

Discrete-time Optimal Covariance Steering via Semidefinite Programming

George Rapakoulias¹ and Panagiotis Tsiotras²

Abstract— This paper addresses the optimal covariance steering problem for stochastic discrete-time linear systems subject to probabilistic state and control constraints. A method is presented to efficiently attain the exact solution of the problem based on a lossless convex relaxation of the original non-linear program using semidefinite programming. Both the constrained and the unconstrained versions of the problem with either equality or inequality terminal covariance boundary conditions are addressed. We first prove that the proposed relaxation is lossless for all of the above cases. Numerical examples are then provided to illustrate the proposed method. Finally, a comparative study is performed on systems of various sizes and steering horizons to illustrate the advantages of the proposed method in terms of computational resources compared to the state of the art.

I. INTRODUCTION

The Covariance Control (CC) problem for linear systems was initially posed by A. Hotz and R. Skelton in [1]. It was studied in an infinite horizon setting for both continuous and discrete-time systems and the authors provided a parametrization for all linear state feedback controllers that achieve a specified system covariance. Later, the authors in [2] provided analytical solutions for the minimum effort controller that achieves a specified steady-state system covariance in the same setting.

Its finite horizon counterpart, the Covariance Steering (CS) problem, gained attention only recently. Although similar ideas can be traced back in the Stochastic Model Predictive Control literature [3], [4], in the sense that these methods also try to address constraints in the system covariance, they achieve this objective by using conservative approximations or by solving computationally demanding non-linear programs. Covariance Steering theory, on the other hand, offers a more direct approach, often providing tractable algorithms for the solution in real time.

The first formal treatment of the CS problem was provided in [5], [6] for continuous-time systems, by studying the minimum-effort finite horizon covariance steering problem in continuous time. Later, in [7] the author provided a numerical approach for solving the discrete version of the problem with a relaxed terminal covariance boundary condition using semidefinite programming. In [8] the authors introduced a constrained version of the original problem where the state and control vectors are required to remain within specified

bounds in a probabilistic sense. Finally, its connections to Stochastic Model Predictive control were cemented in [9].

The recently developed covariance steering theory has been applied to a variety of problems ranging from path planning for linear systems under uncertainty [10], control of linear systems with multiplicative noise [11], distributed robot control [12], as well as for control of non-linear [13], [14] and non-Gaussian [15], [16] systems. In our previous work [17], we presented a new method of solving the optimal covariance steering problem in discrete time based on an exact convex relaxation of the original non-linear programming formulation of the problem. At the same time, but independently, the authors of [18] used the same relaxation to solve the optimal covariance steering problem with an inequality terminal boundary condition for a system with multiplicative noise.

The contributions of this paper are two-fold. First, we extend our previous results and prove that the proposed lossless convex relaxation presented in [17] also holds under state and control chance constraints, as well as for the case of inequality terminal boundary covariance constraint. The motivation for this extension is straightforward; many practical applications of covariance steering theory require probabilistic constraints to characterize the feasible part of state space or limit the control effort applied to the system. Furthermore, the inequality terminal covariance boundary condition might better reflect the desire to limit the uncertainty of the state, rather than driving it to an exact value. In this paper, we establish that the proposed method can handle all variants of the optimal covariance steering problem for linear systems encountered in the literature. Finally, we show that the proposed method outperforms other approaches for solving the CS problem, such as [7] and [9], by over an order of magnitude in terms of run-times, while also having much better scaling characteristics with respect to the steering horizon and model size.

II. PROBLEM STATEMENT

Let a stochastic, discrete, time-varying system be described by the state space model

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

where $k = 0, 1, \dots, N-1$ denotes the time step, $A_k \in \mathbb{R}^{n \times n}$ is the system matrix, $B_k \in \mathbb{R}^{n \times p}$ is the input matrix and $D_k \in \mathbb{R}^{n \times q}$ is the disturbance matrix. The system's state, input, and stochastic disturbance are denoted by x_k , u_k and w_k , respectively. The first two statistical moments of the state vector are denoted by $\mu_k = \mathbb{E}[x_k] \in \mathbb{R}^n$ and $\Sigma_k = \mathbb{E}[(x_k - \mu_k)(x_k - \mu_k)^T] \in \mathbb{R}^{n \times n}$. We assume that the process noise

¹ Ph.D. student, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, 30332, USA. Email: grap@gatech.edu

