

Technical Notes

F-DOC: Feedback Desensitized Optimal Control

Joshua Pilipovsky*[✉] and Panagiotis Tsiotras[†][✉]
Georgia Institute of Technology, Atlanta, Georgia 30332
and

Joseph Hart[‡] and Bart van Bloemen Waanders[§]
Sandia National Laboratories, Albuquerque,
New Mexico 87123

<https://doi.org/10.2514/1.G008468>

I. Introduction

THE standard approach [1–3] to solving trajectory optimization problems under parametric (e.g., model) uncertainties involves two sequential steps. In the first step, one designs an optimal feed-forward controller under nominal conditions, in which the parameters are assumed to be known and fixed to their nominal values. Then, in the second step, the system is linearized about this nominal solution, and an optimal feedback controller is obtained (say, by solving a linear quadratic [LQ] problem), in which the state and control variations are appropriately penalized. This two-step procedure, which we refer to in this work as the *decoupled* design, has a clear purpose: the open-loop control generates the nominal trajectory, and the feedback control is used to mitigate the effect of uncertainties on the state. In this decoupled design, the open-loop control plays no part in reducing the sensitivity with respect to the uncertainties. As such, several works have focused on reducing the dispersion of the trajectories by the suitable selection of the *nominal* trajectory, i.e., something akin to robust *open-loop* control design. Early works on this *sensitivity* reduction design include Winsor and Roy [4], who were among the first to design controllers that achieve similar performance under model uncertainties for linear systems. Other techniques for sensitivity reduction include augmented system analysis [5], augmented cost function [6–8], and even feedback [9,10]. This theory was later extended to nonlinear systems with constraints by the work of Seywald et al. [11–13], who also coined the term “desensitized optimal control” (DOC) for this open-loop robust control approach.

Regardless of the nominal controller used, and since the linearized system is a function of the nominal solution, the optimal gains obtained by solving the linearized problem are dependent on the nominal controller. Therefore, if the nominal solution is sensitive to parametric variations, the feedback controller might require very

high control effort (leading to input saturation in some cases) to minimize the trajectory dispersion under parametric uncertainties, thereby incurring high feedback costs. One method to avoid this situation is to design a nominal controller that is insensitive or robust to parameter variations. However, it might also be beneficial to integrate the feedback control design process with the selection of the optimal feed-forward control to handle model uncertainties more efficiently. One expects that a coupled design that *simultaneously* optimizes the nominal and the feedback control law would be more robust. For a more thorough overview of this coupled feedback–feedforward design framework, we refer interested readers to [14].

In this work, we follow this key idea and propose to choose a nominal controller that also minimizes a feedback cost along with the nominal performance metric. This feedback, or sensitivity cost, is a function of the feedback gains of the chosen feedback subproblem. As a result, the feedback gains are optimized simultaneously with the open-loop, nominal controller. This, consequently, closes the control system loop within the nominal optimization problem itself.

The specific feedback design methodology chosen to desensitize the trajectories of the linearized system is up to the control system designer and the problem at hand; no single feedback method is suitable for all cases. For example, when the disturbance is Gaussian, the LQG (or $\mathcal{H}_2$ synthesis) problem guarantees a feedback control that minimizes the $\mathcal{H}_2$ norm of the closed-loop system [15]. However, when the disturbance is *non-Gaussian*, or some bounded disturbance, one may opt for the more robust $\mathcal{H}_\infty$ feedback control design that provides guarantees on the worst-case disturbance [16]. In this work, we consider $\mathcal{H}_2$, $\mathcal{H}_\infty$, and the more general mixed $\mathcal{H}_2/\mathcal{H}_\infty$ feedback design problems. To this end, there are many ways to solve these feedback problems, ranging from standard constrained optimization techniques [17] to linear matrix inequalities and semidefinite programming [18] and algebraic Riccati equation solutions through the bounded real lemma [19,20]. It has been shown that it is possible to solve each of these three feedback problems through the lens of differential game theory, and each of these problems is a specific case of a more general, non-zero-sum Nash game [21–23]. We follow this approach and provide the $\mathcal{H}_2$, $\mathcal{H}_\infty$, and mixed $\mathcal{H}_2/\mathcal{H}_\infty$ solutions as special cases of a single differential game control design problem involving two competing players (the control designer and nature).

