

Distributionally Robust Covariance Steering with Optimal Risk Allocation

Venkatraman Renganathan

Joshua Pilipovsky

Panagiotis Tsiotras

Abstract—This article extends the optimal covariance steering (CS) problem for discrete time linear stochastic systems modeled using moment-based ambiguity sets. To hedge against uncertainty in the state distributions while performing covariance steering, distributionally robust risk constraints are employed during the optimal allocation of the risk. Specifically, a distributionally robust iterative risk allocation (DR-IRA) formalism is used to solve the optimal risk allocation problem for the CS problem using a two-stage approach. The upper-stage of DR-IRA is a convex problem that optimizes the risk, while the lower-stage optimizes the controller with the new distributionally robust risk constraints. The proposed framework results in solutions that are robust against arbitrary distributions in the considered ambiguity set. Finally, we demonstrate our proposed approach using numerical simulations. Addressing the covariance steering problem through the lens of distributional robustness marks the novel contribution of this article.

I. INTRODUCTION

Intelligent and adaptive systems of the “smart world” that work under operational constraints seek to solve some instance of a constrained optimal control problem for optimizing their performance. Such constrained optimal control problems can now be increasingly solved efficiently using several numerical optimization techniques. For instance, robot path planning in uncertain environments [1]–[6] has gained the attention of researchers worldwide as robots are being increasingly deployed to solve many real-world problems. Apart from realistic constraints, reliability of operation of these systems is often thwarted by the ineffective handling of system uncertainties, which can be either deterministic or stochastic.

Control of stochastic systems can be best formulated as a problem of controlling the distribution of trajectories over time subject to constraints. Recently, the finite horizon covariance steering (CS) problem, namely, the problem of steering the first two moments of an initial distribution to the corresponding moments of a final distribution at a specific final time step subject to linear time varying dynamics has been explored [7]. Specifically, the control problem in the CS problem setting involves steering the mean and the covariance to the desired terminal values. When the decision-making process relies blindly on the functional form of the process that models the stochastic uncertainty, it is known to result in potentially severe miscalculation of risk. For instance, Gaussianity assumptions made in the name of

tractability in several modeling regimes are actually rarely justifiable, as the true distribution that governs the uncertain data might be non-Gaussian. Such shortcomings can be mitigated with risk-based stochastic optimization where the risk of wrong decisions can be appropriately handled to result in risk-averse decision making. One such tool is the Distributionally Robust Optimization (DRO) advocated in [8], [9] which enables modelers to explicitly incorporate ambiguity in probability distributions into the optimization problem.

Control of stochastic systems often involves optimizing the system’s objective subject to chance constraints, where one assumes that the system uncertainties follow a known distribution, and one enforces that the constraints hold with high probability as a function of the decision variables. The number of constraint violations, called the total risk budget, is usually a user-defined a priori specification that is a natural metric to assess risk. Hence, one may consider the problem of finding a risk allocation procedure that will allocate the probability of violating each individual chance constraint at each time step. Given a number of chance constraints across a finite horizon, the total risk budget has to be allocated for all chance constraints across all time steps [10]. It is a common practice to consider a uniform risk allocation, i.e., allocate the same risk for all constraints and across all time steps. However, risk allocation can be optimized as in [11], [12] to reduce the conservatism resulting from a uniform risk allocation. If the probability distribution of the system uncertainties is known exactly, then non-uniform risk allocation can be performed effectively. However, risk allocation optimization with arbitrary distribution of the system uncertainties has not yet been explored till now. This article addresses this shortcoming.

Contributions: This article extends the work of [13] to the case of arbitrary distributions using the theory of distributional robustness (DR). To the best of our knowledge, this article is the first one to extend the CS problem using distributionally robust optimization techniques for both polytopic and convex conic state constraint sets. Our main contributions in this article are as follows:

- 1) We extend the covariance steering problem tailored between Gaussian distributions to arbitrary distributions modeled using moment based ambiguity sets.
- 2) We enforce distributionally robust risk constraints for both polytopic and convex cone state constraint satisfaction, while solving the covariance steering problem, and obtain the optimal risk allocation through a distributionally robust iterative risk allocation (DR-IRA)

V. Renganathan is with the Department of Automatic Control LTH, Lund University, Lund, Sweden. J. Pilipovsky and P. Tsiotras are with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150 USA. E-mail: venkatraman.renganathan@control.lth.se, jpilipovsky3@gatech.edu, tsiotras@gatech.edu.

algorithm.

- 3) We demonstrate our approach using simulation examples, and show the effectiveness of the proposed generalization for covariance steering problems between arbitrary distributions in moment-based ambiguity sets.

Following a short summary of notation and preliminaries, the rest of the paper is organized as follows: The main problem statement of distributionally robust covariance steering problem with iterative risk allocation is presented in Section II. Then, the proposed Distributionally Robust Iterative Risk Allocation (DR-IRA) algorithm is discussed in Section III. Subsequently, the proposed approach is demonstrated using simulation results in Section IV. Finally, the paper is concluded in Section V along with directions for future research.

