Convex Chance Constrained Predictive Control without Sampling*

Lars Blackmore[†] and Masahiro Ono[‡]

In this paper we consider finite-horizon optimal 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 finite-horizon 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 finite-horizon 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, as long as the prescribed maximum probability of constraint violation is below 0.5, the resulting optimization is convex and hence amenable to online control design.

I. Introduction

In this paper we are concerned with finite-horizon optimal control of dynamics systems subject to state and control constraints.[1–3] This classical problem consists of choosing a sequence of control inputs that minimizes some cost functional, defined over the control inputs and the state trajectory, while ensuring that constraints on the control inputs and state are satisfied. These problems have been investigated extensively since the early 60's, beginning with the work of Pontryagin[2] and Bellman.[3] Applications include optimal guidance for spacecraft, where we must find the sequence of thrust inputs that take the spacecraft from its initial state to a desired terminal condition, while avoiding collisions.[4] In the presence of inequality constraints on the state, finite-horizon optimal control problems are typically solved by parameterizing the control sequence with a finite number of parameters. A common parameterization is to assume that the control sequence is piecewise constant over a small, fixed timestep Δt. Parameterization enables the infinite-parameter problem to be converted into a constrained optimization problem, which can then be solved using existing optimization tools.

We are interested in the case where the dynamic system to be controlled is subject to stochastic uncertainty. 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.[5] 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 failure due to 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).

*US Government sponsorship acknowledged.

[†]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.

[‡]M. Ono is with the Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 USA e-mail: hiro_ono@mit.edu.

In this paper we present a new approach to solving the chance-constrained finite-horizon control problem for the case of Gaussian uncertainty and linear discrete-time systems. This problem is challenging for the key reason that evaluating an integral over a multivariate Gaussian is not possible in closed form. This means that a table lookup must be used to evaluate the integral. While this is tractable for univariate Gaussians, the size of the table scales exponentially with the dimension of the distribution. Table lookup is hence not a practical approach for multivariate Gaussians. Previous approaches have used approximate sampling techniques[6–8] or conservative bounding methods[9–12] to avoid this problem. The key idea behind our new approach is to use Boole's inequality to split multivariate chance constraints into many univariate chance constraints, which can be evaluated efficiently. By also optimizing the risk of each univariate constraint being violated, we ensure that the conservatism introduced is very small. Furthermore, we show that the resulting optimization is convex, and can be solved efficiently to global optimality using existing nonlinear solvers.

II. Related Work

A number of approaches to chance-constrained finite-horizon control have been proposed in recent years, in the context of model predictive control (MPC). MPC is a closed-loop control approach that, at each time step, solves a finite-horizon optimal control problem from the current initial state, executes the first step in the resulting optimal control sequence, and then resolves at the next time step. We are not concerned with such *receding-horizon* approaches in the present paper, and to the authors' knowledge, no results exist that guarantee the satisfaction of chance constraints for a closed-loop receding-horizon scheme. For early results in this area, we refer the reader to Ref. 13. However there are a number of results in the MPC literature that address the *finite-horizon* chance-constrained optimal control problem. In the case of Gaussian uncertainty distributions, linear system dynamics and convex feasible regions, Ref. 14 considered chance constraints on individual scalar values. The extension from scalar random variables to joint random variables is essential if we wish to constrain the probability of failure over the entire planning horizon. The work of Refs. 6, 7 considered chance constraints on joint random variables, using the result of Ref. 15 to show that the optimization resulting from the chance-constrained finite-horizon 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.

Ref. 9 used a conservative ellipsoidal set bounding approach to ensure that the chance constraints are satisfied without the need for the evaluation of multivariate Gaussian integrals. The key idea is to characterize a region around the state mean that the state is guaranteed to be in with a certain probability (the '99%' region) and ensure that this deterministic set satisfies the constraints. The approach of Ref. 9 explicitly optimized over feedback laws as well as feedforward controls[a]. An alternative bounding approach is to use Boole's inequality to split joint chance constraints over N variables into N univariate chance constraints, and to ensure that the probability of violation of each of these is at most δ/N, where δ is the specified maximum probability of failure. This approach was suggested by Ref. 11 for convex feasible regions and by Ref. 10 for nonconvex feasible regions. Another bounding approach was proposed by Ref. 12, which does not require that all uncertainty is Gaussian. In this work, the authors draw n samples or 'scenarios' from the random variables and ensure that the constraints are satisfied for all of the samples. The bound in the scenario approach is stochastic, in the sense that n is chosen to ensure that the chance constraint is satisfied with probability $1 - \beta$, where β is small and chosen by the user.

