

Constrained Covariance Steering Based Tube-MPPI

Isin M. Balci

Efstathios Bakolas

Bogdan Vlahov

Evangelos A. Theodorou

Abstract—In this paper, we present a new trajectory optimization algorithm for stochastic linear systems which combines Model Predictive Path Integral (MPPI) control with Constrained Covariance Steering (CSS) to achieve high performance with safety guarantees (robustness). Although MPPI can be used to solve complex nonlinear trajectory optimization problems, it may not always handle constraints effectively and its performance may degrade in the presence of unmodeled disturbances. By contrast, CCS can handle probabilistic state and / or input constraints (e.g., chance constraints) and also steer the state covariance of the system to a desired positive definite matrix (control of uncertainty) which both imply that CCS can provide robustness against stochastic disturbances. CCS, however, suffers from scalability issues and cannot handle complex cost functions in general. We argue that the combination of the two methods yields a class of trajectory optimization algorithms that can achieve high performance (a feature of MPPI) while ensuring safety with high probability (a feature of CCS). The efficacy of our algorithm is demonstrated in an obstacle avoidance problem and a circular track path generation problem.

I. INTRODUCTION

Many real-world tasks for autonomous systems can be associated with finite-horizon trajectory optimization problems, which have received significant attention in control engineering and robotics. In the presence of unmodeled disturbances, model uncertainties, and random exogenous inputs from the environment, one has to deal with *stochastic* trajectory optimization problems in which the goal is to find control input sequences or control policies that minimize the expected value of a relevant cost function while satisfying state and/or input constraints with a given confidence level.

In this work, we present a novel algorithm for constrained stochastic trajectory optimization problems subject to safety constraints. Our proposed algorithm combines Constrained Covariance Steering (CCS) theory for discrete-time stochastic linear systems with Model Predictive Path Integral (MPPI) to achieve robustness to uncertainties and variations of the different parameters of the proposed controllers as well as improved performance, scalability, and ability for real-time implementation.

Literature Review: Optimization-based methods treat the stochastic trajectory optimization problem as a nonlinear program (NLP) which can be solved by specialized NLP solvers. However, these NLP based approaches rely on a good initial guess to achieve high performance and may

suffer from the lack of convergence guarantees [1]. Successive convexification-based methods provide convergence guarantees, but they may still suffer from scalability issues, if the underlying system dynamics are stochastic [2], [3], [4].

Dynamic programming-based algorithms have been proposed for unconstrained stochastic trajectory optimization problems to alleviate the scalability issue in [5], [6] and for constrained problems in [7], [8]. These methods, however, cannot guarantee safety and their applicability is limited to smooth objective functions. On the other hand, sampling-based stochastic optimization algorithms deal with non-smooth objective functions but they often cannot handle unmodeled disturbances and model mismatch [9], [10].

The two papers that are most closely related to our approach are [11], [12]. In [11], the authors use Covariance Steering for path planning for linear systems with chance constrained obstacle avoidance. However, the approach in [11] does not scale well due to the fact that the number of decision variables increases quadratically with the problem horizon because of the feedback terms. Furthermore, integer variables that are used to encode obstacle avoidance constraints add to computational complexity. In addition, the quality of the solution depends on the decomposition of the safe region into convex polytopes, which is in general a complex problem. In [12], the authors use unconstrained covariance steering to take sample trajectories from the low cost regions of the state space to enhance the performance and to avoid local optima. However, this method requires the terminal mean and covariance as design parameters, which can be hard to tune, and the safety constraints are not explicitly satisfied.

Main Contributions: This paper presents a novel trajectory optimization algorithm (CCSMPPI) for stochastic linear systems with a non-convex safe state space. CCSMPPI solves the stochastic trajectory optimization problem with guaranteed satisfaction of the safety constraints, which cannot be achieved by standard MPPI. It does so by combining the standard MPPI with the CCS. In particular, MPPI is used to generate a reference trajectory, which is then used to generate a convex safe region, by solving an unconstrained stochastic trajectory optimization problem whereas CCS generates a control policy which minimizes the divergence from the reference trajectory while satisfying the safety constraints. In this way, CCSMPPI endows the MPPI algorithm with robustness to stochastic disturbances by leveraging the framework of CCS. These disturbances may cause the system to diverge from the computed trajectory and may cause violation of safety constraints. This problem is more frequently observed when the cost function is non-smooth such as the sum of indicator functions [10]. CCSMPPI satisfies the probabilistic safety constraints by taking the distribution of disturbances

This research has been supported in part by NSF award CMMI-1937957. I. M. Balci (PhD student) and E. Bakolas (Associate Professor) are with the Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas 78712-1221, USA, Email: isinmerbalci@utexas.edu, bakolas@austin.utexas.edu. B. Vlahov (PhD student) and E. Theodorou (Associate Professor) are with the Department of Aerospace Engineering, Georgia Tech, Atlanta, Georgia, USA, Email: evangelos.theodorou@gatech.edu

into account.

The final improvement of the CCSMPPI over the standard MPPI is the robustness and the safety guarantees against poorly designed cost functions and incorrect tuning of algorithm parameters which are common issues in MPPI. The parameter values and the cost functions that are tuned for one scenario by the designer may not work well in other scenarios and MPPI may return unsafe control inputs. The CCS procedure of the algorithm filters the unsafe inputs that are computed by MPPI and corrects them by means of a feedback control law. This technique makes the cost function design and parameter tuning tasks less time-consuming and allows designers to experiment with cost functions and parameters more freely without compromising safety.

Finally, the practicality of CCSMPPI is demonstrated in 2 different trajectory optimization problems in which we compare our results with those obtained by using the standard MPPI [9] and tube-MPPI [10]. It is shown in numerical experiments that our approach is superior to standard MPPI and tube-MPPI in terms of providing safety against both stochastic disturbances and poorly designed cost functions.

II. PROBLEM STATEMENT AND PRELIMINARIES

A. Notation

We denote by $\mathbb{R}^n$ the set of n -dimensional real vectors, by $\mathbb{R}^+$ the set of non-negative real numbers and by $\mathbb{N}^+$ the set of non-negative integers. We write $\mathbb{E}[\cdot]$ to denote the expectation functional. We denote the probability of the random event E as $\mathbb{P}(E)$. Given a vector x , its 2-norm is denoted by $\|x\|_2$ and given a matrix $\mathbf{A}$, we denote its Frobenius norm by $\|\mathbf{A}\|_F$ if $\mathbf{A}$ is a square matrix, we denote its trace by $\text{tr}(\mathbf{A})$. We use $\mathbf{0}$ and I_n to denote the zero matrix with suitable dimensions and the $n \times n$ identity matrix, respectively. We will denote the convex cone of $n \times n$ symmetric positive semi-definite (symmetric positive definite) matrices by $\mathbb{S}_n^+$ ($\mathbb{S}_n^{++}$). We write $\text{bdiag}(A_1, \dots, A_\ell)$ to denote the block diagonal matrix formed by the matrices A_i , $i \in \{1, \dots, \ell\}$. We use $\cup_{i=1}^N \mathcal{O}_i$ to denote the union of sets $\mathcal{O}_i$ indexed by $i \in \{1, \dots, N\}$. We denote by μ_z and var_z the mean and the variance of a random vector z , respectively, that is, $\mu_z := \mathbb{E}[z]$ and $\text{var}_z := \mathbb{E}[(z - \mu_z)(z - \mu_z)^T] = \mathbb{E}[zz^T] - \mu_z \mu_z^T$. We use $\mathcal{N}(\mu, \Sigma)$ to denote normally distributed random variable with mean μ covariance Σ .