² David and Andrew Lewis Chair and Professor, School of Aerospace Engineering, and Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA, 30332, USA. Email: tsiotras@gatech.edu

w_k has zero mean and unitary covariance. The discrete-time finite horizon optimal covariance steering problem can be expressed as the following optimization problem:

$$\min_{x_k, u_k} J = \mathbb{E} \left[\sum_{k=0}^{N-1} x_k^\top Q_k x_k + u_k^\top R_k u_k \right], \quad (2a)$$

such that, for all $k = 0, 1, \dots, N-1$,

$$x_{k+1} = A_k x_k + B_k u_k + D_k w_k, \quad (2b)$$

$$x_0 \sim \mathcal{N}(\mu_i, \Sigma_i), \quad (2c)$$

$$x_N \sim \mathcal{N}(\mu_f, \Sigma_f), \quad (2d)$$

$$\mathbb{P}(x_k \in \mathcal{X}) \geq 1 - \epsilon_1, \quad (2e)$$

$$\mathbb{P}(u_k \in \mathcal{U}) \geq 1 - \epsilon_2. \quad (2f)$$

For the rest of this paper, we will assume that $R_k \succ 0$, $Q_k \succeq 0$ and that A_k is invertible for all $k = 0, 1, \dots, N-1$.

The decision variables for problem (2) are stochastic random variables, rendering it hard to solve using numerical optimization methods. As shown in [17], in the absence of the chance constraints (2e), (2f) this problem is solved optimally with a linear state feedback law of the form

$$u_k = K_k(x_k - \mu_k) + v_k, \quad (3)$$

where $K_k \in \mathbb{R}^{p \times n}$ is a feedback gain that controls the covariance dynamics and $v_k \in \mathbb{R}^p$ is a feedforward term controlling the system mean. Using (3), the cost function can be written, alternatively, in terms of the first and second moments of the state as follows

$$\mathbb{E} \left[\sum_{k=0}^{N-1} x_k^\top Q_k x_k + u_k^\top R_k u_k \right] = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k K_k \Sigma_k K_k^\top) + \mu_k^\top Q_k \mu_k + v_k^\top R_k v_k.$$

If the initial distribution of the state is Gaussian and a linear feedback law as in (3) is used, the state distribution remains Gaussian. This allows us to write the constraints (2c) and (2d) as

$$\mu_0 = \mu_i, \quad \Sigma_0 = \Sigma_i, \quad \mu_N = \mu_f, \quad \Sigma_N = \Sigma_f.$$

In contrast to previous works such as [7] and [10], we choose to keep the intermediate states in the steering horizon as decision variables, handling them in terms of their first and second moments. To this end, we replace (2b) with the mean and covariance propagation equations

$$\mu_{k+1} = A_k \mu_k + B_k v_k, \quad (4a)$$

$$\Sigma_{k+1} = (A_k + B_k K_k) \Sigma_k (A_k + B_k K_k)^\top + D_k D_k^\top. \quad (4b)$$

Omitting the chance constraints (2f) and (2e) for the moment, the problem is recast as a standard non-linear program

$$\min_{\Sigma_k, K_k, \mu_k, v_k} J = \sum_{k=0}^{N-1} \text{tr}(Q \Sigma_k) + \text{tr}(R K_k \Sigma_k K_k^\top) + \mu_k^\top Q \mu_k + v_k^\top R v_k, \quad (5a)$$

such that, for all $k = 0, 1, \dots, N-1$,

$$\Sigma_{k+1} = A_k \Sigma_k A_k^\top + B_k K_k \Sigma_k A_k^\top + A_k \Sigma_k K_k^\top B_k^\top$$

$$+ B_k K_k \Sigma_k K_k^\top B_k^\top + D_k D_k^\top, \quad (5b)$$

$$\Sigma_0 = \Sigma_i, \quad (5c)$$

$$\Sigma_N = \Sigma_f, \quad (5d)$$

$$\mu_{k+1} = A_k \mu_k + B_k v_k, \quad (5e)$$

$$\mu_0 = \mu_i, \quad (5f)$$

$$\mu_N = \mu_f. \quad (5g)$$

In the following sections, we will convert this problem to an equivalent convex one.

III. UNCONSTRAINED COVARIANCE STEERING

It is well established in the covariance steering literature that under no coupled mean-covariance constraints, problem (5) can be decoupled into the mean steering problem and the covariance steering problem [7], [8]. The solution to the mean steering is trivial, therefore, we focus solely on the covariance steering, which corresponds to the optimization problem

$$\min_{\Sigma_k, K_k} J_\Sigma = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k K_k \Sigma_k K_k^\top), \quad (6)$$

subject to (5b)-(5d),

Using the change of variables $U_k = K_k \Sigma_k$ and the convex relaxation proposed in [17] one can transform Problem (6) into a linear semidefinite program

$$\min_{\Sigma_k, U_k, Y_k} J_\Sigma = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k Y_k) \quad (7a)$$