NOTATION AND PRELIMINARIES

The sets of real and natural numbers are denoted by $\mathbb{R}$ and $\mathbb{N}$, respectively. The subset of natural numbers between and including a and b with $a < b$ is denoted by $[a, b]$. An identity matrix of dimension n is denoted by I_n . For a non-zero vector $x \in \mathbb{R}^n$ and a matrix $P \in \mathbb{S}_{++}^n$ (equivalently, $P \succ 0$), we denote $\|x\|_P^2 = x^\top P x$, where $\mathbb{S}_{++}^n$ is the set of all positive definite matrices. We say $P \succ Q$ or $P \succeq Q$ if $P - Q \succ 0$ or $P - Q \succeq 0$ respectively. A probability distribution with mean μ and covariance Σ is denoted by $\mathbb{P}(\mu, \Sigma)$ and, specifically $\mathcal{N}_d(\mu, \Sigma)$, if the distribution is normal in $\mathbb{R}^d$. The cumulative distribution function (cdf) of the normally distributed random variable is denoted by Φ . A uniform distribution over a compact set A is denoted by $\mathcal{U}(A)$. Given $q \geq 1$, the set of probability measures in $\mathcal{P}(\mathbb{R}^d)$ with finite q -th moment is denoted by $\mathcal{P}_q(\mathbb{R}^d) := \{\mu \in \mathcal{P}(\mathbb{R}^d) \mid \int_{\mathbb{R}^d} \|x\|^q d\mu < \infty\}$.

II. PROBLEM FORMULATION

Consider a linear, stochastic, discrete and time-varying system as follows

$$x_{k+1} = A_k x_k + B_k u_k + D_k w_k, \quad k = [0, N-1], \quad (1)$$

where $x_k \in \mathbb{R}^n$ and $u_k \in \mathbb{R}^m$ is the system state and input at time k , respectively, and N denotes the total time horizon. Further, $A_k \in \mathbb{R}^{n \times n}$ is the dynamics matrix, $B_k \in \mathbb{R}^{n \times m}$ is the input matrix and $D_k \in \mathbb{R}^{n \times r}$ is the disturbance matrix. The process noise $w_k \in \mathbb{R}^r$ is a zero-mean random vector that is independent and identically distributed across time. We assume that the system state and the future process noise are independent of each other at all time steps, meaning that $\mathbb{E}[x_k w_j^\top] = 0$, for $0 \leq k \leq j \leq N$. The distribution $\mathbb{P}_w$ of w_k is unknown but is assumed, for all $k \in [0, N-1]$, to belong to a moment-based ambiguity set of distributions, $\mathcal{P}^w$ given by

$$\mathcal{P}^w = \{\mathbb{P}_w \mid \mathbb{E}[w_k] = 0, \mathbb{E}[w_k w_k^\top] = \Sigma_w\}. \quad (2)$$

Note that there are infinitely many distributions in the set $\mathcal{P}^w$. For instance, both multivariate Gaussian and multivariate Laplacian distributions with zero mean and covariance Σ_w

belong to $\mathcal{P}^w$ with the latter having heavier tail than the former.

We assume that the system (1) is controllable under zero process noise meaning that given any $x_0, x_f \in \mathbb{R}^n$, there exists a sequence of control inputs $\{u_k\}_{k=0}^{N-1}$ that steers the system state from x_0 to x_f . Since the system (1) is stochastic, the initial condition x_0 is also subject to an uncertainty model, with the distribution belonging to a moment-based ambiguity set, $\mathbb{P}_{x_0} \in \mathcal{P}^{x_0}$, where

$$\mathcal{P}^{x_0} = \{\mathbb{P}_{x_0} \mid \mathbb{E}[x_0] = \mu_0, \mathbb{E}[(x_0 - \mu_0)(x_0 - \mu_0)^\top] = \Sigma_0\}. \quad (3)$$

Provided that the control law u_k at any time k is selected as an affine function of state history $(x_0, x_1, \dots, x_k)$, similar moment-based ambiguity sets $\mathcal{P}^{x_k}$ can be written for the distribution of states at any time $k \in [1, N]$ using the propagated mean and covariance at time k . Note that $\mathcal{P}^{x_k}$ would *not* be empty at all time steps $k \in [1, N]$ as a Gaussian distribution is a member of that set. This is because, Gaussianity is preserved under linear transformations defined by the dynamics. However, the initial state x_0 need *not* be Gaussian, and thereby it would be *inappropriate* to assume that the final state x_N would also be Gaussian. Hence, the terminal state x_N is subject to a similar uncertainty model as x_0 , with its distribution $\mathbb{P}_{x_N}$ belonging to a moment-based ambiguity set $\mathcal{P}^{x_N}$, given by

$$\mathcal{P}^{x_N} = \{\mathbb{P}_{x_N} \mid \mathbb{E}[x_N] = \mu_f, \mathbb{E}[(x_N - \mu_f)(x_N - \mu_f)^\top] = \Sigma_f\}. \quad (4)$$

The objective is to steer the trajectories of (1) from $x_0 \sim \mathbb{P}_{x_0} \in \mathcal{P}^{x_0}$ to $x_N \sim \mathbb{P}_{x_N} \in \mathcal{P}^{x_N}$ in N time steps. This covariance steering objective is usually achieved by minimizing a cost function under specified convex state and input constraints. We define the concatenated variables required for the problem formulation as follows $\mathbf{X} = [x_0^\top \ x_1^\top \ \dots \ x_N^\top]^\top$, $\mathbf{U} = [u_0^\top \ u_1^\top \ \dots \ u_{N-1}^\top]^\top$, $\mathbf{W} = [w_0^\top \ w_1^\top \ \dots \ w_{N-1}^\top]^\top$, $E_k = [0_{n, kn} \ I_n \ 0_{n, (N-k)n}]$, $\bar{\mathbf{X}} = [\mu_0^\top \ \mu_1^\top \ \dots \ \mu_N^\top]^\top$. Using these concatenated variables, system (1) can be compactly written as

$$\mathbf{X} = \mathcal{A}x_0 + \mathcal{B}\mathbf{U} + \mathcal{D}\mathbf{W}, \quad (5)$$

where the matrices $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{D}$ are of appropriate dimensions containing the time varying system matrices A_k , B_k , and D_k respectively. See [13] for more details on this transformation. It is evident that $\mathbb{P}_{\mathbf{X}}$ is not known exactly but an ambiguity set can be constructed for $\mathbb{P}_{\mathbf{X}}$ from the ambiguity sets $\mathcal{P}^w$ and $\mathcal{P}^{x_0}$ ¹. That is,

$$\mathcal{P}^{\mathbf{X}} := \{\mathbb{P}_{\mathbf{X}} \mid \mathbb{E}[\mathbf{X}] = \bar{\mathbf{X}}, \mathbb{E}[(\mathbf{X} - \bar{\mathbf{X}})(\mathbf{X} - \bar{\mathbf{X}})^\top] = \Sigma_{\mathbf{X}}\}. \quad (6)$$

