

Mixed-Strategy Chance Constrained Optimal Control

Masahiro Ono, Yoshiaki Kuwata, and J. (Bob) Balaram

Abstract—This paper presents a novel chance constrained optimal control (CCOC) algorithm that chooses a control action *probabilistically*. A CCOC problem is to find a control input that minimizes the expected cost while guaranteeing that the probability of violating a set of constraints is below a user-specified threshold. We show that a probabilistic control approach, which we refer to as a mixed control strategy, enables us to obtain a cost that is better than what deterministic control strategies can achieve when the CCOC problem is nonconvex. The resulting mixed-strategy CCOC problem turns out to be a convexification of the original nonconvex CCOC problem. Furthermore, we also show that a mixed control strategy only needs to “mix” up to two deterministic control actions in order to achieve optimality. Building upon an iterative dual optimization, the proposed algorithm quickly converges to the optimal mixed control strategy with a user-specified tolerance.

I. INTRODUCTION

A. Illustrative Example

One may better off by deciding his/her action by a coin flip in a chance constrained optimal control (CCOC) problem. More formally, a *mixed control strategy*, which chooses a control input probabilistically, can achieve a cost that is better than the one obtained by any deterministic control strategy.

To illustrate this concept, we first present an example. Consider a system that accepts three discrete options of control inputs, u_A , u_B , and u_C , as shown in Table I. We assume that the system is subject to uncertainty. Hence, each of the control actions results in an uncertain cost and involves a risk of constraint violation. The expected cost and the risk of constraint violation of each control option is shown in Table I. We solve the following CCOC problem:

Problem (illustrative example): Find the control input that minimizes the expected cost while limiting the probability of constraint violation to 1%.

If only a deterministic control strategy is allowed, the optimal solution is u_B , with the expected cost being 30 and the probability of constraint violation being 1%.

The optimal solution of a mixed control strategy, however, is to choose u_A with a 50% probability and u_C with a 50% probability. With this mixed control strategy, the probability of constraint violation is still 1% but the expected cost is 25,

Masahiro Ono (ono@appi.keio.ac.jp) is with Keio University, Kanagawa, Japan. Yoshiaki Kuwata (kuwata@alum.mit.edu) and J. (Bob) Balaram (J.Balaram@jpl.nasa.gov) are with Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA. Copyright 2012 California Institute of Technology. Government sponsorship acknowledged.

TABLE I

AVAILABLE CONTROL OPTIONS FOR THE ILLUSTRATIVE EXAMPLE PRESENTED IN SECTION I-A.

Control input	Expected cost	Risk
u_A	40	0.5 %
u_B	30	1.0 %
u_C	10	1.5 %

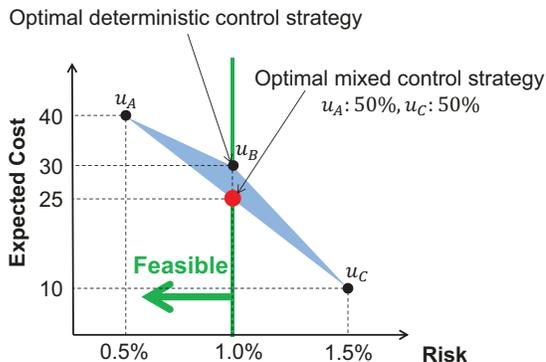


Fig. 1. Graphical interpretation of a mixed control strategy. Three black dots show the deterministic control options given in Table I. The set of the costs and risks of all the possible mixed control strategies, which is represented by the shadowed triangle, is the convex hull of the set of the costs and risks of all the deterministic control strategies, shown as the three black dots. The optimal mixed control strategy is found at the intersection of the 1% risk line and the bottom edge of the triangle, which corresponds to the 50% – 50% mix of u_A and u_C .

which is less than that of the optimal deterministic control strategy.

Figure 1 graphically explains this result. The three deterministic three control options are plotted in dots in the expected cost-risk space. The control options that satisfy the chance constraint lie on the left side of the vertical line with a 1% risk. The optimal deterministic control is u_B because it has the lowest cost among the feasible options.

The set of the costs and risks that can be realized by all possible mixed control strategies is the convex hull of the set of the costs and risks of the deterministic control strategies. This is because the cost and risk of a mixed control strategy is a convex combination of the costs and risks of the deterministic control options that are “mixed” in the mixed strategy. In Figure 1, the convex hull is the shadowed triangular area. The optimal mixed control strategy is found by inspection at the intersection of the 1% risk line and the bottom edge of the triangle.

B. Related Work

A *mixed strategy*, which chooses an action probabilistically, is an essential component in game theory. This idea has been imported to the operation research community and applied for chance-constrained programming problems [1]. A mixed strategy is also commonly assumed in constrained Markov Decision Process (MDP) researches [2], where it is usually referred to as a mixed policy.

However, such a stochastic decision making has been rarely discussed in the context of control theory. To the best of the authors knowledge, there is no prior work in the domain of CCOC that considers a mixed control strategy. This is probably because, as we shown in Theorem 1 of this paper, a mixed control strategy can outperform pure control strategies only for nonconvex CCOC problems, while the majority of existing researches have focused on convex problems [3], [4], [5], [6], [7], [8]. Although there are relatively less numbers of studies on nonconvex CCOC due to the complexity of the problem, it has important applications such as path planning with obstacles [9], [10] and planetary entry descent and landing [11], [12].

The contribution of this paper is twofold. Firstly, we establish several theorems that are useful to analyze and solve a mixed strategy CCOC problem. In addition to the theorem that a mixed control strategy can result in a less cost than deterministic control strategies when a CCOC problem is nonconvex, we also show that an optimal mixed control strategy only needs to choose from up to two control actions (Theorem 2). Secondly, we develop the first algorithm that optimizes a mixed control strategy for CCOC (Algorithm 1). The algorithm returns a mixed control strategy whose suboptimality is within a user-specified tolerance.