B. Problem Statement

We consider a discrete-time stochastic linear system with dynamics:

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \quad (1)$$

where $A_k \in \mathbb{R}^{4 \times 4}$, $B_k \in \mathbb{R}^{4 \times 2}$, $x_k \in \mathcal{X} \subseteq \mathbb{R}^4$ is the state which is decomposed as $x_k = [p_k^T, v_k^T]^T$, where $p_k = [p_k^x, p_k^y]^T \in \mathbb{R}^2$ is the position and $v_k = [v_k^x, v_k^y]^T \in \mathbb{R}^2$ is the velocity, $u_k = [u_k^x, u_k^y]^T \in \mathbb{R}^2$ is the control input and $w_k \in \mathbb{R}^4$ is the disturbance. We assume that $w_k \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_k)$ and $\mathbb{E}[w_k w_l^T] = \mathbf{0}$ for all $k \neq l$. In our problem formulation, we consider that the obstacles lie in the position space $\mathbb{R}^2$.

The choice of the objective function will be determined by the particular application. However, in this paper, we

use the objective function utilized in the formulation of the information-theoretic MPPI [9], which is defined as follows:

$$\mathcal{L}(X^N, U^{N-1}) = \Phi(x_N) + \sum_{k=0}^{N-1} \left(q(x_k) + \lambda u_k^T R_k u_k \right), \quad (2)$$

where $X^N = \{x_0, x_1, \dots, x_N\}$ is the state sequence, $U^{N-1} = \{u_0, u_1, \dots, u_{N-1}\}$ is the control input sequence, $q : \mathbb{R}^n \rightarrow \mathbb{R}^+$ is the state-dependent term of the running cost function, $\Phi : \mathbb{R}^{2n} \rightarrow \mathbb{R}^+$ is the terminal cost function, $R_k \in \mathbb{S}_n^+$ and $\lambda \in \mathbb{R}^+$. While the term $\lambda u_k^T R_k u_k$ penalizes the control input, $q(x_k)$ and $\Phi(x_N)$ are task dependent and each of them can either be a smooth function as in [9] or a sum of indicator functions for obstacle avoidance as in [10].

We assume that the position space is populated by N_{obs} obstacles. The i -th obstacle is parametrized by its position, $s_i \in \mathbb{R}^n$, and its radius, $r_i \in \mathbb{R}^+$, and the region it occupies is denoted as $\mathcal{O}_i$, where

$$\mathcal{O}_i := \{p \in \mathbb{R}^n \mid \|x - s_i\|_2 \leq r_i\}, \quad i \in \{1, \dots, N_{\text{obs}}\}. \quad (3)$$

The position space is defined as $\mathcal{X} \subseteq \mathbb{R}^n$ and the safe region (or obstacle-free region) of the position space is defined as $\mathcal{X}_{\text{safe}} := \mathcal{X} \setminus \mathcal{O}$ where $\mathcal{O} = \cup_{j=1}^{N_{\text{obs}}} \mathcal{O}_j$. Then, the safe trajectory optimization problem can be formally stated as follows:

$$\begin{aligned} & \underset{X^N, U^{N-1}}{\text{minimize}} && \mathbb{E}[\mathcal{L}(X^N, U^{N-1})] \\ & \text{subject to} && x_0 = \bar{x}_0 \end{aligned} \quad (4)$$

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \quad \forall k \in \mathcal{I}^{N-1} \quad (5)$$

$$\mathbb{P}(x_k \in \mathcal{X}_{\text{safe}}) \geq 1 - \text{P}_{\text{fail}}, \quad \forall k \in \mathcal{I}^N \quad (6)$$

where $\mathcal{I}^N = \{0, 1, \dots, N\}$, $\bar{x}_0 \in \mathbb{R}^{2n}$ and $\text{P}_{\text{fail}} \in (0, 0.5]$ is the acceptable level of the probability of the failure.

It should be highlighted here that since the dynamics of the system include the random noise term w_k , the state sequence X^k will be a random process, and the problem that we will be trying to solve will correspond to an instance of a stochastic trajectory optimization problem.

III. MPPI AND TUBE-MPPI REVIEW

MPPI is a sampling based stochastic Model Predictive Control (MPC) algorithm [9]. It works by taking K samples of control sequences from a Gaussian distribution, and finding the corresponding state trajectories and costs. Each sequence is then weighted by an exponential transform of its cost and the optimal control sequence is found as the weighted sum, written as:

$$\mathbf{u}^{\text{MPPI}} = \frac{1}{\eta} \sum_{i=1}^K \omega_i \mathbf{u}^{(i)}, \quad \eta = \sum_{i=1}^K \omega_i, \quad (7)$$

where $\omega_i = \exp(-\frac{1}{\lambda}(C_i - \beta))$, $\beta = \min_{i=1, \dots, K} C_i$, $\mathbf{u}^{(i)} = \mathbf{v} + \epsilon^{(i)}$ and $\epsilon^{(i)} = [\epsilon_0^{(i)}, \dots, \epsilon_{T_{\text{MPPI}}}^{(i)}]$, $\epsilon_k^{(i)} \sim \mathcal{N}(\mathbf{0}, \nu I)$ and C_i is defined as follows:

$$\begin{aligned} C_i = & \Phi(x_{T_{\text{MPPI}}}^{(i)}) + \sum_{k=0}^{T_{\text{MPPI}}} \left(q(x_k^{(i)}) + \frac{1}{2} v_k^T R_k v_k \right. \\ & \left. + \frac{1}{2} (v_k^T R_k \epsilon_k^{(i)} + (1 - \nu^{-1}) \epsilon_k^{(i)T} R_k \epsilon_k^{(i)}) \right), \end{aligned} \quad (8)$$

C_i is the total path cost induced by i th sample trajectory, $\mathbf{v} = [v_0, \dots, v_{T_{\text{MPPI}}-1}]^T$ is the optimal control sequence obtained by the previous iteration of the MPPI, $\epsilon^{(i)}$ is the control sampling noise, $u^{(i)}$ is the i th control sequence sample and $\nu \in \mathbb{R}^+$ is the control sampling covariance parameter. At the next iteration, the previous optimal control sequence is shifted in time and used as the mean of the Gaussian distribution to sample controls from.