The cost function to optimize is given as follows

$$J(\mathbf{U}) := \mathbb{E}[\mathbf{X}^\top \bar{Q} \mathbf{X} + \mathbf{U}^\top \bar{R} \mathbf{U}], \quad (7)$$

¹Note that $\mathbb{P}_{\mathbf{X}}$ is the joint distribution formed with $\mathbb{P}_{x_k}$, $k = 0, \dots, N$ being its marginal distributions.

where $\bar{Q} = \text{diag}(Q_0, Q_1, \dots, Q_{N-1})$ is the state penalty matrix and $\bar{R} = \text{diag}(R_0, R_1, \dots, R_{N-1})$ is the control penalty matrix with each $Q_k \succeq 0, R_k \succ 0$ for all $k \in [0, N-1]$.

Over the whole time horizon, we want the states to respect certain state constraints. Since the state is stochastic, constraint violation can be imposed through a chance constraint formulation. Specifically, a distributionally robust risk constraint formulation is preferred in this setting as the system state need not be Gaussian at any time step k . Hence, we impose the following joint distributionally robust risk constraint that limits the worst case probability defined by $\mathbb{P}_{\mathbf{X}}$ of state constraint violation to be less than a pre-specified threshold,

$$\sup_{\mathbb{P}_{\mathbf{X}} \in \mathcal{P}_{\mathbf{X}}} \mathbb{P}_{\mathbf{X}} \left(\bigwedge_{k=0}^N E_k \mathbf{X} \notin \mathcal{X}_p \right) \leq \Delta, \quad (8)$$

where, $\Delta \in (0, 0.5]$ denotes the pre-specified total risk budget and $\mathcal{X}_p$ denotes the state constraint set to be satisfied. The feasible set $\mathcal{X}_p$ is assumed to be a convex polytope and so can be represented using intersection of finite number of half-spaces as

$$\mathcal{X}_p := \bigcap_{i=1}^M \{x \in \mathbb{R}^n \mid a_i^\top x \leq b_i\}, \quad (9)$$

where $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$. Later in this article, we will extend the theory to the case when the feasible set is a more general convex cone constraint.

Problem 1: Given the system (5) and a total risk budget Δ , we seek an optimal feedback control policy $\Pi^* = [\pi_0^*, \dots, \pi_{N-1}^*]$ such that, for all $k \in [0, N-1]$, the control inputs $u_k^* = \pi_k^*(x_0, x_1, \dots, x_k)$, steers the system from $\mathbb{P}_{x_0}$ belonging to (3) to $\mathbb{P}_{x_N}$ belonging to (4), while minimizing the finite-horizon cost function (7) and by respecting the distributionally robust joint risk constraint given in (8).

III. COVARIANCE STEERING WITH DISTRIBUTIONALLY ROBUST RISK ALLOCATION

In this section, we describe how to convert the joint distributionally robust risk constraint into individual distributionally robust risk constraints and use this result to steer both the mean and the covariance of the initial state to the desired terminal mean and covariance.

A. Propagation of Mean and Covariance

We adopt the following control policy from [13] as follows,

$$\mathbf{U} = \mathbf{V} + \mathbf{K}\mathbf{Y}, \quad \text{where} \quad (10)$$

$$\mathbf{Y} = \mathcal{A}(\underbrace{x_0 - \mu_0}_{:=y_0}) + \mathcal{D}\mathbf{W}. \quad (11)$$

Here $\mathbf{V} = [v_0^\top \ v_1^\top \ \dots \ v_{N-1}^\top]^\top \in \mathbb{R}^{Nm}$ and $v_k \in \mathbb{R}^m$ for all $k = [0, N-1]$. The concatenated gain matrix $\mathbf{K} \in \mathbb{R}^{Nm \times (N+1)n}$ contains the individual gain matrices $K_k \in \mathbb{R}^{m \times n}$. The mean and covariance of $\mathbf{Y}$ are given by

$$\bar{\mathbf{Y}} = \mathbb{E}[\mathbf{Y}] = \mathbb{E}[\mathcal{A}y_0 + \mathcal{D}\mathbf{W}] = 0, \quad (12)$$

$$\Sigma_{\mathbf{Y}} = \mathcal{A}\Sigma_0\mathcal{A}^\top + \mathcal{D}\Sigma_{\mathbf{W}}\mathcal{D}^\top, \quad (13)$$

where $\Sigma_{\mathbf{W}} = \text{diag}(\Sigma_w, \dots, \Sigma_w) \in \mathbb{R}^{Nr \times Nr}$. Substituting (12) and (13) into (10), we can infer that

$$\bar{\mathbf{U}} = \mathbb{E}[\mathbf{U}] = \mathbb{E}[\mathbf{V}] + \mathbb{E}[\mathbf{K}\mathbf{Y}] = \mathbf{V}, \quad (14)$$

$$\Sigma_{\mathbf{U}} = \mathbf{K}\Sigma_{\mathbf{Y}}\mathbf{K}^\top. \quad (15)$$