II. PROBLEM STATEMENT

We first review a general problem formulation of a CCOC problem with a deterministic control strategy. We then extend the formulation to allow a mixed control strategy. A deterministic control strategy can be viewed as a special case of mixed control strategies, where a single control input is chosen with probability one. Therefore, we use the extended problem formulation in the following sections in this paper.

A. Formulation with a pure control strategy

Let $u \in \mathcal{X}$ be a control input, where \mathcal{X} is assumed to be a non-empty compact set. We define the control input u in a broad sense; when considering a model-predictive control and a finite-horizon optimal control, u is a *sequence* of control inputs over a finite prediction horizon; when solving a dynamic programming problem, u represents a sequence of *control policy*, which is a map from a state space to \mathcal{X} . Let (Ω, \mathcal{M}, P) be a probability space and $w \in \Omega$ be a disturbance vector. Like u , w may also represent a *sequence* of disturbance spanning over a finite horizon. We consider a set of constraints, $g(u, w) \sim 0$, where $g : \mathcal{X} \subseteq \Omega \times \mathbb{R}^N$ is a constraint function and \sim is the element-wise inequality. A cost function, $f : \mathcal{X} \subseteq \Omega \times \mathbb{R}$, is given. We assume the existence of the expectation of $f(u, w)$ with respect to the

probability measure P . A CCOC problem is to find a control input u that minimizes the expected cost while guaranteeing that the probability of violating the given set of constraints is below the risk bound, $\Delta \in (0, 1]$. A general formulation of the CCOC problem is given as follows:

Problem 1: Pure-Strategy CCOC

$$\min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)] \quad (1)$$

$$\text{s.t.} \quad P(g(u, w) \leftarrow 0) \geq \Delta, \quad (2)$$

where $\mathbb{E}_P[\cdot]$ is an expected value with respect to the probability measure P . In game theory, a deterministic strategy is referred to as a pure strategy. Following this convention, we call a deterministic control strategy as *pure control strategy*. Likewise, we call Problem 1 a *pure-strategy CCOC problem*.

B. Formulation with a mixed control strategy

A pure control strategy u is a regular (i.e., deterministic) variable. In order to represent a mixed control strategy, we now consider u as a random variable. Let $U \subseteq \mathcal{X}$ be a finite or at most countably infinite set, whose elements are represented as:

$$U = \{u_1, u_2, \dots\}.$$

We define an extended probability space, $(U \subseteq \Omega, \mathcal{M}, P')$, whose outcome is denoted by $(u, w) \in U \subseteq \Omega$. The probability measure P' represents the joint probability distribution of the control input and the disturbance. We use a simplified notation of a marginal probability such as $P'(u_i) := P'(\{u_i\} \subseteq \Omega)$. We assume that the control input and the disturbance are independent. Furthermore, we assume that the marginal probability of the disturbance in P' is equal to the one in P :

$$P'(w) = P(w). \quad (3)$$

Intuitively, this simply means that the probability distribution of the disturbance is not influenced by the choice of control strategies.

Let

$$p_i = P'(u_i).$$

In other words, p_i represents the probability that the control input u_i is chosen. Let \mathbf{p} be a vector consisting of all the marginal probabilities, $\mathbf{p} = [p_1, p_2, \dots]$. Then, a mixed control strategy is formally defined as a 2-tuple, $\langle U, \mathbf{p} \rangle$. We denote by \mathcal{X}_M the set of all possible mixed control strategies. Note that $\langle U, \mathbf{p} \rangle \in \mathcal{X}_M$ must satisfy

$$\sum_i p_i = 1, \quad 0 \leq p_i \leq 1, \exists i, \quad (4)$$

since otherwise P' is not an admissible probability measure.

For example, the optimal mixed control strategy of the illustrative example in Section I-A is represented as follows:

$$U = \{u_A, u_C\}, \quad \mathbf{p} = [0.5, 0.5].$$

A pure control strategy can be viewed as a mixed control strategy with U consisting of a single element. For example, the optimal pure control strategy of the illustrative example in Section I-A is described as:

$$U = \{u_B\}, \quad \mathbf{p} = [1].$$

With a mixed control strategy, a CCOC problem is to find the elements of U as well as the marginal probabilities \mathbf{p} that minimizes the expected cost while guaranteeing that the probability of violating a given set of constraints is below the a user-specified risk bound, $\Delta \in (0, 1]$. Intuitively, an optimizer chooses a set of control inputs to be mixed, as well as the mixture ratio (i.e., probability). A CCOC problem with a mixed control strategy is formally stated as follows:

Problem 2: Mixed-Strategy CCOC

$$\min_{(U, \mathbf{p}) \in \mathcal{U}_M} \mathbb{E}_{P'} [f(u, w)] \quad (5)$$

$$\text{s.t.} \quad P'(g(u, w) \leftarrow 0) \geq \Delta. \quad (6)$$

We refer to this problem as a *mixed-strategy CCOC problem*. In the following sections the existence of an optimal solution for Problems 1 and 2 are assumed.

III. THEORY

The ultimate goal of this paper is to develop an algorithm to obtain an optimal mixed control strategy. In this section we derive several important theorems that are necessary to construct such an algorithm.

A. Convexity of a Mixed-Strategy CCOC Problem

Recall that, in the illustrative example in Section I-A, the set of the costs and risks of all the possible mixed control strategies (the shadowed triangle in Figure 1) is the convex hull of the set of the costs and risks of all the pure control strategies (the three black dots in Figure 1). In this subsection we now formally prove this statement.