Bounding approaches are, however, prone to excessive conservatism whereby the true probability of constraint violation is far lower than the specified allowable level. Conservatism leads to excess cost and can prevent the optimization from finding a feasible solution at all. Our previous work introduced an approach called *risk allocation*, whereby the risk of constraint violation is not fixed at δ/N but is made an explicit optimization parameter.[16] In Ref. 16 optimization takes place in two stages; in the first stage the allocation of risks to univariate constraints is optimized, while in the second the control sequence is optimized subject to the univariate chance constraints being satisfied at the allocated risk level. In the present paper we extend this work to show that for convex polytopic state constraints the problem can be solved as a single convex

[a]Note however, that the chance constraints are only guaranteed to hold in finite-horizon (open loop) execution, rather than in receding horizon (closed loop).

optimization problem[b]. Furthermore, we show empirically that risk allocation ensures 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 many orders of magnitude.

III. 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. The discrete-time system (1) is typically generated by performing a finite parameterization of a continuous-time system, for example assuming a piecewise-constant control input sequence.

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.

In finite-horizon optimal 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 Ref. 9. The chance-constrained optimal control problem can now be stated.

Problem 1. *Chance-Constrained Control Problem.*

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

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.

IV. 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.

[b] This is in contrast to Ref. 11, who state that, in the case of *general* convex constraints, the risk allocation problem is nonconvex (Remark 2.1 in Ref. 11).

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.*

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

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 work used this result, together with the convexity result of Ref. 15, to show that Problem 2 is convex.[6,7] 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 Ref. 6 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 Ref. 6 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.

V. New Approach: Convex Risk Allocation

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 VII 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.*

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

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 $A_1, \dots, A_N$:

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

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

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

For a feasible solution to Problem 3 we know that $\sum_i P(\mathbf{a}_i'\mathbf{x} > b_i) \le \delta$, and hence $P(\mathbf{x} \notin F) \le \delta$. 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. Our previous work proposed this idea of *risk allocation*,[16] however in that work the optimization of the ϵ_i was carried out in a separate optimization step. Other related work[10,11] 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 cdf(·) only once. While cdf(·) 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.*

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

In the following two lemmas we show that the constraints in Problem 4 are convex. The idea for this result comes from our previous work in Ref. 17.

Lemma 2. *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:

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

Proof: Concavity comes from the fact that 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:*

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

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:

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

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). □

Theorem 1. *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 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 VII that the conservatism introduced is small.

VI. Comparison with Set Conversion Techniques

Alternative conservative approximations of the chance constrained problem (Problem 1) have been proposed previously, for example Ref. 9. 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 Ref. 18. 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 Ref. 9, for details. We now discuss how this set conversion approach relates to our new convex optimization approach. In Section VII 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.

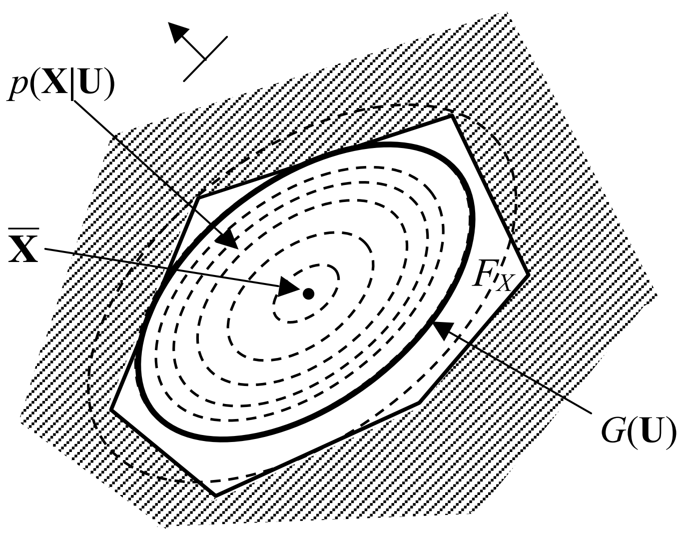

Figure 1. Chance constrained optimization problem approximated using ellipsoidal set conversion. Shown is a two-dimensional Gaussian distribution for $\mathbb{X}$, represented using contours of the pdf (dashed). The thick ellipse is the set $G(\mathbb{U})$ containing 99% of the probability mass. In this 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 $G(\mathbb{U})$.