Similarly, substituting (10) and (12) into (5), the dynamics of the state mean $\bar{\mathbf{X}} := \mathbb{E}[\mathbf{X}]$, and the state covariance $\Sigma_{\mathbf{X}} = \mathbb{E}[(\mathbf{X} - \bar{\mathbf{X}})(\mathbf{X} - \bar{\mathbf{X}})^\top]$ can be written as

$$\bar{\mathbf{X}} = \mathcal{A}\mu_0 + \mathcal{B}\mathbf{V}, \quad (16)$$

$$\Sigma_{\mathbf{X}} = (I + \mathcal{B}\mathbf{K})\Sigma_{\mathbf{Y}}(I + \mathcal{B}\mathbf{K})^\top. \quad (17)$$

It is evident from (16) and (17) that the component $\mathbf{V}$ of the control law (10) steers the mean of the system from μ_0 to μ_f and the component $\mathbf{K}$ of the control law (10) steers the covariance from Σ_0 to Σ_f . Note that the initial and the terminal state moments can be expressed as follows

$$\mu_0 = E_0\bar{\mathbf{X}}, \quad \Sigma_0 = E_0\Sigma_{\mathbf{X}}E_0, \quad (18)$$

$$\mu_f = E_N\bar{\mathbf{X}}, \quad \Sigma_f = E_N\Sigma_{\mathbf{X}}E_N. \quad (19)$$

In order to make the problem convex, we relax the terminal covariance constraint in (19) as an inequality constraint $\Sigma_f \succeq E_N\Sigma_{\mathbf{X}}E_N$ and subsequently reformulate it as a linear matrix inequality (LMI) using the Schur complement as

$$\begin{bmatrix} \Sigma_f & E_N(I + \mathcal{B}\mathbf{K})\Sigma_{\mathbf{Y}}^{\frac{1}{2}} \\ \Sigma_{\mathbf{Y}}^{\frac{1}{2}}(I + \mathcal{B}\mathbf{K})^\top E_N^\top & I \end{bmatrix} \succeq 0. \quad (20)$$

Note that the cost given by (7) can be decoupled into the cost on the mean and the cost on the covariance as follows

$$J(\mathbf{V}, \mathbf{K}) = \underbrace{\bar{\mathbf{X}}^\top \bar{Q} \bar{\mathbf{X}} + \bar{\mathbf{U}}^\top \bar{R} \bar{\mathbf{U}}}_{:=J_\mu} + \underbrace{\text{tr}(\bar{Q}\Sigma_{\mathbf{X}} + \bar{R}\Sigma_{\mathbf{U}})}_{:=J_\Sigma}. \quad (21)$$

Substituting (16) and (17), we get

$$J(\mathbf{V}, \mathbf{K}) = (\mathcal{A}\mu_0 + \mathcal{B}\mathbf{V})^\top \bar{Q} (\mathcal{A}\mu_0 + \mathcal{B}\mathbf{V}) + \mathbf{V}^\top \bar{R} \mathbf{V} + \text{tr} \left[((I + \mathcal{B}\mathbf{K})^\top \bar{Q} (I + \mathcal{B}\mathbf{K}) + \mathbf{K}^\top \bar{R} \mathbf{K}) \Sigma_{\mathbf{Y}} \right]. \quad (22)$$

B. Distributionally Robust Polytopic Joint Risk Constraints

Given that the state constraint set $\mathcal{X}_p$ is assumed to be a convex polytope, the worst case joint probability of violating any of the M state constraints over the horizon N given by (8) can be equivalently written as

$$\sup_{\mathbb{P}_{\mathbf{X}} \in \mathcal{P}_{\mathbf{X}}} \mathbb{P}_{\mathbf{X}} \left(\bigwedge_{k=1}^N \bigwedge_{i=1}^M a_i^\top E_k \mathbf{X} > b_i \right) \leq \Delta. \quad (23)$$

Using Boole's inequality, the above joint distributionally robust risk constraint can be decomposed into individual distributionally robust risk constraints at each time step

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} (a_i^\top E_k \mathbf{X} > b_i) \leq \delta_{i,k}, \quad (24)$$

with $\delta_{i,k}$ denoting the individual risk bound² representing the worst case probability of violating the i^{th} state constraint

²The first and second subscript in $\delta_{i,k}$ denote the constraint defining the state constraint set $\mathcal{X}_p$ and the time step respectively.




Fig. 1: Comparison of tightening constant for the Gaussian case and the distributionally robust case using the Cantelli's inequality are shown here for the risk parameter $\delta \in (0, 0.5]$.

at time step k . That is, for each time step $k \in [1, N]$, and for each half-space $i \in [1, M]$ defining the constraint set $\mathcal{X}_p$, we have

$$\inf_{x_k \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k}(a_i^\top E_k \mathbf{X} \leq b_i) \geq 1 - \delta_{i,k}, \quad (25)$$

$$\sum_{k=1}^N \sum_{i=1}^M \delta_{i,k} \leq \Delta. \quad (26)$$

Using Cantelli's inequality [14], the individual distributionally robust risk constraint in (25) can be equivalently reformulated as deterministically tightened convex second-order cone constraint on the state mean as described in [15]. That is,

$$a_i^\top E_k \bar{\mathbf{X}} \leq b_i - \underbrace{\sqrt{\frac{1 - \delta_{i,k}}{\delta_{i,k}}}}_{:= Q(1 - \delta_{i,k})} \left\| \Sigma_Y^{\frac{1}{2}} (I + \mathcal{B}\mathbf{K})^\top E_k^\top a_i \right\|_2, \quad (27)$$

where the DR quantile function $Q(\delta_{i,k}) := \sqrt{\delta_{i,k}/(1 - \delta_{i,k})}$ plays a similar role to that of Φ^{-1} corresponding to the Gaussian case. Note that, $Q(1 - \delta_{i,k})$ is also a monotonically increasing function of the risk, just like Φ^{-1} , as shown in Figure 1 and the deterministic constraint tightening defined using it leads to a stronger tightening than the constant associated with the Gaussian chance constrained tightening. This stronger tightening will ensure that the worst case probability of state constraint violation is satisfied for any arbitrary distribution in the ambiguity set. It is clear from (27) that Problem 1 can now be converted into the following convex programming problem.