We denote by \mathcal{F} the set of the costs and risks of all the pure control strategies, that is,

$$\mathcal{F} := \{(\mathbb{E}_P [f(u, w)], P(g(u, w) \leftarrow 0)) \mid u \in \mathcal{X}\}.$$

For example, in the illustrative example in Section I-A, $\mathcal{F} = \{(40, 0.005), (30, 0.01), (10, 0.015)\}$. We also denote by \mathcal{N} the set of the costs and risks of all the mixed control strategies, that is,

$$\mathcal{N} := \{(\mathbb{E}_{P'} [f(u, w)], P'(g(u, w) \leftarrow 0)) \mid (U, \mathbf{p}) \in \mathcal{X}_M\}.$$

We assume that \mathcal{F} and \mathcal{N} are nonempty and compact. This is a sufficient condition for the existence of optimal solutions.

The following theorem holds:

Theorem 1:

$$\mathcal{N} = \text{conv}(\mathcal{F}),$$

where $\text{conv}(X)$ is the convex hull of a set X .

Proof: We denote the risk and the cost of a pure control strategy $u \in \mathcal{X}$ by $\mathbf{z}(u) \in \mathbb{R}^2$ as follows:

$$\begin{aligned} \mathbf{z}(u) &:= (r(u), c(u)) \\ r(u) &:= P(g(u, w) \leftarrow 0) \\ c(u) &:= \mathbb{E}_P [f(u, w)]. \end{aligned}$$

With a slight abuse of notation, we also denote the risk and the cost of a mixed control strategy $(U, \mathbf{p}) \in \mathcal{X}_M$ by $\mathbf{z}(U, \mathbf{p}) := (r(U, \mathbf{p}), c(U, \mathbf{p})) \in \mathbb{R}^2$. Then,

$$\begin{aligned} r(U, \mathbf{p}) &:= P'(g(u, w) \leftarrow 0) \\ &= \sum_{i=1}^{|U|} p_i P'(g(u_i, w) \leftarrow 0) \\ &= \sum_{i=1}^{|U|} p_i P(g(u_i, w) \leftarrow 0) = \sum_{i=1}^{|U|} p_i r(u_i). \end{aligned}$$

Note that the third equality follows from (3). Likewise,

$$\begin{aligned} c(U, \mathbf{p}) &:= \mathbb{E}_{P'} [f(u, w)] \\ &= \sum_{i=1}^{|U|} p_i \mathbb{E}_{P'} [f(u_i, w)] \\ &= \sum_{i=1}^{|U|} p_i \mathbb{E}_P [f(u_i, w)] = \sum_{i=1}^{|U|} p_i c(u_i). \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbf{z}(U, \mathbf{p}) &\in \mathcal{N} \\ &\Rightarrow \mathcal{D}\mathbf{z}(u_1, \mathbf{z}(u_2), \dots, \mathcal{F}, p_1, p_2, \dots) [0, 1] : \\ &\quad \mathbf{z}(U, \mathbf{p}) = \sum_i p_i \mathbf{z}(u_i), \sum_i p_i = 1 \\ &\Rightarrow \mathbf{z}(U, \mathbf{p}) \in \text{conv}(\mathcal{F}). \end{aligned}$$

Hence, $\mathcal{N} = \text{conv}(\mathcal{F})$. ■

This theorem is illustrated in Figure 2, where \mathcal{F} is represented by the area surrounded by the thick bold lines while \mathcal{N} is represented by the shadowed area.

A very important implication of Theorem 1 is that, if a CCOC problem is convex (i.e., \mathcal{F} is convex), then the cost of the optimal pure control strategy coincides with the cost of the optimal mixed control strategy because $\mathcal{F} = \mathcal{N}$. Let c_D^* be the optimal cost function value of a pure-strategy CCOC problem (Problem 1) and c_M^* be the optimal cost function value of the corresponding mixed-strategy CCOC problem (Problem 2). The following lemma holds.

Lemma 1: If \mathcal{F} is convex, then

$$c_D^* = c_M^*.$$

Otherwise,

$$c_D^* \leq c_M^*.$$

In other words, we need to consider a mixed control strategy only when a CCOC problem is nonconvex.

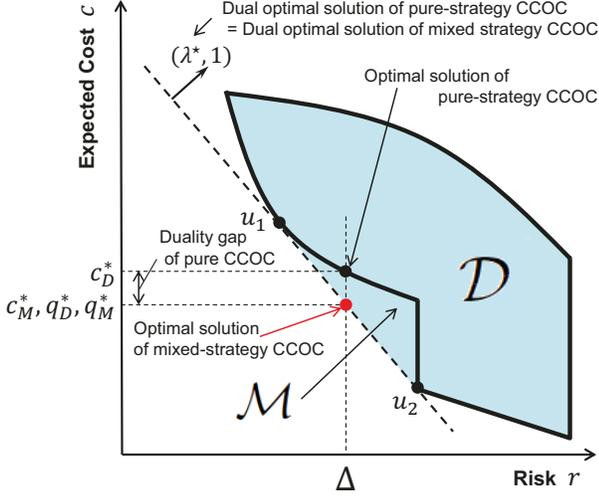


Fig. 2. The cost-risk space of a mixed-strategy CCOC problem, $\text{set}M$, is the convex hull of the cost-risk space of the corresponding pure-strategy CCOC problem, $\text{set}D$ (Theorem 1). Therefore, the optimal mixed control strategy can be obtained by solving the dual optimization problem of the pure-strategy CCOC (Lemmas 2 and 3).

B. Duality of a Mixed-Strategy CCOC

In this subsection we show that an optimal solution to a mixed-strategy CCOC problem is obtained by solving the dual optimization problem of a pure-strategy CCOC.