MPPI may not perform well when the system has a disturbance that causes control sampling distribution to only sample high cost trajectories. Tube-MPPI [10] addresses this by running MPPI from two starting states, a nominal and a real state. The real state is taken from the state of the real system as before whereas the nominal state is found by propagating the previous nominal state and control through the noise-free dynamics model without the disturbance term, w_k . Once the MPPI optimization step is done, the optimal trajectory of the real system is pushed to follow the optimal trajectory of the nominal system through the use of a feedback controller. The nominal state and control sequence can then be reset to the real system's when the difference in the free energy of the real and nominal systems is below a user-defined threshold. This means that Tube-MPPI should perform similarly to MPPI in most cases as the nominal and real states will be nearly equivalent but in cases where a large disturbance affects the real system, the nominal system should provide a control sequence unaffected by the disturbance that the real system can then track to get out of a high-cost region.

IV. LINEAR COVARIANCE STEERING THEORY

The main objective of the constrained covariance steering (CCS) problem is to steer the mean and the covariance of a stochastic linear system to desired values while minimizing the expected value of an objective function subject to state and/or input constraints [13], [14]. The general form of the discrete-time CCS problem can be formally stated as follows:

$$\underset{\pi \in \Pi}{\text{minimize}} \quad J(\mathbf{x}, \mathbf{u}) = \mathbb{E} \left[\sum_{k=0}^{N-1} x_k^T Q_k x_k + u_k^T R_k u_k \right] \quad (9a)$$

$$\text{subject to} \quad x_{k+1} = A_k x_k + B_k u_k + w_k, \quad (9b)$$

$$\mathbb{P}[x_k \in \mathcal{X}] \geq 1 - \epsilon_x, \quad \mathbb{P}[u_k \in \mathcal{U}] \geq 1 - \epsilon_u \quad (9c)$$

$$x_0 \sim \mathcal{N}(\mu_0, \Sigma_0), \quad x_N \sim \mathcal{N}(\mu_d, \Sigma_d) \quad (9d)$$

where Π denotes the set of causal policies, $\mathcal{X} \subseteq \mathbb{R}^{n_x}$ and $\mathcal{U} \subseteq \mathbb{R}^{n_u}$ are arbitrary sets corresponding to state and input constraints. The constraints defined in (9d) represents the initial state distribution and the desired terminal state distribution. We should point out that the terminal covariance constraint in (9d) is dropped in our constrained covariance steering formulation since it is not useful to specify a desired covariance for safety as long as constraints in (9c) are satisfied.

Discrete-time formulations of the CCS problems can be cast as finite-dimensional deterministic optimization problems by restricting the class of admissible policies to those which admit the affine state history feedback parametrization [15] or the disturbance feedback parametrization [16] (both

tailored to the covariance steering problem). In this work, we will be utilizing the latter parametrization according to which the control input at each discrete stage can be expressed as follows:

$$v_k = \begin{cases} \bar{v}_0 + H_0(x_0 - \mu_0) & k = 0, \\ \bar{v}_k + H_k(x_0 - \mu_0) + K_{k-1}w_{k-1} & k > 0, \end{cases} \quad (10)$$

where $H_k, K_k \in \mathbb{R}^{n \times n}$ and $\bar{v}_k \in \mathbb{R}^n$ for all $k \in \{0, 1, \dots, N-1\}$. This parametrization allows us to cast the covariance steering problem as a finite-dimensional (deterministic) optimization problem in terms of the following decision variables: $\{\bar{v}_k, H_k, K_k\}_{k=0}^{N-1}$. In order to do that, the decision variables are represented in a more compact form as follows: $\mathcal{H} := [H_0^T, H_1^T, \dots, H_{N-1}^T]^T$, $\mathbf{K} = \text{bdiag}(K_0, K_1, \dots, K_{N-2}, \mathbf{0}^{n \times n})$, $\mathcal{K} := [\mathbf{0} \quad \mathbf{0}]^T$ and $\bar{\mathbf{v}} = [\bar{v}_0^T, \bar{v}_1^T, \dots, \bar{v}_{N-1}^T]^T$.

Now, let $\mathbf{x} = [x_0^T, x_1^T, \dots, x_N^T]^T$, $\mathbf{v} = [v_0^T, v_1^T, \dots, v_{N-1}^T]^T$ and $\mathbf{w} = [w_0^T, w_1^T, \dots, w_{N-1}^T]^T$. Then, it follows from (1) that

$$\mathbf{x} := \mathbf{\Gamma}x_0 + \mathbf{G}_u\mathbf{v} + \mathbf{G}_w\mathbf{w}, \quad (11a)$$

$$\mathbf{v} := \bar{\mathbf{v}} + \mathcal{H}\tilde{\mathbf{x}}_0 + \mathcal{K}\mathbf{w}, \quad (11b)$$

where $\tilde{x}_0 := x_0 - \mu_0$. Equations (11a)-(11b) are derived from (1) and (10) respectively. The reader can refer to [15], [16] for the details of the previous derivations.

We can compute the mean of the vectors $\mathbf{x}$ and $\mathbf{u}$ by taking the expectation of both sides of (11a) and (11b). Then, we compute the deviation of concatenated vectors $\tilde{\mathbf{x}} = \mathbf{x} - \mu_x$ and $\tilde{\mathbf{v}} = \mathbf{v} - \mu_v$. Finally, we compute the variances $\text{var}_x := \mathbb{E}[\tilde{\mathbf{x}}\tilde{\mathbf{x}}^T]$ and $\text{var}_v := \mathbb{E}[\tilde{\mathbf{v}}\tilde{\mathbf{v}}^T]$ as follows:

$$\mu_x = \mathbf{\Gamma}\mu_0 + \mathbf{G}_u\bar{\mathbf{u}}, \quad (12a)$$

$$\text{var}_x = (\mathbf{\Gamma} + \mathbf{G}_u\mathcal{H})\Sigma_0(\mathbf{\Gamma} + \mathbf{G}_u\mathcal{H})^T + (\mathbf{G}_w + \mathbf{G}_u\mathcal{K})\mathbf{W}(\mathbf{G}_w + \mathbf{G}_u\mathcal{K})^T, \quad (12b)$$

$$\mu_v = \bar{\mathbf{v}}, \quad (12c)$$

$$\text{var}_v = \mathcal{H}\Sigma_0\mathcal{H}^T + \mathcal{K}\mathbf{W}\mathcal{K}^T, \quad (12d)$$

where $\mu_0 = \mathbb{E}[x_0]$, $\Sigma_0 = \mathbb{E}[\tilde{x}_0\tilde{x}_0^T]$ and $\mathbf{W} = \text{bdiag}(\mathbf{W}_0, \dots, \mathbf{W}_{N-1})$. Furthermore, the state and the control at the discrete stage k can be recovered from the concatenated state and input vector as $x_k = F_k^x \mathbf{x}$ and $v_k = F_k^v \mathbf{v}$, where F_k^x and F_k^v denote the block matrices whose k th block is equal to the identity matrix and the other blocks are equal to zero. Thus, the mean and the covariance of x_k and v_k are given by:

$$\mu_{x_k} = F_k^x \mu_x, \quad \text{var}_{x_k} = F_k^x \text{var}_x F_k^{xT}, \quad (13a)$$

$$\mu_{v_k} = F_k^v \mu_v, \quad \text{var}_{v_k} = F_k^v \text{var}_v F_k^{vT}. \quad (13b)$$

So, we can express the mean of x_k and v_k as affine functions of the decision variable $\bar{\mathbf{v}}$ whereas the covariance matrices of x_k and v_k can be expressed as convex quadratic functions of decision variables $\mathcal{H}$ and $\mathcal{K}$. This allows us to cast various forms of the (constrained) covariance steering problem as convex optimization problems which can be solved with highly efficient solvers. The reader can refer to [15], [16], [17], [18] for more details.