The main idea of the proposed optimal control architecture is twofold: first, to reduce the effect of parametric disturbances on the state by minimizing the *sensitivity* of the state to the system parameters; second, to reduce the magnitude of feedback control effort by transferring part of the desensitization task to the open-loop controller. Both of these outcomes result in nominal controllers that are robust to model uncertainties. The previous work in [24] accomplishes the first result by augmenting the cost with sensitivity functions that relate the variation of the state to the uncertain parameters. This is done in a purely open-loop setting, however, and may result in large feedback gains and input saturation. The current work builds off [24] not only by including the sensitivities but also by considering the feedback gains in the original problem formulation. This gives a doubly robust solution that mitigates uncertainties *directly* via sensitivity minimization and *indirectly* via feedback minimization. The balance between open-loop and closed-loop sensitivity reduction is achieved by the solution of an optimization problem.

II. Uncertain Trajectory Optimization Problem

A. Problem Formulation

We consider the problem of finding the controller $u: [t_0, t_f] \rightarrow \mathcal{U} \subset \mathbb{R}^m$ and (perhaps) the final time t_f to minimize the cost

Received 30 April 2024; accepted for publication 24 February 2025; published online 31 March 2025. Copyright © 2025 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. All requests for copying and permission to reprint should be submitted to CCC at www.copyright.com; employ the eISSN 1533-3884 to initiate your request. See also AIAA Rights and Permissions <https://aiaa.org/publications/publish-with-aiaa/rights-and-permissions/>.

*Ph.D. Graduate, Daniel Guggenheim School of Aerospace Engineering; jpilipovsky3@gatech.edu.

[†]David & Lewis Chair and Professor, Daniel Guggenheim School of Aerospace Engineering and the Institute for Robotics & Intelligent Machines; tsiotras@gatech.edu.

[‡]Senior Member of Technical Staff, Scientific Machine Learning; joshart@sandia.gov.

[§]Distinguished Member of Technical Staff, Scientific Machine Learning; bartv@sandia.gov.

$$\bar{\mathcal{J}}(u, t_f) = \Gamma(x_{t_f}, t_f) + \int_{t_0}^{t_f} \mathcal{L}(x_t, u_t, t) dt \quad (1)$$

subject to the dynamics

$$\dot{x}_t = f(x_t, u_t, p, t), \quad x_0 = x_{t_0} : \text{given} \quad (2)$$

along with the trajectory constraints

$$g(x_t, u_t, p, t) \leq 0, \quad t_0 \leq t \leq t_f \quad (3)$$

and the terminal constraints

$$\Psi(x_{t_f}, t_f) = 0 \quad (4)$$

where $x_t \in \mathbb{R}^n$ is the state at time $t \in [t_0, t_f]$, $u_t \in \mathcal{U}$ with $\mathcal{U}$ compact and convex is the control input at time $t \in [t_0, t_f]$, and $p \in \mathcal{P} \subset \mathbb{R}^q$ with $\mathcal{P}$ compact is the vector of the unknown parameters. In the previous problem formulation, $f: \mathbb{R}^n \times \mathbb{R}^q \times \mathbb{R}^m \times [t_0, t_f] \mapsto \mathbb{R}^n$ are the dynamics, $\Gamma: \mathbb{R}^n \times [t_0, t_f] \mapsto \mathbb{R}$ is the terminal cost, and $\mathcal{L}: \mathbb{R}^n \times \mathbb{R}^m \times [t_0, t_f] \mapsto \mathbb{R}$ is the running cost. The vector-valued functions $g: \mathbb{R}^n \times \mathbb{R}^q \times \mathbb{R}^m \times [t_0, t_f] \mapsto \mathbb{R}$ and $\Psi: \mathbb{R}^n \times [t_0, t_f] \mapsto \mathbb{R}$ define the trajectory constraints, and the terminal condition at the final time, respectively. For convenience, and without loss of generality, we assume that all these functions satisfy the required regularity assumptions to ensure the existence of an optimal control [25].