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

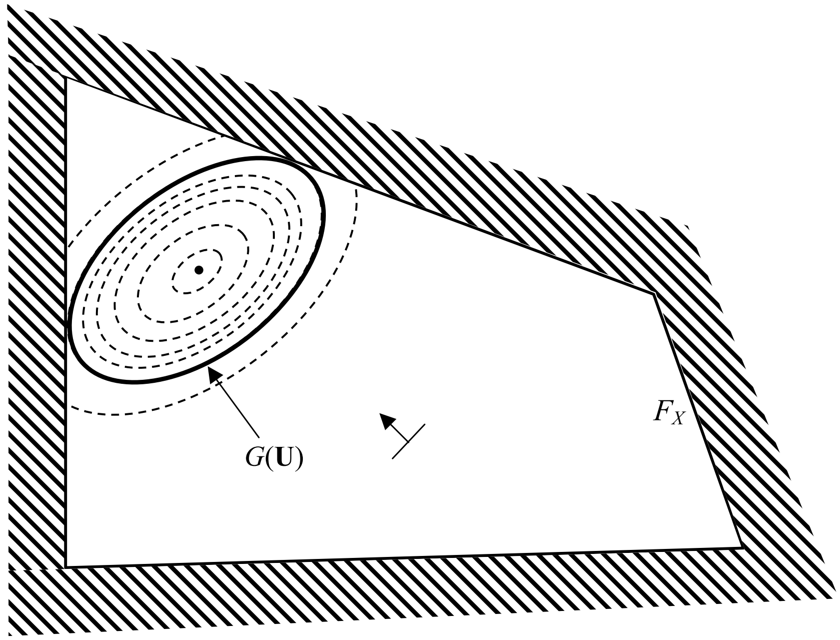

Figure 2. Chance constrained optimization problem with general feasible region (geometry dissimilar to $G(\mathbb{U})$). In the optimal solution returned by the set conversion method, the set $G(\mathbb{U})$ is constrained to lie within F_X. The conservatism introduced by the set bounding approach corresponds to the integral of the pdf over the region that is inside F_X but outside the ellipsoid $G(\mathbb{U})$. 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.

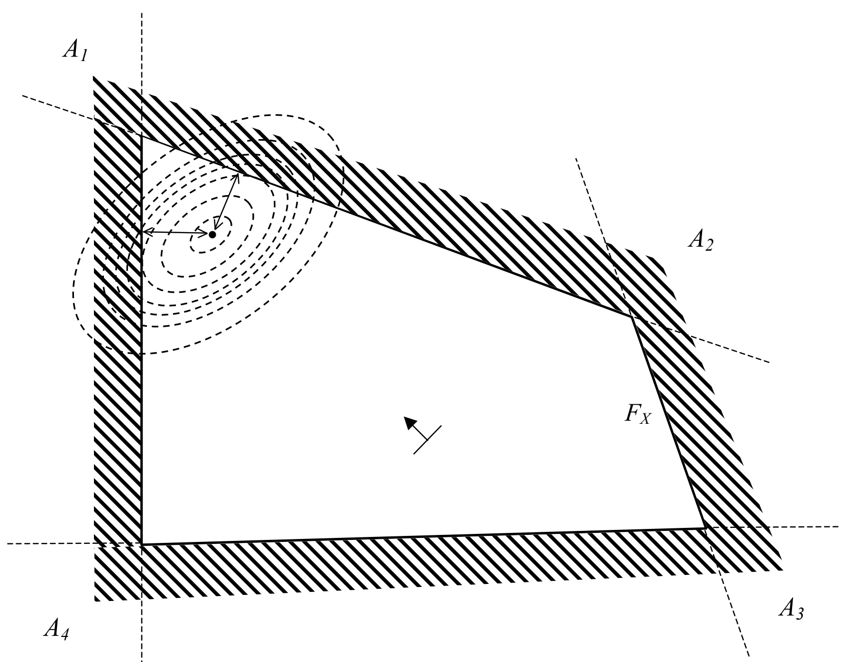

Figure 3. Chance constrained optimization using new convex risk allocation approach. The approach optimally allocates the contribution to the failure probability from each of the constraints. In the solution shown, two of the constraints have approximately zero probability of violation, so the algorithm 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 regions A_1 through A_4, which is small in practice.