such that, for all $k = 0, 1, \dots, N-1$,

$$C_k \triangleq U_k \Sigma_k^{-1} U_k^\top - Y_k \preceq 0, \quad (7b)$$

$$G_k \triangleq A_k \Sigma_k A_k^\top + B_k U_k A_k^\top + A_k U_k^\top B_k^\top + B_k Y_k B_k^\top + D_k D_k^\top - \Sigma_{k+1} = 0, \quad (7c)$$

$$\Sigma_N - \Sigma_f = 0, \quad (7d)$$

where the constraint (7b) can be expressed as an LMI using the Schur complement as

$$\begin{bmatrix} \Sigma_k & U_k^\top \\ U_k & Y_k \end{bmatrix} \succeq 0.$$

Theorem 1. *The optimal solution to the relaxed problem (7) satisfies $C_k = 0$ for all $k = 0, 1, \dots, N-1$ and therefore also optimally solves (6) [17].*

Remark. A different approach that results in the same formulation is that of a randomized feedback control policy presented in [19]. Therein, the injected randomness on the control policy can be interpreted as a slack variable converting (7b) to equality. In [19] it is shown that for the soft-constrained version of the problem, the value of this slack variable is zero. In our work, we tackle directly the hard-constrained version, instead, with equality or inequality terminal covariance constraints as well as chance constraints. In this case, strong duality is not apparent and the technique of the proof of [19] is not directly applicable.

Next, consider Problem (6) but with an inequality terminal covariance boundary condition instead, and its corresponding relaxed version, namely,

$$\min_{\Sigma_k, U_k, Y_k} J_\Sigma = \sum_{k=0}^{N-1} \text{tr}(Q_k \Sigma_k) + \text{tr}(R_k Y_k), \quad (8a)$$

such that, for all $k = 0, 1, \dots, N-1$,

$$C_k \triangleq U_k \Sigma_k^{-1} U_k^\top - Y_k \preceq 0, \quad (8b)$$

$$\Sigma_N - \Sigma_f \preceq 0, \quad (8c)$$

$$G_k \triangleq A_k \Sigma_k A_k^\top + B_k U_k A_k^\top + A_k U_k^\top B_k^\top + B_k Y_k B_k^\top + D_k D_k^\top - \Sigma_{k+1} = 0. \quad (8d)$$

Theorem 2. Assuming that the exact covariance steering problem (6) is feasible, problem (8) satisfies $C_k = 0$ for all $k = 0, 1, \dots, N-1$ and therefore also optimally solves (6) with an inequality terminal covariance boundary condition, instead of (5d).

Proof. Using matrix Lagrange multipliers $M_k^{(1)}$, $M_k^{(2)}$, Λ_k for the constraints (8b), (8c), (8d), respectively, we define the Lagrangian function

$$\mathcal{L}(\cdot) = J_\Sigma + \text{tr}(M^{(2)}(\Sigma_N - \Sigma_f)) + \sum_{k=0}^{N-1} \text{tr}(M_k^{(1)} C_k) + \text{tr}(\Lambda_k G_k).$$

The relevant first-order optimality conditions are [20]:

$$\frac{\partial \mathcal{L}}{\partial U_k} = 2M_k^{(1)} U_k \Sigma_k^{-1} + 2B_k^\top \Lambda_k A_k = 0, \quad (9a)$$

$$\frac{\partial \mathcal{L}}{\partial Y_k} = R_k - M_k^{(1)} + B_k^\top \Lambda_k B_k = 0, \quad (9b)$$

$$\text{tr}(M_k^{(1)} C_k) = 0, \quad (9c)$$

where $k = 0, 1, \dots, N-1$. Note that we can choose Λ_k to be symmetric because of the symmetry of the constraint (8d), while $M_k^{(1)}$ and $M_k^{(2)}$ are symmetric by definition. We will prove that the optimal solution to problem (8) satisfies $C_k = 0$ for all $k = 0, 1, \dots, N-1$. To this end, assume that C_k has at least one nonzero eigenvalue for some k . Equation (9c) then yields that $M_k^{(1)}$ has to be singular [17]. The optimality condition (9a) can then be rewritten as $B_k^\top \Lambda_k = -M_k^{(1)} U_k \Sigma_k^{-1} A_k^{-1}$. Substituting to (9b) yields

$$R_k = M_k^{(1)} (I_p + U_k \Sigma_k^{-1} A_k^{-1} B_k). \quad (10)$$

Calculating the determinants of both sides of (10), we obtain

$$\det(R_k) = \det(M_k^{(1)}) \det(I_p + U_k \Sigma_k^{-1} A_k^{-1} B_k) = 0.$$

This clearly contradicts the fact that $R_k \succ 0$. Therefore, at the optimal solution, the matrix C_k has all its eigenvalues equal to zero. This, along with the fact that C_k is symmetric, yields that $C_k = 0$ for all $k = 0, 1, \dots, N-1$. The final step to conclude the proof is to show that the KKT conditions (9) for the relaxed problem (8) are sufficient for the optimal solution, or in other words, the duality gap for the relaxed problem is zero. We have already proved that strong duality holds for the exact covariance steering problem in [17]. Since