V. MAIN ALGORITHM

The main components of the algorithm are the MPPI controller, the half-space generator, and the Constrained Covariance Steering module. The MPPI controller solves the unconstrained stochastic trajectory optimization problem and returns a state and an input sequence of length T_{MPPI} . The state sequence generated by the MPPI module is used to generate the half-space constraints. The state and the input sequences and half-space constraints are used in the CCS module to solve for a policy that is guaranteed to be safe with high probability.

Algorithm 1: CCSMPPI

Require: $T_{\text{max}}, T_{\text{CS}}, T_{\text{MPPI}}, \mathcal{M}_{\text{obs}}, \{A_k, B_k, \mathbf{W}_k\}_{k=0, \dots, T_{\text{max}}}, \sigma_{\text{max}}$

```

1  $\bar{x}_0 \leftarrow x_0;$ 
2  $\Sigma_k \leftarrow \mathbf{0};$ 
3 for  $k \in \{0, 1, \dots, T_{\text{max}}\}$  do
4    $X^{T_{\text{MPPI}}}, U^{T_{\text{MPPI}}} \leftarrow \text{MPPI}(\bar{x}_k);$ 
5    $\mathcal{S}_k^{\text{obs}} \leftarrow \text{HSGen}(X^{T_{\text{MPPI}}}, \mathcal{M}_{\text{obs}});$ 
6    $\mu_0 \leftarrow \bar{x}_k; \Sigma_0 \leftarrow \Sigma_k;$ 
7    $\bar{\mathbf{v}}, \mathcal{H}, \mathcal{K} \leftarrow \text{CCS}(\mu_0, \Sigma_0, X^{T_{\text{MPPI}}}, U^{T_{\text{MPPI}}}, \mathcal{S}_k^{\text{obs}});$ 
8    $\bar{u}_k \leftarrow \bar{v}_0, L_k \leftarrow H_0;$ 
9    $u_k \leftarrow \bar{u}_k + L_k(x_k - \bar{x}_k);$ 
10   $\text{SendToActuators}(u_k);$ 
11   $\bar{x}_{k+1} \leftarrow A_k \bar{x}_k + B_k \bar{u}_k;$ 
12   $\Sigma_{k+1} \leftarrow (A_k + B_k L_k) \Sigma_k (A_k + B_k L_k)^T + \mathbf{W}_k;$ 
13  if  $\lambda_{\text{max}}(\Sigma) > \sigma_{\text{max}}$  then
14     $\bar{x}_{k+1} \leftarrow x_k; \Sigma_{k+1} \leftarrow \mathbf{0};$ 
```

A. MPPI

In this paper, we follow the procedures described in [9] to use MPPI algorithm. The MPPI module requires system dynamics and initial state x_0 to sample trajectories. The MPPI horizon T_{MPPI} , the input sampling covariance parameter ν , the number of trajectory samples K are required as algorithm parameters. The input cost matrix $R_k \succ \mathbf{0}$, the state-dependent term of the running cost function $q(x_k)$, and the terminal cost function $\Phi(x)$ are taken as problem data, and they are chosen according to the task at hand.

B. Half-space Generation

Safe half-spaces are generated by the ‘‘HSGen’’ procedure which takes the obstacle information tuple $\mathcal{M}_{\text{obs}}$, the reference state sequence $X^{T_{\text{MPPI}}}$, and the T_{CS} as inputs. It takes the first T_{CS} states of the sequence $X^{T_{\text{MPPI}}}$ and projects the position vectors p_ℓ onto each obstacle $\mathcal{O}_j$. Then, it computes the supporting hyperplane:

$$\mathcal{H}_{\ell,j} := \{p \in \mathbb{R}^2 \mid a_{\ell,j}^T p - b_{\ell,j} = 0\}, \quad (14)$$

at the point of projection such that

$$a_{\ell,j}^T p_\ell - b_{\ell,j} \geq 0 \Rightarrow p_\ell \notin \mathcal{O}_j. \quad (15)$$

The procedure of half-space generation is illustrated in Figure 2. The projection of position p_ℓ at time ℓ onto the obstacle $\mathcal{O}_j$ is denoted as $z_{\ell,j}$ and is defined as follows:

$$z_{\ell,j} := s_j + h_{\ell,j} r_j, \quad (16)$$

where $h_{\ell,j} := (p_\ell - s_j) / \|p_\ell - s_j\|_2$ is the unit normal vector to obstacle $\mathcal{O}_j$ at the point $z_{\ell,j}$ pointing towards p_ℓ . We set $a_{\ell,j} = h_{\ell,j}$ and $b_{\ell,j}$ should satisfy $a_{\ell,j}^T z_{\ell,j} - b_{\ell,j} = 0$. So, we can express $a_{\ell,j}$ and $b_{\ell,j}$ in terms of p_ℓ, s_j and r_j as follows:

$$a_{\ell,j} = \frac{(p_\ell - s_j)}{\|p_\ell - s_j\|_2}, \quad b_{\ell,j} = \frac{(p_\ell - s_j)^T s_j}{\|p_\ell - s_j\|_2} + r_j. \quad (17)$$

The halfspace generation process is repeated for each obstacle $\mathcal{O}_j$ where $j \in \{1, \dots, N_{\text{obs}}\}$ and every time step $\ell \in \{0, 1, \dots, T_{\text{CS}}\}$. The halfspace parameters are gathered in the set of tuples $\mathcal{S}_k^{\text{obs}} = \{(a_{\ell,j}, b_{\ell,j})\}_{\ell=0, \dots, T_{\text{CS}}; j=1, \dots, N_{\text{obs}}}$ to be used in Constrained Covariance Steering.

We should also point out that $p_\ell \notin \mathcal{O}_j$ in Figure 2 but this condition is not necessary for the half-space generation procedure. Even if $p_\ell \in \mathcal{O}_j$, the procedure described by equations in (17) generates a half-space $\mathcal{H}_{\ell,j}$ such that $a_{\ell,j}^T p - b_{\ell,j} \leq 0$ holds for all $p \in \mathcal{O}_j$.