Let q_D^* be the optimal dual objective value of a pure-strategy CCOC problem, that is,

$$q_D^* := \max_{\lambda} \min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)] + \lambda (P(g(u, w) \leftarrow 0) - \Delta).$$

Likewise, we denote by q_M^* the optimal dual objective value of a mixed-strategy CCOC problem:

$$q_M^* := \max_{\lambda} \min_{(U, \mathbf{p}) \in \mathcal{U}_M} \mathbb{E}'_P [f(u, w)] + \lambda (P'(g(u, w) \leftarrow 0) - \Delta).$$

By the weak duality theorem, $q_D^* \geq c_D^*$. Strong duality holds for a mixed-strategy CCOC problem since it follows from Theorem 1 that \mathcal{N} is convex. Hence, the following lemma holds.

Lemma 2: A mixed-strategy CCOC problem does not have a duality gap, that is,

$$c_M^* = q_M^*.$$

Since $\mathcal{N} = \text{conv}(\mathcal{F})$, the dual optimization problem of a pure-strategy CCOC is equivalent to the dual optimization problem of the mixed-strategy CCOC. Therefore, the following lemma holds.

Lemma 3: The dual optimal solution of a pure-strategy CCOC is equal to the dual optimal solution of the corresponding mixed-strategy CCOC. Hence,

$$q_M^* = q_D^*.$$

These lemmas suggest that an optimal solution to a mixed-strategy CCOC, whose cost is c_M^* , can be computed through the dual solution to a pure-strategy CCOC. This approach allows us to solve a mixed-strategy CCOC problem with almost the same computation cost as a pure-strategy CCOC problem.

Graphical Interpretation: Lemmas 2 and 3 can be geometrically interpreted through the min common/max crossing (MC/MC) framework [13]. The primal optimal solution to the pure-strategy CCOC (i.e., the optimal pure control strategy) is the point that has the least cost among $D \{ \} (r, c) \mid R^2 \|r\| \geq \Delta \}$. In Figure 2, the black dot on $r = \Delta$ represents the optimal pure control strategy. Likewise, the optimal mixed control strategy is represented by the red dot.

In order to graphically interpret the dual optimization problems of the pure-strategy CCOC, consider a nonvertical line that contains \mathcal{F} in its upper closed halfspace. The slope of the line corresponds to λ . The dual optimization problem is to find the maximum crossing point of such a line with the vertical line $r = \Delta$. In Figure 2, the maximum crossing point for \mathcal{F} is represented by the red dot. Note that the maximum crossing point for \mathcal{N} is the same red dot. Furthermore, recall that this red dot also represents the optimal mixed control strategy. Therefore, the optimal mixed control strategy can be obtained by solving the dual optimization problem of the pure-strategy CCOC.

Note that the vertical distance between the black dot and the red dot in Figure 2 corresponds to the duality gap of the pure-strategy CCOC problem. When the pure-strategy CCOC problem is nonconvex, a mixed control strategy can outperform pure control strategies because it eliminates the duality gap by convexifying \mathcal{F} .

C. Structure of an Optimal Mixed Control Strategy

A mixed control strategies can consist of countably infinitely many pure control strategies. However, a practical algorithm can hardly handle such a mixed control strategy. A good news is that we only need to consider mixed control strategies consisting of just two pure control strategies because the following theorem holds.

Theorem 2: An optimal mixed control strategies can be composed of up to two pure control strategies.

Proof: Let $[r_M^*, c_M^*] \in \mathcal{N}$ be an optimal solution to a mixed-strategy CCOC problem. Note that $[r_M^*, c_M^*]$ always lies at the boundary of \mathcal{N} . Also, recall that both \mathcal{F} and \mathcal{N} are compact set in \mathbb{R}^2 and $\mathcal{N} = \text{conv}(\mathcal{F})$. Therefore, by the supporting hyperplane theorem, there exists a line (i.e., one-dimensional linear space), L , that passes through $[r_M^*, c_M^*]$ and contains \mathcal{F} in one of its closed half-spaces. Hence, $[r_M^*, c_M^*] \in \text{conv}(L \{ \mathcal{F} \})$. Since $L \{ \mathcal{F}$ is obviously in a one-dimensional space, it follows from Carathéodory's Theorem that $[r_M^*, c_M^*]$ is in the convex combination of two or fewer points in $L \{ \mathcal{F}$. ■

For example, in Figure 2, the optimal mixed control strategy is composed of the two pure control strategy, u_1

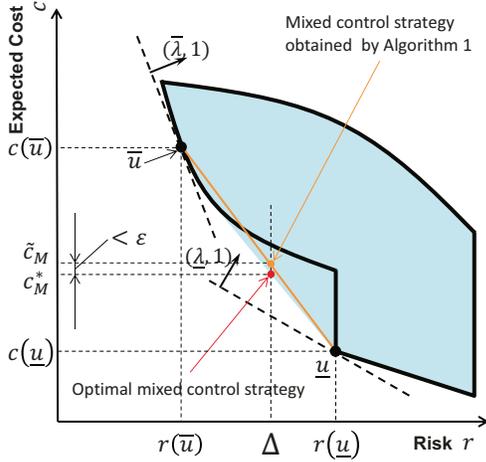


Fig. 3. Graphical interpretation of Algorithm 1.

and u_2 . Note that the supporting hyperplane L in the proof corresponds to the dashed line representing the dual optimal solution (i.e., the solution to the max crossing point problem).

IV. ALGORITHM

A. Overview

As is discussed in Section III-B, we obtain an optimal mixed control strategy by solving the dual of the corresponding pure-strategy CCOC problem. A dual optimal solution is obtained by finding the zero-crossing of the derivative (more precisely, the subgradient) of the dual objective function using a zero-finding method called Brent's method. Like the bisection method, Brent's method starts from an interval that contains the zero-crossing, and iteratively narrows the interval until a certain error measure is below a given tolerance [14].