assigned to each constraint, denoted ϵ_i. In doing so it can greatly reduce the conservatism of the solution, as illustrated in Figure 3. The new approach does introduce conservatism in the use of Boole's inequality. However we show empirically in Section VII that this conservatism is small, while the set conversion approach is conservative by many orders of magnitude.

VII. Simulation Results

In this section we show simulation results demonstrating the new approach to chance-constrained optimal control. 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 SNOPT.[19] Optimization was performed on a MacBook Pro with a 2.4GHz processor and 4GB RAM.

A. Example Solutions

Figure 4 shows the solution to the finite-horizon optimal control problem using the new risk allocation 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. The non-negligible values correspond to the bound $y_k \leq 0.25$ at time steps $k = 1, \ldots, 5$. 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 Ref. 9. 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 10^7 Monte Carlo simulations and estimated the true probability of constraint violation. Table 1 compares the conservatism of the new risk allocation approach with that of elliptical set conversion and the scenario approach of Ref. 12. In this case $\mathbf{x}^r = [1 \; 0]'$, $N = 20$ and $\delta = 0.1$. For elliptical set conversion and the scenario approach, optimization was performed using SDPT3,[20] which is able to more efficiently exploit the structure of the Second Order Cone Programs that result from the approach. For the scenario approach we used $\beta = 0.001$, resulting in 2615 scenarios (samples) being used to ensure that $P(P(\mathbb{X} \notin F_X) > 0.1) \leq 0.001$. Table 1 shows that the risk allocation approach is orders of magnitude less conservative that the elliptical set conversion approach, and as a result the cost of the optimal solution is reduced by 40%. Compared to the scenario approach, risk allocation is both less conservative and far less computationally expensive. It should be noted, however, that the scenario approach applies to arbitrary uncertainty distributions, whereas both risk allocation and the elliptical set bounding approaches are restricted to Gaussian distributions.

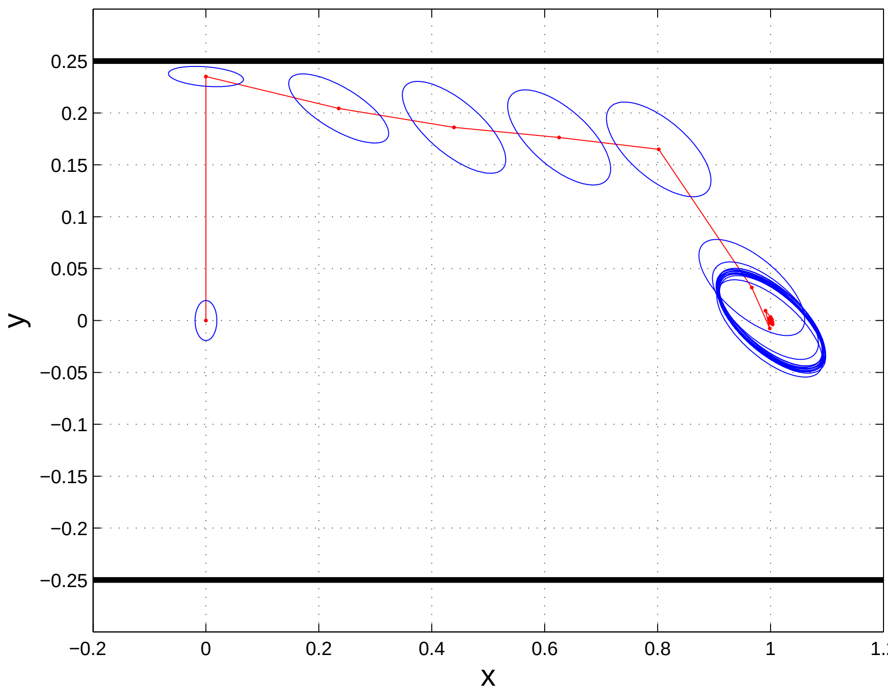

Figure 4. Single solution using new convex optimization approach for $\mathbf{x}^r = [1\ 0]'$, $T = 20$ and $\delta = 0.01$. The red dots show the state mean $\bar{\mathbf{x}}_k$ for $k = 0, \ldots, T$. The blue ellipses show the 3-sigma ellipses for $\mathbf{x}_k$. The state constraints are shown as 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.