C. Constrained Covariance Steering

The goal of the Constrained Covariance Steering Module is to minimize the deviation of the actual state and control sequence from the reference state and control sequence which is computed by the MPPI algorithm while satisfying the safety constraints. This problem can be formally stated as the following stochastic optimal control problem:

$$\begin{aligned} \underset{\pi \in \Pi}{\text{minimize}} \quad & J(\pi) := \mathbb{E} \left[\sum_{\ell=0}^{T_{\text{CS}}-1} \delta x_\ell^T Q_\ell \delta x_\ell + \delta u_\ell^T R_\ell \delta u_\ell \right. \\ & \left. + \delta x_{T_{\text{CS}}}^T Q_{T_{\text{CS}}} \delta x_{T_{\text{CS}}} \right] \end{aligned} \quad (18a)$$

$$\text{subject to} \quad x_{\ell+1} = A_\ell x_\ell + B_\ell u_\ell + w_\ell, \quad \forall \ell \in \mathcal{I}_t \quad (18b)$$

$$u_\ell = \pi(x_0, \dots, x_\ell), \quad \forall \ell \in \mathcal{I}_t \quad (18c)$$

$$\mathbb{P}[a_{\ell,j}^T p_\ell - b_{\ell,j} \geq 0] \geq 1 - P_{\text{fail}}, \quad \forall \{\ell, j\} \in \mathcal{I} \quad (18d)$$

where Π denotes the set of all admissible control policies, $\delta x_\ell = x_\ell - x_\ell^{\text{MPPI}}, \delta u_\ell = u_\ell - u_\ell^{\text{MPPI}}, \mathcal{I}_t := \{0, \dots, T_{\text{CS}}\}, \mathcal{I}_o := \{1, \dots, N_{\text{obs}}\}, \mathcal{I} = \mathcal{I}_t \times \mathcal{I}_o$.

The stochastic optimal control problem defined in (18a), (18b) and (18d) can be cast as a deterministic optimization problem by fixing the policy as in (10) and concatenating the states $\{x_k\}_{k=0}^N$, the inputs $\{u_k\}_{k=0}^{N-1}$ and the random noise vectors $\{w_k\}_{k=0}^{N-1}$ as explained in Section IV. The resulting finite dimensional deterministic optimization problem is given by:

$$\begin{aligned} \underset{\bar{\mathbf{u}}, \mathcal{H}, \mathcal{K}}{\text{minimize}} \quad & \mathcal{J}(\bar{\mathbf{u}}, \mathcal{H}, \mathcal{K}) := \delta \bar{\mathbf{x}}^T \mathbf{Q} \delta \bar{\mathbf{x}} + \delta \bar{\mathbf{u}}^T \mathbf{R} \delta \bar{\mathbf{u}} \\ & + \text{tr}(\mathbf{Q} \text{var}_{\mathbf{x}}) + \text{tr}(\mathbf{R} \text{var}_{\mathbf{u}}) \end{aligned} \quad (19a)$$

$$\text{subject to} \quad a_{\ell,j}^T \mathbf{P}_\ell \mu_{\mathbf{x}} - b_{\ell,j} \geq \alpha \|\zeta^T \mathbf{P}_\ell^T a_{\ell,j}\|_2, \quad \forall \{\ell, j\} \in \mathcal{I} \quad (19b)$$

$$\zeta = [(\mathbf{G}_0 + \mathbf{G}_u \mathcal{H}) (\mathbf{G}_w + \mathbf{G}_w \mathcal{K})] \mathbf{R} \quad (19c)$$

where $\mathbf{R} \mathbf{R}^T = \text{bdiag}(\Sigma_0, \mathbf{W})$, $\mathbf{x}$ and $\mathbf{u}$ are defined as in Section IV, $\delta \bar{\mathbf{x}} = \mu_{\mathbf{x}} - \mathbf{x}^{\text{MPPI}}, \delta \bar{\mathbf{u}} = \mu_{\mathbf{u}} - \mathbf{u}^{\text{MPPI}}, \mathbf{Q} = \text{bdiag}(Q_0, \dots, Q_{T_{\text{CS}}}), \mathbf{R} = \text{bdiag}(R_0, \dots, R_{T_{\text{CS}}-1})$. $\mathbf{P}_\ell$ is defined such that $p_\ell = \mathbf{P}_\ell \mathbf{x}$ and $\alpha = \varphi^{-1}(1 - P_{\text{fail}})$ where φ




Fig. 1: Main Algorithm Flowchart




Fig. 2: Half-space Generation: Grey area illustrates the region occupied by the obstacle. Red dot shows the reference position p_ℓ at time step ℓ . Black arrow pointing towards p_ℓ shows the unit normal vector $h_{\ell,j} = a_{\ell,j}$. Dashed line and red zone illustrate the supporting hyperplane $\mathcal{H}_{\ell,j} := \{p \in \mathbb{R}^2 \mid a_{\ell,j}^T p - b_{\ell,j} = 0\}$ and restricted unsafe zone.

is the cumulative density function of normal random variable with zero mean and unit variance. Finally, μ_x , μ_u , var_x and var_u are defined as in (12a)-(12d). In addition, we observe that $\text{var}_x = \zeta \zeta^T$ where ζ is given by (19c).

To see the equivalence of optimization problems in (19) and (18), First, observe that the objective function in (18) can be written as $\mathbb{E}[\delta x^T \mathcal{Q} \delta x + \delta u^T \mathcal{R} \delta u]$ where

$$\delta x = x - x^{\text{MPPI}}, \quad \delta u = u - u^{\text{MPPI}}. \quad (20)$$

Using the cyclic permutation property of trace operator, the linearity of expectation and the equalities $\text{var}_{\delta x} = \text{var}_x$ and $\text{var}_{\delta u} = \text{var}_u$, it follows readily that the objective functions in (18a) and (19a) are equivalent.

We use Proposition 1 along with the expressions of $\mu_{p_\ell} = \mathbf{P}_\ell \mu_x$ and $\text{var}_{p_\ell} = \mathbf{P}_\ell \zeta \zeta^T \mathbf{P}_\ell^T$ to show that (19b), (19c) $\Leftrightarrow$ (18d).

Proposition 1. Let $p \sim \mathcal{N}(\mu_p, \Sigma_p)$, where $a, \mu_p \in \mathbb{R}^n$, $b \in \mathbb{R}$, $\Sigma_p \in \mathbb{S}_n^+$, and $P \in (0, 1/2]$. Then, $\mathbb{P}[a^T p - b \geq 0] \geq 1 - P$ if and only if $a^T \mu_p - b \geq \varphi^{-1}(1 - P) \|\mathbf{R}^T a\|_2$

where $\varphi : \mathbb{R} \rightarrow (0, 1)$ is the cumulative density function of normally distributed random variable with zero mean and unit variance and, finally $\mathbf{R}$ is such that $\mathbf{R} \mathbf{R}^T = \Sigma_p$.