After the convergence of Brent's method, our algorithm outputs a mixed control strategy consisting of the two pure control strategies corresponding to the both ends of the interval. The probability distribution, \mathbf{p} , is proportional to the difference between the risk of failure of each pure control strategy to the risk bound, Δ .

For example, in Figure 3, the optimal dual solution is λ^* and the optimal mixed control strategy is shown as a red dot. Assume that the zero-finding algorithm converges to an interval $[\bar{\lambda}, \lambda]$, as shown in the figure. The corresponding pure control strategies are \bar{u} and u , which result in risks $r(\bar{u})$ and $r(u)$ while their costs being $c(\bar{u})$ and $c(u)$. The algorithm returns the mixed control strategy shown as the orange dot in the figure, which mixes \bar{u} and u with the probabilities

$$\mathbf{p} = \left[\frac{r(u) - \Delta}{r(u) - r(\bar{u})}, \frac{\Delta - r(\bar{u})}{r(u) - r(\bar{u})} \right].$$

As a result, the risk of the mixed control strategy is exactly Δ . Note that its cost, represented as \tilde{c}_M , is greater than the cost of the optimal mixed control strategy, c_M^* . Our proposed algorithm provides a bound on this solution error (Theorem 3).

B. Review of Dual Optimization

Before presenting the algorithm, we briefly review the theory of dual optimization. A reader who is familiar with dual optimization may skip this subsection. The dual objective function at λ is obtained by minimizing the Lagrangian as follows:

$$q_D(\lambda) = \min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)] + \lambda (P(g(u, w) \leftarrow 0) - \Delta). \quad (7)$$

The dual optimization problem is the maximization problem of this dual objective function. We denote by λ^* the optimal dual value:

$$\lambda^* = \arg \max_{\lambda \geq 0} q_D(\lambda). \quad (8)$$

The maximum is attained when the subgradient of q_D contains zero, that is,

$$0 \in \partial q_D(\lambda^*).$$

An element of the subgradient is obtained as follows:

$$P(g(u, w) \leftarrow 0) - \Delta = r(u(\lambda)) - \Delta \in \partial q(\lambda),$$

where $u(\lambda)$ is the optimal solution to the minimization problem (7). It is known that a dual optimization function is always concave, hence its subgradient is monotonically decreasing. Therefore, $r(u(\lambda)) - \Delta$ has a sole zero-crossing, which corresponds to the dual optimal solution. The proposed algorithm solves the dual optimization problem (8) by finding the zero-crossing using a zero-finding algorithm. Intuitively, this corresponding to finding a control strategy u whose risk, $r(u)$, is exactly Δ .

C. Algorithm

The full description of the algorithm is provided below.

Algorithm 1 Mixed-Strategy Chance Constrained Optimal Control

- 1: $u \rightarrow \arg \min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)]$
 - 2: $\underline{u} \rightarrow u$
 - 3: **if** $r(u) - \Delta \geq 0$ **then**
 - 4: **return** $\langle u, [1] \rangle$
 - 5: **end if**
 - 6: $[\underline{\lambda} \ \bar{\lambda}] \rightarrow [0 \ \lambda^+]$
 - 7: $\bar{u} \rightarrow \arg \min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)] + \lambda^+ P(g(u, w) \leftarrow 0)$
 - 8: **while** $\frac{(r(\bar{u}) - \Delta)(\Delta - r(\underline{u}))}{r(\underline{u}) - r(\bar{u})} (\bar{\lambda} - \underline{\lambda}) > \epsilon$ **do**
 - 9: $\lambda \rightarrow$ Brent's method with $[\underline{\lambda} \ \bar{\lambda}]$
 - 10: $u \rightarrow \arg \min_{u \in \mathcal{U}} \mathbb{E}_P [f(u, w)] + \lambda P(g(u, w) \leftarrow 0)$
 - 11: **if** $r(u) - \Delta = 0$ **then**
 - 12: **return** $\langle u, [1] \rangle$
 - 13: **else if** $r(u) - \Delta < 0$ **then**
 - 14: $\bar{\lambda} \rightarrow \lambda, \bar{u} \rightarrow u$
 - 15: **else**
 - 16: $\underline{\lambda} \rightarrow \lambda, \underline{u} \rightarrow u$
 - 17: **end if**
 - 18: **end while**
 - 19: **return** $\left\langle \bar{u}, \underline{u}, \left[\frac{r(\underline{u}) - \Delta}{r(\underline{u}) - r(\bar{u})}, \frac{\Delta - r(\bar{u})}{r(\underline{u}) - r(\bar{u})} \right] \right\rangle$
-

In the algorithm we denote by $[\underline{\lambda} \ \bar{\lambda}]$ the interval on the dual variable of Brent's method. The optimal control inputs corresponding to the upper and lower bound on the interval are denoted by \underline{u} and \bar{u} .

Lines 1-5 are to separate the special case where the dual optimal solution is $\lambda^* = 0$. In such a case, the risk of the optimal pure control strategy is less than Δ . The optimal mixed strategy deterministically choose this pure control strategy (Line 4).

Line 6 initializes the interval of the zero-finding algorithm with 0 and a sufficiently large value λ^+ that results in a risk smaller than Δ . The pure control strategies corresponding to the both ends of the interval are represented as \underline{u} and \bar{u} .