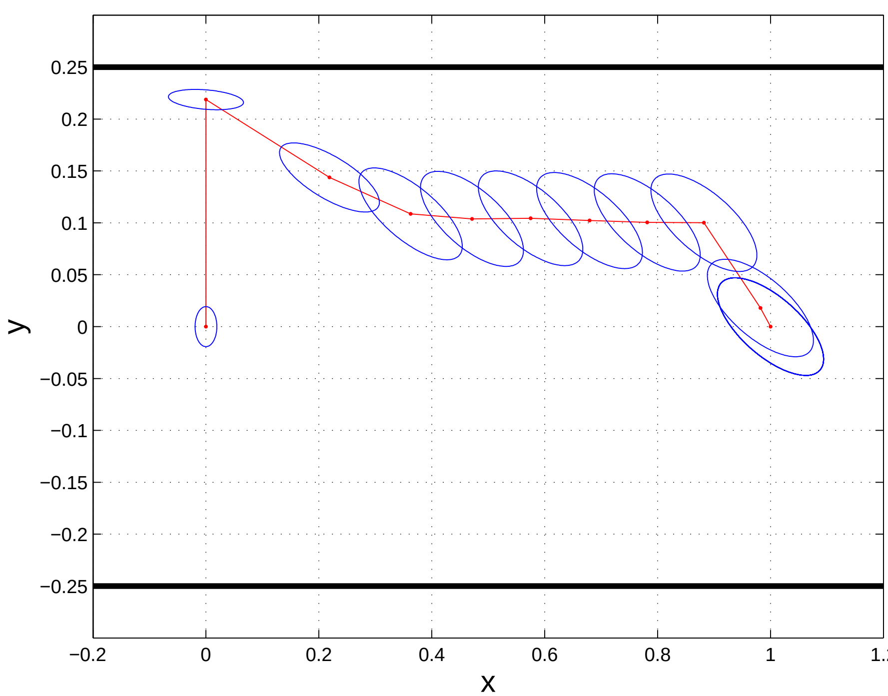

Figure 5. Single solution for $\mathbf{x}^r = [1\ 0]'$, $T = 20$ and $\delta = 0.01$, using elliptical set conversion approach of Ref. 9. The state means are very far from the constraints compared to the solution in Figure 4, indicating a great deal of conservatism.

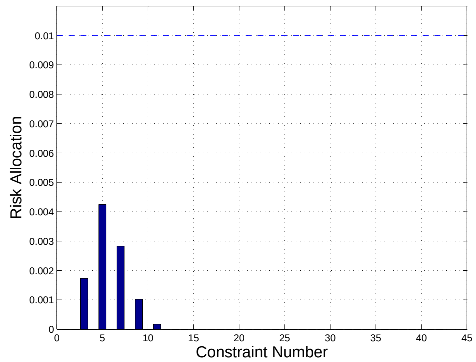

Figure 6. Risk allocation in optimal solution of Figure 4. The bars show the value ϵ_i, i.e. the risk allocation, for each constraint in the problem. The allocated risks add to the maximum probability of failure for the entire horizon, δ, shown as the dashed line. Only a handful of the ϵ_i are non-negligible.

	Risk Allocation	Elliptical Set Conversion	Scenario
Max $P(fail)$ (δ)	0.1	0.1	0.1
Estimated $P(fail)$	0.097	2×10^{-6}	0.0022
Optimal Cost	3.15	5.26	3.76
Optimization Time (s)	3.38	1.76	1.41×10^4

Table 1. Comparison of new risk allocation approach with elliptical set conversion for a single example. The risk allocation approach is orders of magnitude less conservative that the elliptical set conversion approach, and as a result the cost of the optimal solution is reduced by 40%. Also shown are results for the scenario approach, which is more conservative that risk allocation and requires orders of magnitude greater optimization time.