Fig. 1. Example of a convexified domain for a 1-dimensional system

the relaxed terminal boundary condition problem (8) has a domain at least as big as the exact problem (7) and strong duality holds for the exact problem, from Slater's condition strong duality holds for the relaxed problem as well. $\square$

IV. CONSTRAINED COVARIANCE STEERING

Many real-world applications require additional constraints of the form (2f), (2e) to be imposed on the problem to reflect the physical limitations of the system or some other desired behavior. These may include constraints on the total control effort u_k on each time step or physical limits on the state vector x_k . In this work, we assume polytopic state and control constraints of the form

$$\mathbb{P}(\alpha_x^\top x_k \leq \beta_x) \geq 1 - \epsilon_x, \quad (11a)$$

$$\mathbb{P}(\alpha_u^\top u_k \leq \beta_u) \geq 1 - \epsilon_u, \quad (11b)$$

where $\alpha_x \in \mathbb{R}^n$, $\alpha_u \in \mathbb{R}^p$, $\beta_x, \beta_u \in \mathbb{R}$ and $\epsilon_x, \epsilon_u \in [0, 0.5]$ reflects the violation probability of each constraint. To convert the probabilistic constraints (11) into deterministic constraints on the decision variables note that $\alpha_x^\top x_k$ and $\alpha_u^\top u_k$ are univariate random variables with first and second moments given by

$$\mathbb{E}(\alpha_x^\top x_k) = \alpha_x^\top \mu_k, \quad (12a)$$

$$\mathbb{E}(\alpha_u^\top u_k) = \alpha_u^\top v_k, \quad (12b)$$

$$\mathbb{E}(\alpha_x^\top (x_k - \mu_k)(x_k - \mu_k)^\top \alpha_x) = \alpha_x^\top \Sigma_k \alpha_x, \quad (12c)$$

$$\mathbb{E}(\alpha_u^\top K_k (x_k - \mu)(x_k - \mu)^\top K_k^\top \alpha_u) = \alpha_u^\top U_k \Sigma_k^{-1} U_k^\top \alpha_u. \quad (12d)$$

To this end, according to [8], equations (11) are converted to

$$\Phi^{-1}(1 - \epsilon_x) \sqrt{\alpha_x^\top \Sigma_k \alpha_x} + \alpha_x^\top \mu_k - \beta_x \leq 0, \quad (13a)$$

$$\Phi^{-1}(1 - \epsilon_u) \sqrt{\alpha_u^\top U_k \Sigma_k^{-1} U_k^\top \alpha_u} + \alpha_u^\top v_k - \beta_u \leq 0, \quad (13b)$$

where $\Phi^{-1}(\cdot)$ is the inverse cumulative distribution function of the normal distribution. If the Gaussian assumption for the disturbances is dropped, then $\Phi^{-1}(\cdot)$ can be conservatively replaced using Cantelli's concentration inequality with $Q(1 - \epsilon) = \sqrt{\epsilon/(1 - \epsilon)}$ [16].

Using the same relaxation as before to handle the non-linear term $U_k \Sigma_k^{-1} U_k^\top$, equation (13b) is further relaxed to

$$\Phi^{-1}(1 - \epsilon_u) \sqrt{\alpha_u^\top Y_k \alpha_u} + \alpha_u^\top v_k - \beta_u \leq 0. \quad (14)$$

Unfortunately, due to the square root on the decision variables Σ_k and Y_k neither of (13a), (14) are convex. One

conservative option to overcome this issue is to linearize these constraints around some reasonable value of $\alpha_x^\top \Sigma_k \alpha_x$ and $\alpha_u^\top Y_k \alpha_u$, respectively, for a given problem. Because the square root is a concave function, the tangent line can serve as a linear global overestimator, yielding

$$\sqrt{x} \leq \frac{1}{2\sqrt{x_0}}x + \frac{\sqrt{x_0}}{2}, \quad \forall x, x_0 > 0.$$