Proof.

$$\begin{aligned} \mathbb{P}[a^T p \geq b] &= \mathbb{P}[a^T (p - \mu_p) \geq b - a^T \mu_p] \\ &= \mathbb{P}\left[\frac{a^T (p - \mu_p)}{\sqrt{a^T \Sigma_p a}} \geq \frac{b - a^T \mu_p}{\sqrt{a^T \Sigma_p a}}\right] \end{aligned} \quad (21)$$

Let $z = \frac{a^T (p - \mu_p)}{\sqrt{a^T \Sigma_p a}}$. Observe that $z \sim \mathcal{N}(0, 1)$. Then,

$$\mathbb{P}\left[z \geq \frac{b - a^T \mu_p}{\sqrt{a^T \Sigma_p a}}\right] = \varphi\left(\frac{a^T \mu_p - b}{\sqrt{a^T \Sigma_p a}}\right) \quad (22)$$

Let $\gamma = \frac{a^T \mu_p - b}{\sqrt{a^T \Sigma_p a}}$ for the sake of readability. Thus,

$$\mathbb{P}[a^T p \geq b] \geq 1 - P \Leftrightarrow \varphi(\gamma) \geq 1 - P \quad (23)$$

$$\Leftrightarrow \gamma \geq \varphi^{-1}(1 - P) \quad (24)$$

$$\Leftrightarrow a^T \mu_p - b \geq \alpha \sqrt{a^T \Sigma_p a} = \alpha \|\mathbf{R}^T a\|_2 \quad (25)$$

The implication in (23) comes from derivation of (21) and (22). Since the function $\varphi : \mathbb{R} \rightarrow (0, 1)$ is one-to-one and non-decreasing, the implication in (24) holds. The inequality in (25) is obtained by multiplying both sides of (24) by $\sqrt{a^T \Sigma_p a}$. Finally, we show the equality in (25) using the definition of $\|\cdot\|_2$ and the fact that $\Sigma_p = \mathbf{R} \mathbf{R}^T$ thus complete the proof. $\square$

The problem in (19) has a convex quadratic objective function and the constraint in (19c) is affine. Also, the constraint in (19b) corresponds to a second-order cone constraint since $\alpha = \varphi^{-1}(1 - P_{\text{fail}}) \geq 0$ for all $P_{\text{fail}} \in (0, 0.5]$ [19]. Thus, problem 19 can be solved for global optimal solution $(\bar{v}^*, \mathcal{H}^*, \mathcal{K}^*)$ using off-the-shelf solvers such as [20]. Then, the terms $\bar{v}_0$ and H_0 are recovered and assigned to $\bar{u}_k$ and L_k to be used for the nominal and the real system.

The flowchart of the main algorithm is given in Figure 1. First, the algorithm is initialized by setting $\bar{x}_0 = x_0$

and $\Sigma_0 = \mathbf{0}$ where $\bar{x}_0$ and Σ_0 represent the initial nominal state and initial covariance respectively. Then, using $\bar{x}_k$ as the initial state, MPPI generates a pair of reference state and input sequences $(X^{T_{\text{MPPI}}}, U^{T_{\text{MPPI}}})$. The state sequence is used to generate a safe convex region over which the constraints (18d) are satisfied based on the technique that will be described in Section V-B. Then, we formulate a corresponding CCS problem that seeks for a control policy in the form of (10). This control policy will guarantee collision avoidance while minimizing the deviation from the state and input sequences generated by the MPPI module. If the largest eigenvalue of the computed covariance Σ_{k+1} exceeds a predetermined threshold σ_{max} , then the nominal state $\bar{x}_k$ is set equal to the real state x_k and covariance Σ_k is set to $\mathbf{0}$. Next, the nominal state $\bar{x}_{k+1}$ and covariance matrix Σ_{k+1} will be updated as described in lines 9-10 in Algorithm 1.

D. Discussion

It is worth mentioning that the CCS module uses the disturbance noise covariance in its formulation. However, this information is usually unknown in real-world scenarios. But, this can easily be handled by over-approximating the noise covariance, that is, by taking $\mathbf{W}_k \succeq \mathbf{W}_k^{\text{real}}$, where $\mathbf{W}_k^{\text{real}}$ is the actual noise covariance that is acting on the system, and the previous inequality should be understood in the Loewner partial ordering sense. This allows the CCS module to find a policy that satisfies the safety constraints. Although this approach may generate overly conservative policies, system identification techniques can be used to learn the actual noise covariance [21] and hence reduce conservativeness.

The half-space generation module typically under-approximates the safe region for a time step k . That is, the generated safe set is a subset of the actual safe region. If the real covariance $\mathbf{W}_k^{\text{real}}$ of the disturbance action on the system w_k satisfies the matrix inequality $\mathbf{W}_k \succeq \mathbf{W}_k^{\text{real}}$, then the satisfaction of the constraint (18d) implies that the state will stay in the safe region with probability greater than $1 - P_{\text{fail}}$.

If the trajectory that is returned by the stochastic optimization module violates safety constraints, the half-space generation module might return half-spaces that are very different in consecutive time steps, which makes the CCS problem defined in (19) infeasible. To avoid this potential problem, we use the constrained covariance steering formulation for only the first few time steps of the trajectory optimization algorithm ($T_{\text{max}} > T_{\text{CS}}$). If $\bar{x}_k \in \mathcal{X}_{\text{safe}}$, then the safety violations in the first T_{CS} time steps would be small and the half-spaces between consecutive time steps would be close to each other i.e. $\|a_{\ell,j} - a_{\ell+1,j}\|_2 \leq \varepsilon$ and $\|b_{\ell,j} - b_{\ell+1,j}\|_2 \leq \varepsilon$ for some small $\varepsilon > 0$.

The final component of our algorithm is the use of nominal dynamics, which are the same as real dynamics except it is noise-free. Since the CCS module assures that chance constraints are satisfied with high probability and it uses the nominal state as the initial mean state in its formulation, the nominal state will be safer than the real state. This justifies the use of the nominal state as the initial state in the trajectory optimization module. Also, by computing the covariance in line 10, we compute the high probability region where the real state lies. Then, this covariance value is used

as initial covariance in the CCS procedure in the next step to guarantee the satisfaction of the chance constraints with high probability. This procedure also allows us to use the feedback term computed in the CCS because if μ_0 and Σ_0 are set to x_k and $\mathbf{0}$ respectively, then feedback term $\mathcal{H}$ would be equal to $\mathbf{0}$ and consequently the covariance steering would have little effect.

VI. NUMERICAL EXPERIMENTS

In our numerical experiments, we consider a double integrator with dynamics described by (1) with:

$$A_k = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad B_k = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ dt & 0 \\ 0 & dt \end{bmatrix}, \quad (26)$$

$dt = 0.05$, and $\{w_k\}$ is taken to be a white noise process with $w_k \sim \mathcal{N}(\mathbf{0}, \mathbf{W}_k)$ where the noise covariance matrix $\mathbf{W}_k$ varies depending on different problem instances. We show the efficacy of our approach in two trajectory optimization problems: an obstacle avoidance problem and a path generation problem in a circular track.