B. Conservatism Against δ and Horizon Length

Figure 7 shows how the conservatism of the new approach depends on the allowable maximum probability of failure, δ. For this example we used $T = 10$ steps and $\mathbf{x}^r = [1\ 0]'$, and 10^7 Monte Carlo simulations were used to estimate the true probability of failure in the returned solution. Risk allocation yields probabilities of failure very close to the allowable value, indicating very little conservatism. Critically, the conservatism does not increase appreciably as δ decreases. For comparison Figure 7 also shows the estimated probability of failure for elliptical set bounding. Note that this approach is highly conservative, and the conservatism increases as δ decreases. This conservatism prevents elliptical set bounding from finding a feasible solution for $\delta \leq 0.01$, even though feasible solutions exist for much smaller δ.

Figure 8 shows how the conservatism of the new approach depends on the horizon length, T. For this example we used $\delta = 0.1$ and $\mathbf{x}^r = [1\ 0]'$, and 10^7 Monte Carlo simulations were used to estimate the true probability of failure in the returned solution. Risk allocation gives solutions with very little conservatism, and the conservatism increases only slightly as the horizon length, T, increases; for $T = 5$, the estimated probability of failure is 0.0974, while for $T = 20$ the estimated probability of failure is 0.0970. By contrast the set bounding approach is highly conservative, and the conservatism increases as the horizon length increases. This conservatism prevents set bounding from finding a feasible solution for $T > 10$.

C. Average performance

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 $T = 20$ and $\delta = 0.01$. The average results for 20 solutions are shown in Table 2. 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.

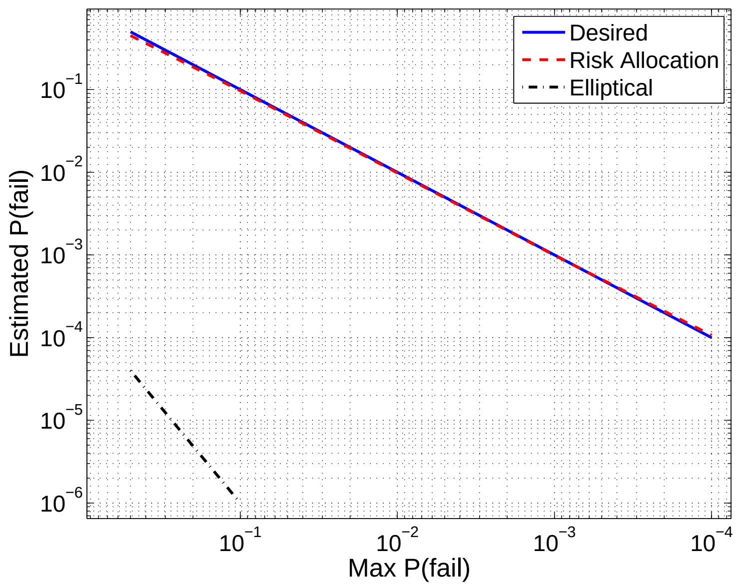

Figure 7. Comparison of maximum allowable probability of failure δ and estimated probability of failure as a function of δ. Risk allocation yields probabilities of failure very close to the allowable value, indicating very little conservatism. Furthermore this conservatism does not increase as δ decreases. By contrast the set bounding approach is highly conservative, and the conservatism increases as δ decreases. In fact set bounding fails to find a solution for $\delta \leq 0.01$, even though feasible solutions exist for much smaller δ.

Algorithm	Time (s)	$P(fail)$
Risk Allocation	1.03	0.0079
Elliptical Set Conversion	0.22	$< 10^{-6}$

Table 2. Optimization time and estimated 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. Conclusion

We have proposed new approaches for chance-constrained finite-horizon optimal 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.

IX. 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.

References

[1]Berkovitz, L. D., *Optimal Control Theory*, Springer Verlag, 1974.

[2]Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., and Mischenko, E. F., *The Mathematical Theory of Optimal Processes*, Interscience Publishers, 1962.

[3]Bellman, R., *Introduction to the The Mathematical Theory of Control Processes*, Academic Press, 1967.

[4]Acikmese, B. and Ploen, S. R., "Convex Programming Approach to Powered Descent Guidance for Mars Landing," *AIAA Journal of Guidance, Control and Dynamics*, Vol. 30, No. 5, 2007, pp. 1353–1366.