Problem 2: Given the system (5) and a total risk budget Δ , we seek an optimal feedback control sequence of inputs $\mathbf{V}^*, \mathbf{K}^*$ that steers the system from the initial state distribution belonging to (3) with first and second moments given by (18) to the final state distribution belonging to (4) with first and second moments as in (19) and (20) by minimizing the finite-horizon cost function (22) and by respecting the DR risk constraint tightening given in (27).

C. Distributionally Robust Risk Allocation Optimization

From Theorem 1 in [13], the optimal cost obtained from solving Problem 2 will be a monotonically decreasing function of the stage risk budget $\delta_{i,k}$. For brevity of notation, we define the vector of all individual risk bounds over the whole time horizon and across all half-spaces defining the state constraint set $\mathcal{X}_p$ as $\boldsymbol{\delta} = [\delta_{1,1} \ \dots \ \delta_{M,N}]^\top \in \mathbb{R}^{MN}$. Recall that in the risk allocation problem, the stage risk budget $\delta_{i,k}$ becomes a decision variable along with $\mathbf{K}$. However, in the distributionally robust risk constraints given by (27), $\delta_{i,k}$ and $\mathbf{K}$ occur in a bilinear form. A better tractable approach would be to concurrently allocate $\delta_{i,k}$ when solving the optimization Problem 2, so as to minimize the total cost given by (22).

1) Two-Stage Optimization Framework: In the following, we show how to optimally allocate the risk across the time steps and across the constraints defining the state constraint set $\mathcal{X}_p$. Following [13], a two-stage optimization framework is proposed here. The upper stage optimization finds the optimal risk allocation $\boldsymbol{\delta}^*$ and the lower stage solves Problem 2 for the optimal controller $\mathbf{U}^*$, equivalently $(\mathbf{V}^*, \mathbf{K}^*)$, given the optimal risk allocation $\boldsymbol{\delta}^*$ from the upper stage.

Let the value of the objective function after the lower stage optimization for a given risk allocation $\boldsymbol{\delta}$ be

$$J^*(\boldsymbol{\delta}) := \min_{\mathbf{V}, \mathbf{K}} J(\mathbf{V}, \mathbf{K}). \quad (28)$$

Then, the upper-stage optimization problem can be formulated as follows:

$$\begin{aligned} & \underset{\boldsymbol{\delta}}{\text{minimize}} && J^*(\boldsymbol{\delta}), \\ & \text{subject to} && \sum_{k=1}^N \sum_{i=1}^M \delta_{i,k} \leq \Delta, \\ & && \delta_{i,k} \geq 0. \end{aligned} \quad (29)$$

Note that (29) is a convex optimization problem, given that the objective function $J^*(\boldsymbol{\delta})$ is convex, and $\Delta \in (0, 0.5]$. Following Theorem 1 in [13], the optimal cost can be reduced with each successive iteration by carefully increasing the risk allocations $\delta_{i,k}$. That is, the risk can be lowered by tightening the constraints that are too conservative, and increased by loosening the constraints that are already active. It now remains to define active and inactive constraints in the context of distributionally robust risk allocation. Note that the distributionally robust risk constraint given by (27) can be equivalently written as

$$\delta_{i,k} \geq \left(1 + \left(\frac{b_i - a_i^\top E_k \bar{\mathbf{X}}^*}{\left\| \Sigma_Y^{\frac{1}{2}} (I + \mathcal{B}\mathbf{K}^*)^\top E_k^\top a_i \right\|_2} \right)^2 \right)^{-1} =: \bar{\delta}_{i,k}. \quad (30)$$

Here the quantity $\bar{\delta}_{i,k}$ represents the true risk experienced by the optimal trajectories, when using $(\mathbf{V}^*, \mathbf{K}^*)$. Clearly, the selected risk need not be equal to the actual risk once the optimization is completed. When $\delta_{i,k} = \bar{\delta}_{i,k}$, we say that the constraint (30) is active, and is inactive otherwise. Solutions are considered good when the true risk is within

a small margin of the allocated risk. Conservative solutions correspond to the case when $\delta_{i,k} > \bar{\delta}_{i,k}$.

2) *Distributionally Robust Iterative Risk Allocation (DR-IRA) Algorithm:* Starting with some feasible risk allocation $\delta^{(j)}$, with j denoting the iteration number, solve Problem 2 to obtain the optimal controller $(\mathbf{V}_{(j)}^*, \mathbf{K}_{(j)}^*)$. It is straightforward to observe that using the above optimal controller at iteration j leads to the optimal mean trajectory $\bar{\mathbf{X}}_{(j)}^*$ respecting the optimal stage risk budget $\delta_{i,k}^{(j)}$. The risk budget is then successively loosened and tightened according to Algorithm 1 as in [13]. This iterative process generates a sequence of risk allocations by continuously lowering the optimal cost.