Lines 8-18 iteratively solves the dual optimization problem of the pure-strategy CCOC through zero-finding. Line 9 computes one step of a zero-finding algorithm, Brent's method, in order to obtain $\lambda \in (\underline{\lambda} \ \bar{\lambda})$. In Line 10, the optimal pure control strategy corresponding to λ is obtained as the minimizer of the Lagrangian,

$$\mathbb{E}_P [f(u, w)] + \lambda(P(g(u, w) \leftarrow 0) - \Delta).$$

Note that Δ is omitted in the algorithm since it is a constant. Lines 11 - 15 updates the interval as well as the corresponding pure control strategies so that the optimal dual solution, λ^* , is within the interval. The algorithm exits from the loop if our error measure,

$$\frac{(r(\underline{u}) - \Delta)(\Delta - r(\bar{u}))}{r(\underline{u}) - r(\bar{u})} (\bar{\lambda} - \underline{\lambda}),$$

is within a given tolerance, $\epsilon > 0$. This choice of error measure allows us to bound the suboptimality of the algorithm, as discussed in the following subsection. Brent's method has a superlinear rate of convergence. Therefore, computation time of the proposed algorithm largely depends on the solution time of the optimization problem in Line 10.

D. Suboptimality Bound

As mentioned previously, Algorithm 1 is suboptimal, meaning that the cost of the mixed control strategy obtained by the algorithm, \tilde{c}_M , is greater than or equal to the cost of the optimal mixed control strategy, c_M^* . The following theorem provides a bound on the suboptimality.

Theorem 3:

$$\tilde{c}_M - c_M^* \geq \epsilon.$$

Proof: If the algorithm exits at either Line 4 or 12, the returned solution is strictly optimal, hence

$$\tilde{c}_M - c_M^* = 0.$$

In such cases, the optimal solution is a pure control strategy.

If the algorithm exits at Line 19, it returns a mixed control strategy. Since the resulting mixed control strategy chooses between \bar{u} and \underline{u} with the probabilities of

$$\left[\frac{r(\underline{u}) - \Delta}{r(\underline{u}) - r(\bar{u})}, \frac{\Delta - r(\bar{u})}{r(\underline{u}) - r(\bar{u})} \right],$$

$$\tilde{c}_M = \frac{r(\underline{u})}{r(\underline{u})} \frac{\Delta}{r(\bar{u})} c(\bar{u}) + \frac{\Delta}{r(\underline{u})} \frac{r(\bar{u})}{r(\bar{u})} c(\underline{u}). \quad (9)$$

Recall that $c_M^* = q_D^*$ (Lemmas 2 and 3), where q_D^* is the dual optimal solution to the pure-strategy CCOC problem. Since the dual optimal solution attains the maximum of the dual objective function,

$$q_D(\bar{\lambda}) = c(\bar{u}) + \bar{\lambda}(r(\bar{u}) - \Delta) \geq q_D^* = c_M^* \quad (10)$$

$$q_D(\underline{\lambda}) = c(\underline{u}) + \underline{\lambda}(r(\underline{u}) - \Delta) \geq q_D^* = c_M^*. \quad (11)$$

The suboptimality bound is obtained as follows.

$$\begin{aligned} \tilde{c}_M - c_M^* &= \frac{r(\underline{u})}{r(\underline{u})} \frac{\Delta}{r(\bar{u})} (c(\bar{u}) - c_M^*) + \frac{\Delta}{r(\underline{u})} \frac{r(\bar{u})}{r(\bar{u})} (c(\underline{u}) - c_M^*) \\ &\geq \frac{r(\underline{u})}{r(\underline{u})} \frac{\Delta}{r(\bar{u})} (\bar{\lambda}(r(\bar{u}) - \Delta) - \Delta) + \frac{\Delta}{r(\underline{u})} \frac{r(\bar{u})}{r(\bar{u})} (\underline{\lambda}(r(\underline{u}) - \Delta) - \Delta) \\ &= \frac{(r(\underline{u}) - \Delta)(\Delta - r(\bar{u}))}{r(\underline{u}) - r(\bar{u})} (\bar{\lambda} - \underline{\lambda}) \geq \epsilon. \end{aligned}$$

The first inequality follows from (10) and (11) and the last inequality follows from the exit condition of the while loop (Line 8). \blacksquare

V. IMPLEMENTATION

A. Mixed-Strategy Chance-constrained Dynamic Programming

Algorithm 1 can be used for any types of chance constrained optimal control problem as long as the optimization problem (7) can be solved. In this paper we present an implementation for a chance-constrained dynamic programming (CCDP). The CCDP algorithm proposed by [12] is build upon a zero-finding-based dual optimization, which is similar to Algorithm 1. Hence, it is straightforward to implement Algorithm 1 for a CCDP. Specifically, there are two changes that must be made on the CCDP algorithm in order to obtain an optimal mixed control strategy. One is to modify the exit condition to the one in Line 8 of Algorithm 1, and the other is to add Line 19. We call the resulting algorithm as *mixed-strategy CCOC*.

B. Application to Path Planning

In this subsection we apply the mixed CCDP to a path planning problem with obstacles. We use the same problem formulation as in [12]. A two-dimensional rectangular state space is discretized into a 100x100 grid. Obstacles are placed as shown in Figure 4. The following dynamics are assumed:

$$\begin{aligned} x_{k+1} &= x_k + u_k + w_k \\ u_k &\geq d_k, \quad w_k \sim \mathcal{U}(0, \sigma^2 I), \end{aligned}$$

where d_k and σ are constant parameters, $\mathcal{U}(0, \Sigma)$ is a zero-mean Gaussian distribution with the covariance matrix Σ , and I is the two-dimensional identity matrix. We set $d_k = 6$ and $\sigma = 1$. The control input and disturbance are also discretized with the same interval as the state variable. The

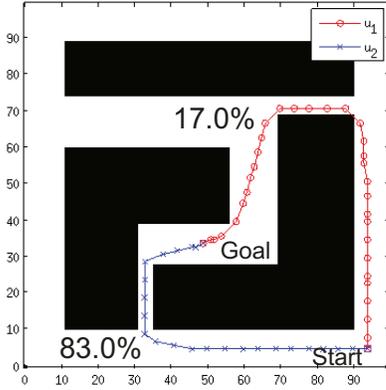


Fig. 4. Application of CCDP to a path planning problem. The optimal mixed control strategy chooses between the two paths, u_1 , u_2 , with probabilities of 17.0% and 83.0%, respectively.