[5]Barr, N. M., Gangsaas, D., and Schaeffer, D. R., "Wind Models for Flight Simulator Certification of Landing and Approach Guidance and Control Systems," Paper FAA-RD-74-206, 1974.

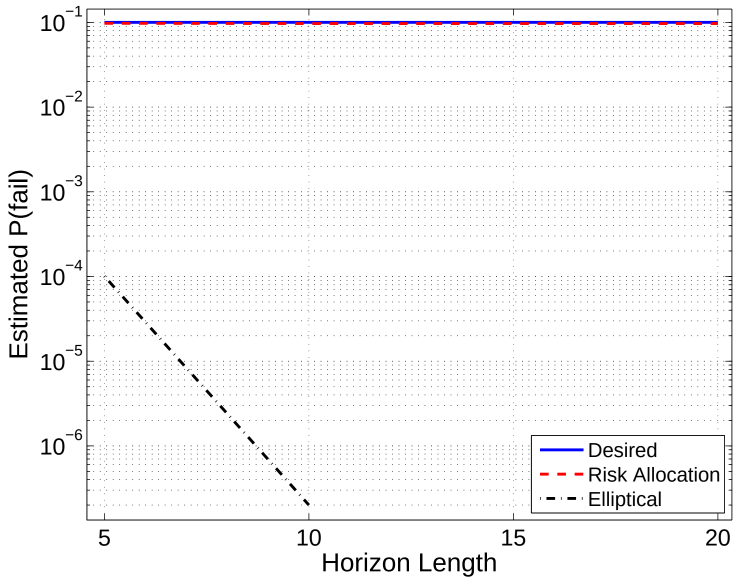

Figure 8. Comparison of maximum allowable probability of failure δ and estimated probability of failure as a function of horizon length. Risk allocation gives solutions with very little conservatism, and the conservatism increases only slightly as the horizon length, T, increases. By contrast the set bounding approach is highly conservative, and the conservatism increases as the horizon length increases. Set bounding fails to find a solution for $T > 10$.

[6]Li, P., Wendt, M., and Wozny, G., "Robust Model Predictive Control under Chance Constraints," *Computers and Chemical Engineering*, Vol. 24, 2000, pp. 829–834.

[7]Li, P., Wendt, M., and Wozny, G., "A Probabilistically Constrained Model Predictive Controller," *Automatica*, Vol. 38, 2002, pp. 1171–1176.

[8]Blackmore, L., "A Probabilistic Particle Control Approach to Optimal, Robust Predictive Control," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2006.

[9]van Hessem, D., Scherer, C. W., and Bosgra, O. H., "LMI-based closed-loop economic optimization of stochastic process operation under state and input constraints," *Proceedings of the Control and Decision Conference*, 2001.

[10]Blackmore, L., Li, H. X., and Williams, B. C., "A Probabilistic Approach to Optimal Robust Path Planning with Obstacles," *Proceedings of the American Control Conference*, 2006.

[11]Nemirovksi, A. and Shapiro, A., "Convex Approximations of Chance Constrained Programs," *SIAM Journal of Optimization*, Vol. 17, No. 4, 2006, pp. 969–996.

[12]Calafiore, G. C. and Campi, M. C., "The Scenario Approach to Robust Control Design," *IEEE Transactions on Automatic Control*, Vol. 51, No. 5, 2006.

[13]Primbs, J. A., "Receding Horizon Control of Constrained Linear Systems with State and Control Multiplicative Noise," *Proceedings of the American Control Conference*, 2007.

[14]Schwarm, A. and Nikolaou, M., "Chance-constrained Model Predictive Control," *AlChE Journal*, Vol. 45, 1999, pp. 1743–1752.

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

[16]Ono, M. and Williams, B. C., "Efficient Motion Planning Algorithm for Stochastic Dynamic Systems with Constraints on Probability of Failure," *Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-08)*, Chicago, IL, July 2008.

[17]Ono, M. and Williams, B. C., "Two-stage Optimization Approach to Robust Model Predictive Control with a Joint Chance Constraint," *Proceedings of 47th IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2008.

[18]Richards, A., *Robust Constrained Model Predictive Control*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2005.

[19]Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131.

[20]Tutuncu, R. H., Toh, K. C., and Todd, M. J., "Solving semidefinite-quadratic-linear programs using SDPT3," *Mathematical Programming B*, Vol. 95, 2003, pp. 189–217.