Algorithm 1 The covariance steering algorithm with DR-IRA

Input: $\delta_k^{(j)} \leftarrow \Delta/(NM), \epsilon, \rho$
Output: $\delta^*, J^*, \mathbf{V}^*, \mathbf{K}^*$
while $|J^* - J_{\text{prev}}^*| > \epsilon$ **do**
 $J_{\text{prev}}^* \leftarrow J^*$
 Solve **Problem 2** with current δ to get $\bar{\delta}$.
 $\hat{N} \leftarrow$ the number of active constraints
 if $\hat{N} = 0$ or $\hat{N} = MN$ **then**
 break
 end if
 for each j^{th} inactive constraint at k^{th} time step **do**
 $\delta_k^{(j)} \leftarrow \rho \delta_k^{(j)} + (1 - \rho) \bar{\delta}_k^{(j)}$
 end for
 $\delta_{\text{res}} \leftarrow \Delta - \sum_{k=1}^N \sum_{j=1}^M \delta_k^{(j)}$
 for each j^{th} active constraint at k^{th} time step **do**
 $\delta_k^{(j)} \leftarrow \delta_k^{(j)} + \delta_{\text{res}}/\hat{N}$
 end for
end while

D. Distributionally Robust Convex Conic Risk Constraints

Feasibility constraints having the form of convex cone are more common than the polytopic constraints in most engineering applications. In the following discussion, we extend our covariance steering formulation to distributionally robust convex cone feasibility constraints. Consider the following convex cone state constraint set

$$\mathcal{X}_c := \{x \in \mathbb{R}^n \mid \|Ax + b\|_2 \leq c^\top x + d\}. \quad (31)$$

We can specify the distributionally robust risk constraint for all time steps $k \in [1, N]$ with conic state constraint set $\mathcal{X}_c$ as

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} [\|Ax_k + b\|_2 \leq c^\top x_k + d] \geq 1 - \delta_k, \quad (32)$$

or

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} [x_k \in \mathcal{X}_c] \geq 1 - \delta_k, \quad (33)$$

$$\sum_{k=1}^N \delta_k \leq \Delta. \quad (34)$$

Notice that (33) is an infinite dimensional constraint and not necessarily convex [13]. Hence, we resort to a convex approximation so that (33) and (34) holds true for all $\Delta \in (0, 0.5]$. We first seek to relax (33) to a DR quadratic risk constraint and then use the reverse union bound approximation as in [16].

1) Relaxing the DR Conic Risk Constraint:

Lemma 1: Given $\delta_k \in (0, 0.5]$ for all $k \in [1, N]$, the following DR quadratic risk constraint

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} [\|Ax_k + b\|_2 \leq c^\top \bar{x}_k + d] \geq 1 - \delta_k, \quad (35)$$

is a relaxation of the original DR conic risk constraint

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} [\|Ax_k + b\|_2 \leq c^\top x_k + d] \geq 1 - \delta_k. \quad (36)$$

Proof: See Appendix 1 of [17]. ■

For each time step $k \in [1, N]$, denote $\psi_{i,k} := a_i^\top x_k + b_i$ and $\kappa_k := c^\top \bar{x}_k + d$ with $a_i^\top$ denoting the i^{th} row of A , and observe that (36) can be equivalently written as

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\left(\sum_{i=1}^n \psi_{i,k}^2 \right)^{\frac{1}{2}} \leq \kappa_k \right] \geq 1 - \delta_k, \quad (37)$$

or

$$\sup_{\mathbb{P}_{x_k} \in \mathcal{P}^{x_k}} \mathbb{P}_{x_k} \left[\sum_{i=1}^n \psi_{i,k}^2 \leq \kappa_k^2 \right] \geq 1 - \delta_k \quad (38)$$

Proposition 1: The DR quadratic constraint (38) is satisfied if the following constraints are satisfied (subscript k dropped for brevity of notation) for some non-negative $f_1, \dots, f_n$ and $\beta_1, \dots, \beta_n$:

$$\sup_{\mathbb{P}_x \in \mathcal{P}^x} \mathbb{P}_x \left[\sum_{i=1}^n |\psi_i| \leq f_i \right] \geq 1 - \beta_i \delta, \quad i = [1, N], \quad (39a)$$

$$\sum_{i=1}^n f_i^2 \leq \kappa^2, \quad (39b)$$

$$\sum_{i=1}^n \beta_i = 1. \quad (39c)$$

Proof: The proof uses the same arguments as in [13] and hence is omitted. ■

2) *Approximation Using Reverse Union Bound:* Note that the constraints given by (39b) can be equivalently written using the mean dynamics given by (16) as

$$\|f_k\|_2 \leq \kappa_k, \quad k = 1, \dots, N, \quad (40)$$

which holds if and only if

$$\|f_k\|_2 \leq c^\top E_k(\mathcal{A}\mu_0 + \mathcal{B}\mathbf{V}) + d, \quad k = 1, \dots, N. \quad (41)$$

Proposition 2: Let $\epsilon_{i,k}^1, \epsilon_{i,k}^2 > 0$ for all $i = 1, \dots, n$ and $k = 1, \dots, N$. Assume that the following convex DR SOC constraints hold true for some $\mathbf{V}, \mathbf{K}$, and $\epsilon_{i,k}^1 + \epsilon_{i,k}^2 \geq 2 -$

$\beta_i \delta_k$.

$$a_i^\top E_k \bar{\mathbf{X}} + \sqrt{\frac{\epsilon_{i,k}^1}{1 - \epsilon_{i,k}^1}} \left\| \Sigma_Y^{\frac{1}{2}} (I + \mathbf{BK})^\top E_k^\top a_i \right\|_2 \leq f_{i,k} - b_i \quad (42a)$$

$$-a_i^\top E_k \bar{\mathbf{X}} + \sqrt{\frac{\epsilon_{i,k}^2}{1 - \epsilon_{i,k}^2}} \left\| \Sigma_Y^{\frac{1}{2}} (I + \mathbf{BK})^\top E_k^\top a_i \right\|_2 \leq f_{i,k} + b_i \quad (42b)$$

Then, the two-sided DR risk constraint given by (39a) holds true as well.

