i\ \forall i$, for the particular controller parameters. However, there is no guarantee that this risk allocation is the optimal risk allocation for the original problem.

B. Lower Stage

Once the risk allocation has been fixed, the program P4.1 simplifies to a standard second order cone program which can be solved efficiently by many standard solvers.

VI. Applications: Motion Planning

The above method is applicable to many different fields and is evaluated on a stochastic motion planning problem. In this example, the system has double integrator dynamics with a 2D position, and $\Delta t = 0.1$ seconds. The noise parameters are $\Sigma_w = \mathrm{diag}(0.0003, 0.0005, 0.0003, 0.0005)$ and $\Sigma_v = \mathrm{diag}(0.001, 0.002)$. The objective function for this problem is quadratic in the final state as well as the control inputs,

$$f(\bar{\mathbb{X}}, \bar{\mathbb{U}}) = \left(x_N - x_{ref}\right)^{\mathrm{T}} Q_{obj} \left(x_N - x_{ref}\right) + \bar{\mathbb{U}}^{\mathrm{T}} R_{obj} \bar{\mathbb{U}} \tag{36}$$

with $Q_{obj} = I$, $R_{obj} = 0.001I$ and $x_{ref} = \begin{bmatrix} 2 & 1 & 0 & 0 \end{bmatrix}^{\mathrm{T}}$. The allowed risk is $\delta = 0.15$.

The results are shown in Figure 2. The solution with a fixed LQR trajectory tracking controller and Kalman filter is shown in Figure 2(a) with an objective function cost of 0.184. The system has to take the suboptimal route through the top; the bottom path is infeasible because the vertical uncertainty is too large for the allowed risk. The solution for the ellipsoidal relaxation technique is shown in Figure 2(b). Here, the system also had to take the suboptimal route through the top of the environment because the relaxation method for the joint constraints was too conservative. The objective function cost is 0.186 with a probability of failure of 0.004 which is significantly less than the allowed risk.

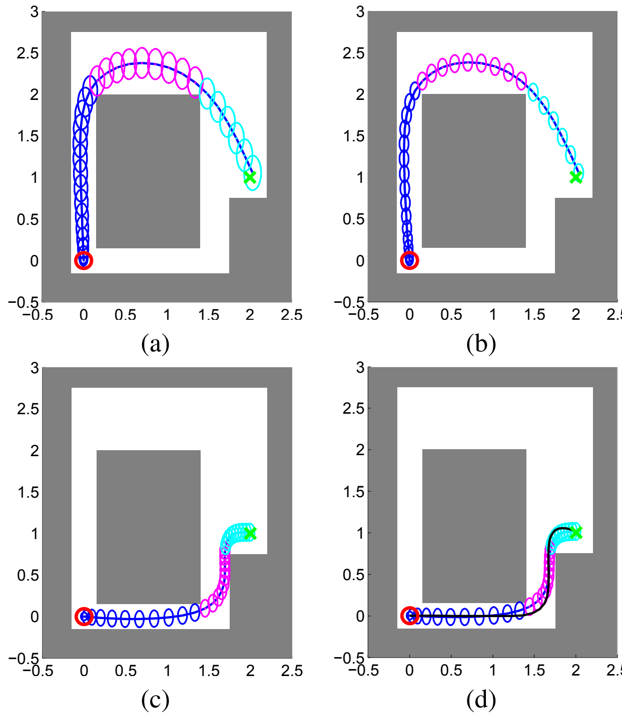

Fig. 2. The solution and the confidence ellipsoids for a nonconvex environment. (a) Solution for an assumed LQR trajectory tracking controller. (b) The solution for the ellipsoidal relaxation method. (c) The solution from the two stage algorithm using the bisection method for the upper stage. (d) The solution for the optimal risk given fixed controller upper stage. The black line is the solution from the SOCP solver assuming a uniform allocation of risk. The blue line, with the confidence ellipsoids around it, is the solution using the controller parameters from the first SOCP stage and optimizing the risk for each constraint.

The solution using the two-stage algorithm for optimizing the output feedback controller parameters is shown in Figure 2(c-d). Figure 2(c) shows the solution from the binary search upper stage algorithm. The algorithm took 16 iterations for an $\epsilon = 1 \times 10^{-5}$ and the objective function's value was 0.120. The solution for the optimal risk given fixed controller upper stage is shown in Figure 2(d) which only took 1 iteration. The algorithm was initialized with a uniform allocation of the risk $\delta_i = \delta / N_{F_{\mathbb{X}}}$, and the trajectory for this allocation is shown as the black line. The final solution is shown as the blue line with the 90% confidence ellipsoids around it. The initial objective cost is 0.146 and the final objective cost is 0.116 which is a 26% relative improvement.

VII. Conclusion

The planning problem for a linear, Gaussian stochastic system with state constraints was formulated as a stochastic optimal control problem. A two-stage solution methodology was proposed that allocates the risk in one stage and optimizes the feedback controller in the other. This methodology reduces the conservatism in the probability of violation calculation and outperforms the other proposed methods in the example shown.

References

[1] P. J. Campo and M. Morari, "Robust model predictive control," in *Proceedings of the 1987 American Control Conference*, (Minneapolis, MN), June 1987.

[2] W. Langson, I. Chryssochoos, S. V. Raković, and D. Mayne, "Robust model predictive control using tubes," *Automatica*, vol. 40, pp. 125–133, January 2004.

[3] A. Charnes and W. Cooper, "Deterministic equivalents for optimizing and satisfying under chance constraints," *Operations Research*, no. 11, pp. 18–39, 1963.

[4] A. Prékopa, *Stochastic Programming*. Kluwer Scientific, 1995.

[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] L. Blackmore, "Robust path planning and feedback design under stochastic uncertainty," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Honolulu, HI), August 2008.

[7] L. Blackmore, "A probabilistic particle control approach to optimal, robust predictive control," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Keystone, Colorado), August 2006.

[8] 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.

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

[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, S. L. Waslander, and C. J. Tomlin, "Locally optimal decomposition for autonomous obstacle avoidance with the tunnel-MILP algorithm," in *Proceedings of the 47th IEEE Conference on Decision and Control*, (Cancun, Mexico), December 2008.

[12] J. Skaf and S. Boyd, "Design of affine controllers via convex optimization," *Automatic Control, IEEE Transactions on*, vol. 55, pp. 2476–2487, November 2010.

[13] M. P. Vitus and C. J. Tomlin, "Closed-loop belief space planning for linear, gaussian systems," in *2011 IEEE International Conference on Robotics and Automation*, (Shanghai, China), May 2011.