The constraints in (13) can therefore be conservatively approximated as

$$\begin{aligned} & \Phi^{-1}(1 - \epsilon_x) \frac{1}{2\sqrt{\alpha_x^\top \Sigma_r \alpha_x}} \alpha_x^\top \Sigma_k \alpha_x + \alpha_x^\top \mu_k \\ & - \left(\beta_x - \Phi^{-1}(1 - \epsilon_x) \frac{1}{2} \sqrt{\alpha_x^\top \Sigma_r \alpha_x} \right) \leq 0, \end{aligned} \quad (15a)$$

$$\begin{aligned} & \Phi^{-1}(1 - \epsilon_u) \frac{1}{2\sqrt{\alpha_u^\top Y_r \alpha_u}} \alpha_u^\top Y_k \alpha_u + \alpha_u^\top v_k \\ & - \left(\beta_u - \Phi^{-1}(1 - \epsilon_u) \frac{1}{2} \sqrt{\alpha_u^\top Y_r \alpha_u} \right) \leq 0, \end{aligned} \quad (15b)$$

where Σ_r , Y_r are some reference values. The linearized constraints now form a convex set, as illustrated in Figure 1. For notational simplicity, next, we consider the more general constraint form of

$$\ell^\top \Sigma_k \ell + \alpha_x^\top \mu_k - \beta_x \leq 0, \quad (16a)$$

$$e^\top Y_k e + \alpha_u^\top v_k - \beta_u \leq 0. \quad (16b)$$

Given the additional constraints in (16) the fundamental question is whether the relaxation proposed in (7) remains lossless. To this end, consider the constrained Covariance Steering problem

$$\min_{\Sigma_k, U_k, Y_k, \mu_k, v_k} J, \quad (17a)$$

such that, for all $k = 0, 1, \dots, N-1$,

$$\mu_{k+1} = A_k \mu_k + B_k v_k, \quad (17b)$$

$$C_k(\Sigma_k, Y_k, U_k) \leq 0, \quad (17c)$$

$$G_k(\Sigma_{k+1}, \Sigma_k, Y_k, U_k) = 0, \quad (17d)$$

$$\ell^\top \Sigma_k \ell + \alpha_x^\top \mu_k - \beta_x \leq 0, \quad (17e)$$

$$e^\top Y_k e + \alpha_u^\top v_k - \beta_u \leq 0. \quad (17f)$$

where J is defined in (5a). Note that an equality terminal covariance condition is implied, by excluding Σ_N from the optimization variables and treating it as constant.

Theorem 3. *The optimal solution to the problem (17) satisfies $C_k = 0$ for all $k = 0, 1, \dots, N-1$.*

Proof. Define again the problem Lagrangian as

$$\begin{aligned} \mathcal{L}_a(\cdot) = & J + \sum_{k=0}^{N-1} \text{tr}(M_k^\top C_k) + \text{tr}(\Lambda_{1,k}^\top G_k) \\ & + \lambda_{1,k}^\top (\mu_{k+1} - A_k \mu_k - B_k v_k) \\ & + \lambda_{2,k} (\ell^\top \Sigma_k \ell + \alpha_x^\top \mu_k - \beta_x) \\ & + \lambda_{3,k} (e^\top Y_k e + \alpha_u^\top v_k - \beta_u). \end{aligned}$$

The relevant first-order optimality conditions for this problem are

$$\frac{\partial \mathcal{L}_a}{\partial U_k} = 2M_k U_k \Sigma_k^{-1} + 2B_k^\top \Lambda_k A_k = 0, \quad (18a)$$

$$\frac{\partial \mathcal{L}_a}{\partial Y_k} = R_k - M_k + B_k^\top \Lambda_k B_k + \lambda_{3,k} e e^\top = 0, \quad (18b)$$

$$\text{tr}(M_k C_k) = 0. \quad (18c)$$

Following the same steps as in the proof of Theorem 1, let C_k have at least one nonzero eigenvalue. From (18c), M_k has to be singular. Solving for $B_k^\top \Lambda_k$ in (18a) and substituting in (18b) we get

$$R_k + \lambda_{3,k} e e^\top = M_k (I_p + U_k \Sigma_k^{-1} A_k^{-1} B_k). \quad (19)$$

Since $\lambda_{3,k} \geq 0$ by definition, and $e e^\top \succeq 0$, it follows that $R_k + \lambda_{2,k} e e^\top \succ 0$. Therefore, taking again the determinant of both sides of (19) leads to a contradiction. $\square$