Proof: See Appendix 2 of [17]. ■

It is important to fix a priori $\epsilon_{i,k}^1$ and $\epsilon_{i,k}^2$ in (42) to avoid bilinearity in the decision variables. Thus, the relaxed DR cone risk constraints given by (35) was approximated using (41) and the alternate approximations (42a) and (42b) of the two-sided DR risk constraint (39a). Note that the constraints (41), (42a) and (42b) are convex and can be solved using the standard SDP solvers. In terms of computational complexity, notice that there are $2n + 1$ SOC constraints per time step and thus $(2n + 1)N$ SOC constraints in total.

IV. NUMERICAL SIMULATIONS


In this section, we demonstrate the proposed approach using two examples. One dealing with spacecraft proximity operation as in [13], and the second one using a simple double integrator based path planning as in [7]. The code is available at <https://github.com/venkatramanrenganathan/Covariance-Steering-With-Optimal-DR-Risk-Allocation>.

A. Double Integrator Path Planning


1) *Simulation Setup:* We consider a path-planning problem for a vehicle modelled using the following time invariant, stochastic double integrator system dynamics:

$$A = \begin{bmatrix} I_2 & \Delta t I_2 \\ 0_{2 \times 2} & I_2 \end{bmatrix}, \quad B = \begin{bmatrix} (\Delta t)^2 I_2 \\ \Delta t I_2 \end{bmatrix}, \quad D = 10^{-3} I_4.$$

We assume $\mu_0 = [-10 \ 1 \ 0 \ 0]$ and $\Sigma_0 = \text{diag}(0.1, 0.1, 0.01, 0.01)$ and the discretization time step to be $\Delta t = 0.2$, with horizon $N = 15$. We wish to steer the distribution from the above initial state to the final mean $\mu_f = 0$ with final covariance $\Sigma_f = 0.25 \Sigma_0$, while minimizing the cost function with penalty matrices $Q = \text{diag}(10, 10, 1, 1)$ and $R = 10^3 I_2$. The state constraints are defined as $0.2(x-1) \leq y \leq -0.2(x-1)$. We impose the joint probability of failure over the whole horizon to be $\Delta = 0.10$, which implies that the worst case probability of violating *any* state constraint over the whole horizon is less than 10%. For the Monte Carlo trials, the disturbances were sampled from a multivariate Laplacian distribution with zero mean and unit covariance. Similarly, the initial state x_0 was sampled from a multivariate Laplacian distribution with mean μ_0 and covariance Σ_0 . For comparison, simulations were performed with both Gaussian and DR chance constraints.



(a) Gaussian chance constraints.



(b) Distributionally-robust chance constraints.

Fig. 2: Comparison of 100 independent Monte Carlo trajectories using Gaussian (a) and DR (b) chance constraints with x_0 and noise history $\{w_k\}$ sampled from multivariate Laplacian distributions.

2) *Results & Discussion:* The results from 500 independent Monte-Carlo trials are shown in Figure 2. Given that the disturbances were sampled from a multivariate Laplacian distribution, it is evident from Figure 2a that some Monte-Carlo trials resulted in a violation of the state constraints as the covariance steering was performed assuming Gaussian disturbances. On the other hand, the distributionally robust risk-constrained covariance steering (Figure 2b) ensured that the total risk budget is respected, despite being more conservative. Moreover, all probabilistic state constraints were satisfied. This shows that, assuming Gaussian chance constraints might potentially lead to severe miscalculation of the risk. It is evident from Figures 3 and 4 that the true risk is always upper bounded by the allocated risk regardless of whether a Gaussian or a DR iterative risk allocation is employed.

B. Spacecraft Proximity Operation

1) *Simulation Setup:* We demonstrate the proposed theory using the spacecraft proximity operations problems in orbit as described in [13], albeit with certain parameter changes. We perform three set of simulations, namely: a) covariance steering with distributionally robust polytopic state risk constraints; b) covariance steering with distributionally robust iterative risk allocation for polytopic state risk constraints; and c) covariance steering with distributionally robust iterative risk allocation for convex conic state risk constraints using reverse union bound approximation. For the case of polytopic state risk constraints, we assume $\mu_0 = [100 \ -120 \ 90 \ 0 \ 0 \ 0]$ and $\Sigma_0 = 0.4 \text{diag}(1, 1, 1, 0.1, 0.1, 0.1)$. We wish to steer the distri-




Fig. 3: Allocated and true risks with Gaussian chance constraints.

bution from the above initial state to the final mean $\mu_f = 0$ with final covariance $\Sigma_f = 0.5 \Sigma_0$, while minimizing the cost function with penalty matrices $Q = \text{diag}(10, 10, 10, 1, 1, 1)$ and $R = 10^3 I_3$. We impose the joint probability of failure over the whole horizon to be $\Delta = 0.15$. As in the previous example, the disturbances were sampled from a multivariate Laplacian distribution with zero mean and unit covariance. Similarly, the initial state x_0 was sampled from a multivariate Laplacian distribution with mean μ_0 and covariance Σ_0 . For the cone constraints, we shift the initial x mean to $\mu_0^{(1)} = 10$.

2) *Results & Discussion:* The Monte Carlo results of DR covariance steering are shown in Figure 5, which show the $x-y$ projection of the state trajectories and the 3σ covariance ellipses for three different risk solutions. Notice that the Gaussian solution in Figure 5c remains too close to the boundary, which implies that under non-Gaussian noises, a CS controller based on Gaussian risk constraints may lead to a significant miscalculation of risk. On the other hand, the DR solution in Figure 5b steers away from the boundary to ensure proper constraint satisfaction; but this too is under suboptimal risk placement, from which Figure 5a arises as the optimal trajectories corresponding to the optimal risk budget. Figure 6 shows the monotonically non-increasing trend of the costs corresponding to the polytopic state risk constraints and the convex conic state risk constraints with each iteration of the DR-IRA scheme. This shows that the optimal risk budget resulting from Algorithm 1 has resulted in the minimum possible cost for the DR-CS problem.