TABLE II
THE COST AND RISK OF THE OPTIMAL PURE AND MIXED CONTROL STRATEGIES FOR THE PATH PLANNING PROBLEM IN FIGURE 4.

Control strategy	Expected path length	Risk
Pure	130.8	0.64%
Mixed	104.2	2.0%

dynamic programming problem is formulated with 50 time steps ($N = 50$). The cost function is the expected length of the resulting path that connects the start and the goal, which are specified as Figure 4.

The simulation result is shown in Figure 4 and Table II. The resulting mixed control strategy consists of two pure control strategies, u_1 and u_2 , which have expected path lengths of 130.8 and 98.7 while the risks of hitting obstacles being 0.64% and 2.28%, respectively. The two pure control strategies are chosen with probabilities of 17.0% and 83.0%, respectively. As a result, the mixed control strategy has a expected path length of 104.2 while the risk of hitting obstacles is exactly 2.0%. On the other hand, the pure-strategy CCOC algorithm results in the same pure control strategy as u_1 , whose expected path length is 130.8. This result shows that, when a chance constraint is imposed, one can minimize the expected path length by choosing a path probabilistically. The computation time of the mixed control strategy is 20.52 seconds while that of the pure control strategy was 20.38 seconds¹. The difference in computation time is small because the only extra computation required to obtain a mixed control strategy is Line 19 of Algorithm 1.

C. Application to Mars Entry, Descent, and Landing

We next apply the mixed-strategy CCDP to a Mars entry, descent, and landing (EDL) scenario [12]. Mars EDL is subject to various source of uncertainties such as atmospheric variability and imperfect aerodynamics model. The resulting dispersions of the landing position typically spans over tens

of kilometers for a 99.9% confidence ellipse [15]. Given such a highly uncertain nature of EDL, a target landing site must be carefully chosen in order to limit the risk of landing on rocky or uneven terrain. At the same time, it is equally important to land near science targets in order to minimize the traverse distance after the landing.

Future Mars lander/rover missions aim to reduce the uncertainty by using several new active control technologies, consisting of the following three stages: entry-phase targeting, powered-descent guidance (PDG) [16], and hazard detection and avoidance (HDA) [17]. Each control stage is capable of making corrections to the predicted landing position by a certain distance, but each stage is subject to execution errors, which deviates the spacecraft away from the planned landing position.

We pose this problem as an optimal sequential decision making under a persisting uncertainty. At the k th control stage, x_k represents the projected landing location without further control. By applying a control at the k th stage, the lander can correct the projected landing location to u_k , which must be within an ellipsoid centered around x_k . At the end of the k th control stage, the projected landing location x_{k+1} deviates from u_k due to a disturbance w_k , which is assumed to have a Gaussian distribution. x_3 is the final landing location. This EDL model is described as follows:

$$x_{k+1} = u_k + w_k$$

$$(u_k \quad x_k)^T D_k (u_k \quad x_k) \geq d_k^2, \quad w_k \succ \mathcal{U}(0, \Sigma_k),$$

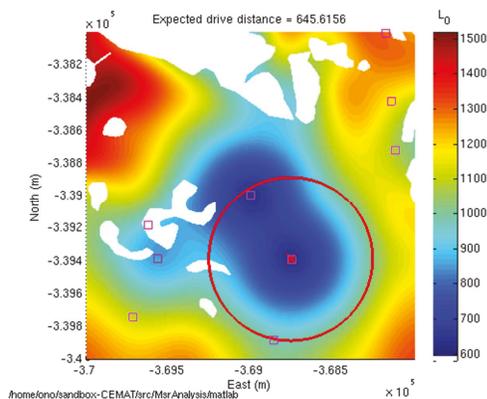
where D_k and Σ_k are positive definite matrices, and d_k is a scalar constant. We use the same parameter settings as [12].

The state space \square is a 2 km-by-2 km square, which is discretized at a one meter resolution. As a result, the problem has four million states at each time step. The control and the disturbance are also discretized at the same resolution. The infeasible areas are specified using the data of HiRISE (High Resolution Imaging Science Experiment) camera on the Mars Reconnaissance Orbiter. We use the real landscape of a site named ‘‘East Margaritifer’’ on Mars.

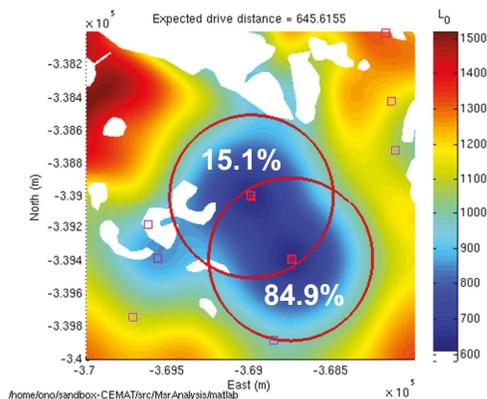
Table III and Figure 5 show the simulation result with a risk bound $\Delta = 0.1\%$. The mixed control strategy obtained by the algorithm chooses between two pure control strategies, \bar{u} and \underline{u} , with the probabilities of 84.9% and 15.1%. The risks of the two pure control strategies are $r(\bar{u}) = 0.016\%$ and $r(\underline{u}) = 0.574\%$ while their costs being $c(\bar{u}) = 645.49$ and $c(\underline{u}) = 641.02$. Note that \bar{u} is equivalent to the optimal pure control strategy shown in Table III. Table III shows that this mixed control strategy results in less cost than the optimal pure control strategy while satisfying the chance constraint. Note that the optimal pure control strategy takes significantly less risk than the risk bound. This is because there is no other solution that is within the risk bound and has less cost. The mixed control strategy improves the cost by mixing this optimal pure control strategy with another pure control strategy that has an excessive risk but a less cost.