V. NUMERICAL EXAMPLES AND RUN-TIME ANALYSIS

To illustrate our method, we will first consider the problem of path planning for a quadrotor in a 2D plane. The lateral and longitudinal dynamics of the quadrotor will be modeled as a triple integrator, with state matrices

$$A = \begin{bmatrix} I_2 & \Delta T I_2 & 0_2 \\ 0_2 & I_2 & \Delta T I_2 \\ 0_2 & 0_2 & I_2 \end{bmatrix}, \quad B = \begin{bmatrix} 0_2 \\ 0_2 \\ \Delta T I_2 \end{bmatrix}, \quad D = 0.1 I_6,$$

a time step of $\Delta T = 0.1$ sec, a horizon of $N = 60$ and boundary conditions

$$\begin{aligned} \Sigma_i &= I_6, \quad \Sigma_f = 0.1 I_6, \\ \mu_i &= \begin{bmatrix} 20 & 0_{1 \times 5} \end{bmatrix}^\top, \quad \mu_f = 0_{6 \times 1}. \end{aligned}$$

The feasible state space and control input space are characterized by bounding boxes expressed in the form of (15) with parameters

$$\alpha_x = \left\{ \begin{bmatrix} \pm 1 & 0 & 0_{1 \times 4} \end{bmatrix}^\top, \begin{bmatrix} 0 & \pm 1 & 0_{1 \times 4} \end{bmatrix}^\top \right\},$$

$$\beta_x = \{22, -3, 7, -7\}$$

$$\alpha_u = \left\{ \begin{bmatrix} \pm 1 & 0 \end{bmatrix}^\top, \begin{bmatrix} 0 & \pm 1 \end{bmatrix}^\top \right\}, \quad \beta_u = \{\pm 25, \pm 25\}.$$

Also, two position waypoints are implemented by constraining the first two components of the state at time steps 20 and 40 of the steering horizon. The chance constraint linearization is performed around $\Sigma_r = 1.2 I_6$ and $Y_r = 15 I_2$. All optimization problems are solved in Matlab using MOSEK [21] and YALMIP [22]. The resulting optimal steering as well as the control effort is illustrated in Figure 2. The feasible set in each figure is denoted with green lines and the mean of each signal with a dashed black line. The 3-sigma confidence level bounds are represented with blue ellipses for the state covariance and by the light-blue area around the mean control signal. Initial and terminal values for the state 3-sigma confidence ellipses as well as the waypoints are denoted with red.

In the previous example, the plant dynamics were assumed linear. In the next example, we will explore the performance




Fig. 2. Left: Resulting trajectory, Right: required control effort. In both Figures, green lines represent the feasible part of the state space and control action space.

of the CS algorithm for controlling the non-linear quadrotor dynamics around an aggressive reference trajectory. The nonlinear system describing a quadrotor is

$$\dot{r} = v, \quad (20a)$$

$$\dot{v} = 1/m(e_3 g + R(q)\hat{e}_3 \tau + w_f), \quad (20b)$$

$$\dot{q} = S(q)(\omega + w_m), \quad (20c)$$

where r , v represent the position and velocity of the quad in an inertial coordinate frame, $q = [\psi \ \phi \ \theta]^T$ represent the attitude, parameterized using ZYX Euler angles. The system inputs are the total thrust τ and the body-frame angular rates ω^T . In this setup, it is assumed that the thrust and rotational velocity commands can be realized by means of some low-level controller running at a higher frequency. The matrix $R(q)$ is the standard rotation matrix and $S(q)$ is the matrix converting the body-frame angular rates ω to Euler angle rates $\dot{q}$. The total quadrotor mass is denoted by m , the acceleration due to gravity by g , and the unit vector in the z direction by e_3 . The vectors w_f, w_m represent disturbances. This system is known to be differentially flat, with a flat output $\zeta = [r_x \ r_y \ r_z \ \psi]^T$ [23]. That is, given any smooth trajectory in the flat output, the commands required to realize this trajectory as well as the values of all the state variables throughout can be described in terms of the flat output and its derivatives. Exploiting this result, and considering $\psi = 0$ throughout the trajectory for simplicity, we generate a smooth, discrete-time nominal path for the flat output using the model of a triple integrator.

For controlling the uncertainty, we consider a linearization of the system around the nominal trajectory. To this end, note that the nonlinear plant (20) can be written in the general form $\dot{x} = f(x, u, w)$ and discretized using a first-order difference approximation, yielding

$$x_{k+1} = F(x_k, u_k, w_k) = x_k + \Delta T f(x_k, u_k, w_k), \quad (21)$$

where ΔT is the sampling step. From there, using a Taylor series expansion, the system can be approximated by

$$x_{k+1} \approx F(\bar{x}_k, \bar{u}_k, 0) + \frac{\partial F}{\partial x}(x_k - \bar{x}_k) + \frac{\partial F}{\partial u}(u_k - \bar{u}_k) + \frac{\partial F}{\partial w} w_k$$

Noticing that $\bar{x}_{k+1} = F(\bar{x}_k, \bar{u}_k, 0)$, the deviation from the nominal trajectory $\tilde{x}_k = x_k - \bar{x}_k$ can be propagated using




Fig. 3. Uncertainty control around nominal trajectory.