Obstacle Avoidance: In the obstacle avoidance case, we compare the performances of CCSMPPI with tube-MPPI [10] under high noise that is acting upon the system to show the robustness of our approach against stochastic disturbances. In the tube-MPPI formulation, an LQG tracking controller is used to track nominal state and input sequences. In our experiments, the LQG cost function parameters $Q_k^{\text{LQG}}, R_k^{\text{LQG}}$ are chosen to be equal to the cost function parameters used in the CCS formulation given in (18a). Also, the failure parameter P_{fail} that is shown in (18d) is taken to be 0.01.

To show that CCSMPPI guarantee safety against stochastic disturbances, we consider the running cost function $q(x_k)$ to be equal to q_{hard} which is defined as follows:

$$q_{\text{hard}}(p_k) = \|p_k - p_{\text{des}}\|_2^2 + 5000 \sum_{j=1}^{N_{\text{obs}}} \mathbb{I}_{\mathcal{O}_j}(p_k) \quad (27)$$

and $\Phi_{\text{hard}}(x_T) = 0$ where $\mathbb{I}_{\mathcal{O}_j} : \mathbb{R}^2 \rightarrow \{0, 1\}$ is the indicator function of set $\mathcal{O}_j$.

The parameters of the MPPI algorithm used in the experiments which produce the results shown in Figure 3 and Figure 4, are $T_{\text{MPPI}} = 40$, $K = 100$, $\lambda = 0.1$, $\nu = 0.1$ and $\epsilon_k \sim \mathcal{N}(\mathbf{0}, 0.001I)$. In addition, the problem horizon parameter $T_{\text{max}} = 200$ and the noise covariance matrix $\mathbf{W}_k = \text{bdiag}(0., 0., 5.0, 5.0)$. In these experiments, the state-dependent term of the running cost function was taken to be $q(x_k) = 10 q_{\text{hard}}(x_k)$ and the desired final position $p_{\text{des}} = [2.0, 10.0]^T$.

Figure 3 illustrates 10 randomly sampled trajectories induced by the CCSMPPI algorithm. Although the intensity of the noise that is acting upon the system is quite high compared to the sampling distribution parameter ν , the CCSMPPI is successfully avoiding obstacles. Figure 4 shows 10 randomly sampled trajectories of the system running under the tube-MPPI algorithm. It can be seen that the agent reaches the goal position but fails to avoid obstacles even though $q_{\text{hard}}(x_k)$ is used as running cost, and the




Fig. 3: Trajectories induced by CCSMPPI




Fig. 4: Trajectories induced by Tube-MPPI

trajectories that collide with obstacles are heavily penalized using indicator functions. In this case, tube-MPPI fails to handle uncertain disturbances and causes collision with the obstacles.

Circular Track: In this scenario, the goal is to keep the position of the system in a circular track with inner radius $R_{in} = R_c - 0.125$ and outer radius $R_{out} = R_c + 0.125$ where $R_c = 2$ while maintaining a desired speed v_{des} in counter-clockwise direction. In the numerical experiments, this goal is encoded 2 different state-dependent running cost functions $q_{t,s}(x_k)$ and $q_{t,h}(x_k)$ which are defined as follows:

$$q_{t,s}(x_k) := (\|v_k\|_2 - v_{des})^2 + 2\|(p_k^x v_k^y - v_k^x p_k^y) - R_c v_{des}\| + 100 \left(\sqrt{p_k^x{}^2 + p_k^y{}^2} - R_c \right)^2 \quad (28a)$$

$$q_{t,h}(x_k) := (\|v_k\|_2 - v_{des})^2 + \|(p_k^x v_k^y - v_k^x p_k^y) - R_c v_{des}\| + 5000 \mathbb{I}_{-C}(p_k) \quad (28b)$$

where $C := \{p \in \mathbb{R}^2 \mid R_c - 0.125 \leq \|p\|_2 \leq R_c + 0.125\}$, $\mathbb{I}_{-C} : \mathbb{R}^2 \rightarrow \{0, 1\}$ is the indicator function such that $\mathbb{I}_{-C}(p) = 1$ if $p \notin C$ and $\mathbb{I}_{-C}(p) = 0$ if $p \in C$.

The safety criterion in this example is to stay within the circular track which is formally defined by set C . This condition is encoded in $q_{t,s}(x_k)$ in (28a) by penalizing the deviation of the position p_k from the mid radius R_c and increasing its weight. This choice makes $q_{t,s}(x_k)$ a smooth function. On the other hand, the safety criterion is encoded




Fig. 5: CCSMPPI, Tube-MPPI and standard MPPI in different scenarios

by using indicator functions in (28b) which is a more clear encoding of the safety constrained however non-smoothness of $q_{t,h}(x_k)$ makes this problem harder to solve.

In Figure 5, trajectories generated by CCSMPPI, Tube-MPPI and standard MPPI are shown from the top to the bottom. Figures 5a, 5c, 5e show the results with running cost $q(x_k) = 100 q_{t,s}(x_k)$ and Figures 5b, 5d, 5f show the results with $q(x_k) = 100 q_{t,h}(x_k)$. The parameters of MPPI algorithm are given as: $T_{MPPI} = 20$, $K = 200$, $\lambda = 0.1$, $\nu = 1.0$, and $\epsilon \sim \mathcal{N}(0, 0.001I)$. In addition, the problem horizon $T_{max} = 300$, $\mathbf{W}_k = \text{bdiag}(0.001, 0.001, 1.0, 1.0)$ and $v_{des} = 6.0$. When $q_{t,s}(x_k)$ is used as a running cost function, the trajectories induced by tube-MPPI and standard MPPI fail to meet the safety criterion (shown in Figures 5c, 5e) due to the poor design of cost function. On the other hand, it is shown in Figure 5a that CCSMPPI manages to keep the position within the track. When the running cost is switched to $q_{t,h}(x_k)$, the non-smoothness of the cost function causes standard MPPI to fail as shown in Figure 5f. Even though tube-MPPI seems to keep the position within the circular track in Figure 5d, there are more violations of the safety constraints than CCSMPPI as is shown in Figure 5b.

In Table I, we compare the performance of CCSMPPI,

tube-MPPI and standard MPPI statistically by sampling $N_{\text{sim}} = 15$ trajectories for both experiments #1 and #2. The running cost function is taken as $100q_{t,s}(x_k)$ and $100q_{t,h}(x_k)$ in experiments #1 and #2, respectively. Also, T_{max} is taken as 200 and 300 in experiments #1 and #2, respectively. In both experiments, $\mathbf{W}_k$ are chosen to be equal to $\text{bdiag}(0.005, 0.005, 0.5, 0.5)$ and $v_{\text{des}} = 6.0$. $\mathbf{Pr}_{\text{fail}}$ represents the probability of failure and it is computed by dividing the number of trajectories that leave the circular track at least once (N_{fail}) by total number of trajectories N_{sim} .