The simulation results of covariance steering with distributionally robust iterative risk allocation for convex conic state risk constraints using reverse union bound approximation are shown in Figure 7, with the cost versus the iteration trade-off shown in Figure 6. The state trajectories and their 3σ dispersions remain well within the cone for all time steps and are robust to any zero mean, unit covariance disturbances.

V. CONCLUSION AND FUTURE DIRECTIONS

We have proposed a DR-IRA strategy to optimize the worst case probability of violating the state chance con-





Fig. 4: Allocated and true risks with distributionally-robust chance constraints.

straints at every time step within the CS problem of a linear stochastic system subject to distributionally robust risk constraints. The use of DR-IRA in the context of CS with distributionally robust risk constraints results in optimal solutions that have a true risk much closer to the intended design requirements, compared to the use of a uniform risk allocation. We also extended the approach to quadratic chance constraints in the form of convex cones. Future work will seek to solve the CS problem with output feedback and also extend the problem setting to handle nonlinear dynamics. The problem of steering the first $k > 2$ (higher order) moments [18] is interesting to investigate as well.

Acknowledgment: The work of the first author has been supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under grant agreement No 834142 (Scalable Control). The work of the second and third author has been supported by NASA University Leadership Initiative award 80NSSC20M0163 and ONR award N00014-18-1-2828.

REFERENCES

- [1] B. Luders, M. Kothari, and J. How, “Chance constrained RRT for probabilistic robustness to environmental uncertainty,” in *AIAA Guidance, Navigation, and Control (GNC) Conference*, Toronto, Ontario, Canada, Aug. 2–5, 2010, p. 8160.
- [2] B. D. Luders, S. Karaman, and J. P. How, “Robust sampling-based motion planning with asymptotic optimality guarantees,” in *AIAA Guidance, Navigation, and Control (GNC) Conference*, Boston, MA, Aug. 19–22, 2013, p. 5097.
- [3] T. Summers, “Distributionally robust sampling-based motion planning under uncertainty,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, Oct. 1–5, 2018, pp. 6518–6523.
- [4] V. Renganathan, I. Shames, and T. H. Summers, “Towards integrated perception and motion planning with distributionally robust risk constraints,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 15 530–15 536, Jul. 12–17, 2020.
- [5] W. Han, A. Jasour, and B. Williams, “Non-Gaussian risk bounded trajectory optimization for stochastic nonlinear systems in uncertain environments,” in *International Conference on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, May. 23–27, 2022, pp. 11 044–11 050.



(a) DR chance constraint solution with optimal risk allocation. (b) DR chance constraint solution with uniform risk allocation. (c) Gaussian chance constraint solution with optimal risk allocation.

Fig. 5: Trajectories of 500 independent Monte Carlo samples for distributionally-robust and Gaussian polytopic chance constraints with 3σ covariance ellipses. Each subplot displays the projection of the full state onto the $x - y$ plane.




Fig. 6: Optimal cost for each DR-IRA algorithm iteration for polytope and cone chance constraints.




Fig. 7: Optimal trajectories for 500 independent Monte Carlo trials under DR cone chance constraints.

control for stochastic systems under chance constraints,” *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 266–271, 2018.

- [8] J. Goh and M. Sim, “Distributionally robust optimization and its tractable approximations,” *Operations Research*, vol. 58, no. 4-part-1, pp. 902–917, 2010.
- [9] W. Wiesemann, D. Kuhn, and M. Sim, “Distributionally robust convex optimization,” *Operations Research*, vol. 62, no. 6, pp. 1358–1376, 2014.
- [10] S. Safaoui, B. J. Gravell, V. Renganathan, and T. H. Summers, “Risk-averse *RRT** planning with nonlinear steering and tracking controllers for nonlinear robotic systems under uncertainty,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Prague, Czech Republic, Sep. 27–Oct. 01, 2021, pp. 3681–3688.
- [11] M. Ono and B. C. Williams, “Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint,” in *47th IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 9–11, 2008, pp. 3427–3432.
- [12] M. P. Vitus and C. J. Tomlin, “On feedback design and risk allocation in chance constrained control,” in *50th IEEE Conference on Decision and Control and European Control Conference*, Orlando, FL, Dec. 12–15, 2011, pp. 734–739.
- [13] J. Pilipovsky and P. Tsiotras, “Covariance steering with optimal risk allocation,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, pp. 3719–3733, 2021.
- [14] I. R. Savage, “Probability inequalities of the tchebycheff type,” *Journal of Research of the National Bureau of Standards-B. Mathematics and Mathematical Physics B*, vol. 65, no. 3, pp. 211–222, 1961.
- [15] G. C. Calafiore and L. E. Ghaoui, “On distributionally robust chance-constrained linear programs,” *Journal of Optimization Theory and Applications*, vol. 130, no. 1, pp. 1–22, 2006.
- [16] D. Hunter, “An upper bound for the probability of a union,” *Journal of Applied Probability*, vol. 13, no. 3, pp. 597–603, 1976.
- [17] V. Renganathan, J. Pilipovsky, and P. Tsiotras, “Distributionally robust covariance steering with optimal risk allocation,” 2022. [Online]. Available: https://github.com/venkatramanrenganathan/Covariance-Steering-With-Optimal-DR-Risk-Allocation/blob/main/DR_Covariance_Steering.pdf
- [18] V. Renganathan, N. Hashemi, J. Ruths, and T. H. Summers, “Higher-order moment-based anomaly detection,” *IEEE Control Systems Letters*, vol. 6, pp. 211–216, 2021.

- [6] M. Ono and B. C. Williams, “An efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure,” in *AAAI*, Chicago, IL, Jul. 13–17, 2008, pp. 1376–1382.
- [7] K. Okamoto, M. Goldshtein, and P. Tsiotras, “Optimal covariance