Fig. 4. Required control effort

the LTV system

$$\tilde{x}_{k+1} = A_k \tilde{x}_k + B_k \tilde{u}_k + (D_k + \tilde{D}) w_k, \quad (22)$$

where A_k , B_k , D_k are the Jacobian matrices of the Taylor expansion evaluated around the nominal trajectory and control input, and zero disturbance, $\tilde{u}_k$ is the deviation from the nominal input and w_k is the external disturbance, assumed here to be zero mean white noise with unitary covariance. Finally, $\tilde{D}$ accounts for all discretization and linearization errors. The control of the deviation from the nominal trajectory can therefore be cast as a covariance steering problem subject to the dynamics of (22). To this end, consider the standard CS problem with cost matrices $Q_k = 10I_6$, $R_k = 0.1I_4$, and initial and final terminal covariances

$$\begin{aligned} \Sigma_i &= \text{blkdiag}(10^{-2}I_3, 10^{-3}I_6), \\ \Sigma_f &= \text{blkdiag}(5 \cdot 10^{-4}I_3, 10^{-3}I_6) \end{aligned}$$

The first 3×3 blocks of the covariance matrices for all time steps after $k = 50$ are constrained to be smaller than $2 \cdot 10^{-3}I_3$, guaranteeing that the trajectory stays within a prescribed tube with respect to the nominal one, which is the minimum jerk path subject to the dynamics of a discrete, triple integrator with initial and terminal boundary conditions

$$\bar{x}_0 = [0, 0.5, 0, 0_{1 \times 6}]^T, \bar{x}_N = [0, -0.5, 0, 0_{1 \times 6}]^T \quad (23)$$

and four position waypoints (see Figure 3). The sampling time was chosen as $\Delta T = 0.01s$ and the steering horizon was $N = 500$ steps. The resulting steering is illustrated in Figure 3, while the required control effort is in Figure 4.

Finally, we present a run-time and resulting optimization problem size comparison between different methods for solving the unconstrained covariance steering problem. To evaluate the performance of each algorithm, random state space models of various sizes were generated using Matlab's `drss()` command. For each system, we use as many noise channels as state variables and half as many input channels as state variables. The analysis was performed for systems of varying size and a fixed steering horizon of 32 time steps, as well as for varying time horizons for an 8×8 system. The results are summarized in Tables I and II respectively. Run times are measured in seconds and the problem size is the number of decision variables in each program. The empty cells are due to the program running out of memory. The simulations were carried out in Matlab 2022 running on an 11th Gen. Intel Core i7-11800H and 16 GB of RAM.

TABLE I
RUN-TIME COMPARISON FOR VARYING STATE SPACE SIZE.

n	Approach 1, [7]		Approach 2, [10]		Proposed approach	
	p. size	r. time	p. size	r. time	p. size	r. time
4	3200	93.28	256	0.20	884	0.03
8	10496	-	1024	2.91	3536	0.18
16	37376	-	4096	138.07	14144	2.59
32	140288	-	16384	-	56576	151.76

TABLE II
RUN-TIME COMPARISON FOR VARYING HORIZON SIZE.

N	Approach 1, [7]		Approach 2, [10]		Proposed approach	
	p. size	r. time	p. size	r. time	p. size	r. time
8	640	3.57	256	0.12	848	0.04
16	2306	76.74	512	0.70	1744	0.08
32	8704	-	1024	3.33	3536	0.17
64	33792	-	2048	19.27	7120	0.36
128	133120	-	4096	-	14288	0.75
256	528384	-	8129	-	28624	1.60

It is clear that the proposed approach outperforms the state-of-the-art algorithms significantly, by over an order of magnitude for almost all cases. Also, it is worth noting that problem (7) is a linear semidefinite program, while the formulations of [7] and [9] result in quadratic semidefinite programs, which need to be converted to linear ones using suitable relaxations, increasing further the number of decision variables needed as well as the complexity of the problem. Finally, Problem (7) involves $N - 1$ LMIs of dimensions $p \times p$ as opposed to a single large LMI of dimensions $(N + 2)n \times (N + 2)n$ for the terminal covariance constraint used in methods [7], [9]. As suggested in [21], multiple smaller LMIs can be solved more efficiently compared to a single larger one due to the resulting sparsity of the constraints. This also explains why although Approach 2 of [9] results in smaller problem sizes compared to the proposed approach, still has significantly larger solution times.