It can be seen from the results of experiment #1 in Table I that standard MPPI performs better in terms of minimizing the cost than both tube-MPPI and CCSMPPI and reaches higher speeds. However, this is due to the poor design of the cost function, and the fact that the control inputs that are corrected by CCS module to guarantee safety are not optimal with respect to the used cost function. When $q_{t,h}(x_k)$ is used as the running cost in experiment #2, standard MPPI performs worse than both tube-MPPI and CCSMPPI due to the presence of random noise w_k . In these experiments, the safety of the trajectory is the first priority, as encoded in the running cost $q_{t,h}(x_k)$. Although tube-MPPI reaches higher speeds, it fails to reach the safety levels of CCSMPPI. Thus, we can conclude that CCSMPPI is superior to standard MPPI and tube-MPPI in terms of minimizing safety violations.

It should be highlighted that the probability of violating the constraint in (18d) at every time step k is less than $P_{\text{fail}} = 0.01$ but still greater than 0. This means that as $T_{\text{max}} \rightarrow \infty$, the failure probability of a trajectory approaches 1. This is the reason why $\mathbf{Pr}_{\text{fail}}$ is non-zero for CCSMPPI in both experiments. $\mathbf{Pr}_{\text{fail}}$ can be reduced by lowering the safety threshold P_{fail} , however it is not possible to make it 0 since w_k is assumed to be normally distributed which is unbounded.

Exp. #1	Av. Speed	Max Speed	$\mathbf{Pr}_{\text{fail}}$	Cost
MPPI	2.46 ± 0.31	3.42 ± 0.35	1.0	44.1 ± 7.0
Tube-MPPI	2.37 ± 0.32	2.95 ± 0.36	0.87	47.9 ± 8.2
CCSMPPI	2.33 ± 0.31	3.03 ± 0.37	0.13	58.6 ± 8.5
Exp. #2	Av. Speed	Max Speed	$\mathbf{Pr}_{\text{fail}}$	Cost
MPPI	1.66 ± 0.24	3.53 ± 0.32	1.0	259.6 ± 73.8
Tube-MPPI	1.77 ± 0.26	2.85 ± 0.35	0.67	95.6 ± 20.9
CCSMPPI	1.65 ± 0.27	2.67 ± 0.29	0.07	66.1 ± 8.3

TABLE I: Performance Comparison Statistics

VII. CONCLUSION

In this paper, we presented a novel framework for safe trajectory optimization for stochastic linear systems. Our method mainly consists of three components which are a stochastic optimization algorithm, a convex safe region generator, and a constrained covariance steering algorithm. In particular, we used Model Predictive Path Integral (MPPI) control for stochastic optimization and a projection-based linearization method for the generation of safe convex regions. In addition, we used a Constrained Covariance Steering algorithm based on the affine disturbance feedback parametrization to safeguard against unmodeled noise disturbances that the MPPI algorithm may not always handle satisfactorily. Our numerical simulations have demonstrated

that our approach can guarantee safety against unmodeled noise uncertainties as well as unsafe outputs generated by the stochastic optimization algorithm. In our future work, we plan to extend our proposed framework to trajectory generation problems for uncertain nonlinear systems based on model-free trajectory optimization algorithms while guaranteeing safety by utilizing nonlinear covariance steering algorithms.

REFERENCES

- [1] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear MPC: bridging the gap via the real-time iteration," *International Journal of Control*, vol. 93, no. 1, pp. 62–80, 2020.
- [2] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Açikmeşe, "Convex optimization for trajectory generation," *arXiv preprint arXiv:2106.09125*, 2021.
- [3] Y. Mao, M. Szmuk, and B. Açikmeşe, "Successive convexification of non-convex optimal control problems and its convergence properties," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 3636–3641, IEEE, 2016.
- [4] F. Augugliaro, A. P. Schoellig, and R. D'Andrea, "Generation of collision-free trajectories for a quadcopter fleet: A sequential convex programming approach," in *2012 IEEE/RSJ international conference on Intelligent Robots and Systems*, pp. 1917–1922, IEEE, 2012.
- [5] E. Theodorou, Y. Tassa, and E. Todorov, "Stochastic differential dynamic programming," in *Proceedings of the 2010 American Control Conference*, pp. 1125–1132, IEEE, 2010.
- [6] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the 2005, American Control Conference*, 2005., pp. 300–306, IEEE, 2005.
- [7] Y. Aoyama, G. Boutselis, A. Patel, and E. A. Theodorou, "Constrained differential dynamic programming revisited," *arXiv preprint arXiv:2005.00985*, 2020.
- [8] Z. Xie, C. K. Liu, and K. Hauser, "Differential dynamic programming with nonlinear constraints," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 695–702, IEEE, 2017.
- [9] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1714–1721, IEEE, 2017.
- [10] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. M. Rehg, and E. A. Theodorou, "Robust sampling based model predictive control with sparse objective information," in *Robotics: Science and Systems*, 2018.
- [11] K. Okamoto and P. Tsotras, "Optimal stochastic vehicle path planning using covariance steering," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2276–2281, 2019.
- [12] J. Yin, Z. Zhang, E. Theodorou, and P. Tsotras, "Improving model predictive path integral using covariance steering," *arXiv preprint arXiv:2109.12147*, 2021.
- [13] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part I," *IEEE Transactions on Automatic Control*, vol. 61, no. 5, pp. 1158–1169, 2015.
- [14] A. Hotz and R. E. Skelton, "Covariance control theory," *International Journal of Control*, vol. 46, no. 1, pp. 13–32, 1987.
- [15] E. Bakolas, "Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints," *Automatica*, vol. 91, pp. 61–68, 2018.
- [16] I. M. Balci and E. Bakolas, "Covariance control of discrete-time gaussian linear systems using affine disturbance feedback control policies," *arXiv preprint arXiv:2103.14428*, 2021.
- [17] E. Bakolas, "Covariance control for discrete-time stochastic linear systems with incomplete state information," in *2017 American Control Conference (ACC)*, pp. 432–437, IEEE, 2017.
- [18] J. Pilipovsky and P. Tsotras, "Chance-constrained optimal covariance steering with iterative risk allocation," in *2021 American Control Conference (ACC)*, pp. 2011–2016, IEEE, 2021.
- [19] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [20] MOSEK ApS, *MOSEK Optimizer API for Python 9.2.40*, 2019.
- [21] B. Feng, M. Fu, H. Ma, Y. Xia, and B. Wang, "Kalman filter with recursive covariance estimation—sequentially estimating process noise covariance," *IEEE Transactions on Industrial Electronics*, vol. 61, no. 11, pp. 6253–6263, 2014.