It may sound unrealistic to decide a landing site probabilistically. However, consider a situation where there are 1,000

¹Simulations are conducted on a machine with the Intel(R) Xenon(R) X5690 CPU clocked at 3.47GHz and 96GB of RAM



(a) Optimal pure control strategy



(b) Optimal mixed control strategy

Fig. 5. (a) The optimal pure control strategy and (b) the optimal mixed control strategy for the Mars EDL scenario with a risk bound $\Delta = 0.01$. The red \times -marks are the optimal EDL target of the pure control strategies, while the red circles represents 3σ of the disturbance in the first stage, w_D . The mixed strategy chooses between the two pure control strategies with the probabilities of 15.1% and 84.9%.

TABLE III

THE COST AND RISK OF THE OPTIMAL PURE AND MIXED CONTROL STRATEGIES FOR THE MARS EDL SCENARIO.

Control strategy	Expected cost (m)	Risk
Pure	645.49	0.0160%
Mixed	644.81	0.1000%

vehicles and we require 999 of them to land successfully while minimizing the total cost. Then our result means that the optimal strategy is to send 849 of them to the first landing site and 151 of them to the other. When having only one vehicle, the interpretation of this result varies with viewpoint. For a person who knows the result of the coin flip in advance of the landing, the resulting action is no more mixed and hence it may violate the given chance constraint. However, if the result of the coin flip is hidden from the observer, like Schrödinger's cat in a box, then this mixed strategy results in the minimum expected cost while the probability of failure is still within the specified bound.

VI. CONCLUSION

We showed that a mixed control strategy can result in a less cost than deterministic control strategies when a CCOC problem is nonconvex. We also showed that an optimal mixed control strategy only needs to choose from up to two control actions. Building upon these theoretical results, we developed an algorithm that optimizes a mixed control strategy for CCOC with a guarantee that the suboptimality is within a user-specified tolerance.

ACKNOWLEDGEMENT

Yusuke Kobayashi provided help on the proof of Theorem 2. 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, and in part at Keio University.

REFERENCES

- [1] S. P. Mukherjee, "Mixed strategies in chance-constrained programming," *The Journal of the Operational Research Society*, vol. 31, no. 11, pp. 1045–1047, 1980. [Online]. Available: <http://www.jstor.org/stable/2581286>
- [2] E. Altman, *Constrained Markov decision processes*, ser. Stochastic modeling. Chapman & Hall/CRC, 1999. [Online]. Available: <http://books.google.com/books?id=3X9S1NM2iOgC>
- [3] A. T. Schwarm and M. Nikolaou, "Chance-constrained model predictive control," *AICHE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999. [Online]. Available: <http://dx.doi.org/10.1002/aic.690450811>
- [4] D. H. van Hessem, "Stochastic inequality constrained closed-loop model predictive control with application to chemical process operation," Ph.D. dissertation, Delft University of Technology, 2004.
- [5] P. Li, M. Wendt, and G. Wozny, "A probabilistically constrained model predictive controller," *Automatica*, vol. 38, pp. 1171–1176, 2002.
- [6] P. Couchman, M. Cannon, and B. Kouvaritakis, "Stochastic MPC with inequality stability constraints," *Automatica*, vol. 42, no. 12, pp. 2169–2174, 2006.
- [7] M. Ono and B. C. Williams, "Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint," in *Proceedings of 47th IEEE Conference on Decision and Control*, 2008.
- [8] M. Ono, "Joint chance-constrained model predictive control with probabilistic resolvability," in *Proceedings of American Control Conference*, 2012.
- [9] L. Blackmore, H. Li, and B. C. Williams, "A probabilistic approach to optimal robust path planning with obstacles," in *Proceedings of American Control Conference*, 2006.
- [10] M. Ono, L. Blackmore, and B. C. Williams, "Chance constrained finite horizon optimal control with nonconvex constraints," in *Proceedings of American Control Conference*, 2010, pp. 1145–1152.
- [11] Y. Kuwata, M. Pavone, and J. B. Balam, "A risk-constrained multi-stage decision making approach to the architectural analysis of planetary missions," in *Proceedings of the IEEE Conference on Decision and Control*, 2012.
- [12] M. Ono, Y. Kuwata, and J. B. Balam, "Joint chance-constrained dynamic programming," in *Proceedings of the IEEE Conference on Decision and Control*, 2012.
- [13] D. P. Bertsekas, *Convex Optimization Theory*. Athena Scientific, 2009.
- [14] K. E. Atkinson, *An Introduction to Numerical Analysis, Second Edition*. John Wiley & Sons, 1989.
- [15] P. C. Knocke, G. G. Wawrzyniak, B. M. Kennedy, P. N. Desai, T. J. Parker, M. P. Golombek, T. C. Duxbury, and D. M. Kass, "Mars exploration rovers landing dispersion analysis," in *Proceedings of AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, 2004.
- [16] B. Acikmese and S. Ploen, "Convex programming approach to powered descent guidance for mars landing," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1353–1366, 2007.
- [17] A. Johnson, A. Huertas, R. Werner, and J. Montgomery, "Analysis of On-Board Hazard Detection and Avoidance for Safe Lunar Landing," in *IEEE Aerospace Conference*, 2008.