ACKNOWLEDGMENT

The authors would like to sincerely thank Dr. Fengjiao Liu for her constructive comments and discussion on the paper, and Ujjwal Gupta for his help with the quadrotor example. Support for this work has been provided

by ONR award N00014-18-1-2828 and NASA ULI award #80NSSC20M0163. This article solely reflects the opinions and conclusions of its authors and not of any NASA entity.

REFERENCES

- [1] A. Hotz and R. E. Skelton, "Covariance control theory," *International Journal of Control*, vol. 46, pp. 13–32, July 1987.
- [2] K. M. Grigoriadis and R. E. Skelton, "Minimum-energy covariance controllers," *Automatica*, vol. 33, no. 4, pp. 569–578, 1997.
- [3] J. A. Primbs and C. H. Sung, "Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise," *IEEE Transactions on Automatic Control*, vol. 54, pp. 221–230, Feb. 2009.
- [4] M. Farina, L. Giulioni, L. Magni, and R. Scattolini, "A probabilistic approach to model predictive control," in *52nd IEEE Conference on Decision and Control*, (Firenze, Italy), pp. 7734–7739, Dec. 2013.
- [5] 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, pp. 1158–1169, May 2015.
- [6] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part II," *IEEE Transactions on Automatic Control*, vol. 61, pp. 1170–1180, May 2015.
- [7] E. Bakolas, "Finite-horizon covariance control for discrete-time stochastic linear systems subject to input constraints," *Automatica*, vol. 91, pp. 61–68, May 2018.
- [8] K. Okamoto, M. Goldshtein, and P. Tsiotras, "Optimal covariance control for stochastic systems under chance constraints," *IEEE Control Systems Letters*, vol. 2, pp. 266–271, July 2018.
- [9] K. Okamoto and P. Tsiotras, "Stochastic model predictive control for constrained linear systems using optimal covariance steering," *arXiv preprint arXiv:1905.13296*, 2019.
- [10] K. Okamoto and P. Tsiotras, "Optimal stochastic vehicle path planning using covariance steering," *IEEE Robotics and Automation Letters*, vol. 4, pp. 2276–2281, July 2019.
- [11] F. Liu and P. Tsiotras, "Optimal covariance steering for continuous-time linear stochastic systems with multiplicative noise," *arXiv preprint arXiv:2206.11735*, 2022.
- [12] A. D. Saravanos, A. Tsolovikos, E. Bakolas, and E. Theodorou, "Distributed covariance steering with consensus ADMM for stochastic multi-agent systems," in *Proceedings of Robotics: Science and Systems*, (Virtual), July 2021.
- [13] J. Ridderhof, K. Okamoto, and P. Tsiotras, "Nonlinear uncertainty control with iterative covariance steering," in *IEEE 58th Conference on Decision and Control (CDC)*, (Nice, France), pp. 3484–3490, Dec. 2019.
- [14] A. D. Saravanos, I. M. Balci, E. Bakolas, and E. A. Theodorou, "Distributed model predictive covariance steering," *arXiv preprint arXiv:2212.00398*, 2022.
- [15] V. Sivaramakrishnan, J. Pilipovsky, M. Oishi, and P. Tsiotras, "Distribution steering for discrete-time linear systems with general disturbances using characteristic functions," in *American Control Conference (ACC)*, (Atlanta, GA, USA), pp. 4183–4190, June 2022.
- [16] V. Renganathan, J. Pilipovsky, and P. Tsiotras, "Distributionally robust covariance steering with optimal risk allocation," *arXiv preprint arXiv:2210.00050*, 2022.
- [17] F. Liu, G. Rapakoulas, and P. Tsiotras, "Optimal covariance steering for discrete-time linear stochastic systems," *arXiv preprint arXiv:2211.00618*, 2022.
- [18] I. M. Balci and E. Bakolas, "Covariance steering of discrete-time linear systems with mixed multiplicative and additive noise," *arXiv preprint arXiv:2210.01743*, 2022.
- [19] I. M. Balci and E. Bakolas, "Exact SDP formulation for discrete-time covariance steering with Wasserstein terminal cost," *arXiv preprint arXiv:2205.10740*, 2022.
- [20] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM Review*, vol. 38, no. 1, pp. 49–95, 1996.
- [21] M. ApS, *Mosek Optimization Toolbox for MATLAB*, 2019.
- [22] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in MATLAB," in *In Proceedings of the CACSD Conference*, (Taipei, Taiwan), 2004.
- [23] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *IEEE International Conference on Robotics and Automation*, (Shanghai, China), pp. 2520–2525, 2011.