The optimal controller that minimizes the objective (1), subject to the dynamics (2), the path constraints (3), and the terminal constraint (4), can be obtained in a feed-forward or a feedback form. In the former case, the controller is just a function of time, that is, u_t , while in the latter case, the controller is a function of both state x_t and time, that is, $u_t(x_t)$. A feedback controller is expected to handle disturbances and model uncertainties better than an open-loop controller and hence is preferable; however, obtaining optimal feedback control in the case of a general nonlinear system is elusive except for some very special cases. Instead, traditionally, the system is linearized along some nominal trajectory resulting from an optimal feed-forward controller, and then a linear-feedback control is obtained using, for example, LQ regulator (LQR) theory to minimize the deviation from the nominal trajectory. We call this traditional approach the *decoupled* method, as first the open-loop control is obtained, and only afterward is the feedback control designed from the linearized system. In this work, we advocate a coupled approach that computes the nominal open-loop and feedback controllers *simultaneously* within the same optimization problem.

B. Sensitivity Dynamics

For the aforementioned optimal control problem (1)–(4), observe that the resulting optimal trajectory depends on both the optimal control input and the *uncertain* parameters of interest. We may explicitly denote this dependence by $x_t(p; u_t)$. Assuming a nominal parameter vector $\bar{p}$, we define the corresponding nominal trajectory as $\bar{x}_t \triangleq x_t(\bar{p}; \bar{u}_t)$, where $\bar{u}_t$ denotes the nominal control input, which is the open-loop control we ultimately wish to optimize. Of course, if one were to naïvely use the nominal trajectory $\bar{x}_t$ under parametric uncertainty, the actual state might deviate significantly from $\bar{x}_t$. Hence, it is useful to quantify how deviations in the parameters from their nominal values will affect the resulting state trajectory. To this end, let $\tilde{x}_t \triangleq x_t(p) - x_t(\bar{p})$ and $\tilde{p} \triangleq p - \bar{p}$. To first order, it can be shown [26] that

$$\tilde{x}_t \approx S_t \tilde{p} \quad (5)$$

where $S_t \in \mathbb{R}^{n \times q}$ is the *sensitivity matrix*. In the purely open-loop setting, the dynamics of the sensitivity matrix are given by

$$\dot{S}_t = A_t S_t + D_t, \quad S_{t_0} = 0 \quad (6)$$

where the matrices A_t and D_t are the derivatives of the dynamics with respect to the state and the parameters, respectively, evaluated along the nominal trajectory $x_t(\bar{p})$, that is,

$$A_t = \frac{\partial f}{\partial x}(\bar{x}_t, \bar{u}_t, \bar{p}, t), \quad D_t = \frac{\partial f}{\partial p}(\bar{x}_t, \bar{u}_t, \bar{p}, t) \quad (7)$$

Crucially, since A_t and D_t are evaluated along the nominal trajectory $x_t(\bar{p})$, which itself is a function of the nominal control input $\bar{u}_t$, it follows that the sensitivity function is affected by the selection of the control inputs. Therefore, the topic of DOC aims to optimize the nominal control input that results in optimal trajectories that are less sensitive to parametric variations [24].

In this work, our objective is to desensitize the nominal OCP through a *coupled* feed-forward/feedback approach. We use the *controlled* sensitivity dynamics to guide the design of a nominal control input that is robust against changes in p . Specifically, we decompose the total control input as $u_t(\tilde{x}_t) = \bar{u}_t + \tilde{u}_t(\tilde{x}_t)$, where $\tilde{u}_t = K_t \tilde{x}_t$ is taken to be a linear-feedback controller. The resulting state deviation dynamics are given by

$$\dot{\tilde{x}}_t = (A_t + B_t K_t) \tilde{x}_t + D_t \tilde{p}, \quad \tilde{x}_{t_0} = 0 \quad (8)$$

and similarly, invoking Eq. (5), the sensitivity matrix S_t obeys

$$\dot{S}_t = (A_t + B_t K_t) S_t + D_t, \quad S_{t_0} = 0 \quad (9)$$

where

$$B_t = \frac{\partial f}{\partial u}(\bar{x}_t, \bar{u}_t, \bar{p}, t) \quad (10)$$

In order to derive DOC strategies, one must assume some statistics on the uncertain parameter vector p . We will focus on two cases. In the first case, we will assume that $\tilde{p}$ in Eq. (8) is distributed according to a zero-mean Gaussian distribution with covariance Σ_p , that is, $\tilde{p} \sim \mathcal{N}(0, \Sigma_p)$.[†] In the second case, we assume that $\tilde{p}$ is a constant within a given compact set $\mathcal{P}$, equivalently, is uniformly distributed within $\mathcal{P}$, that is, $\tilde{p} \sim U(\mathcal{P})$. We remark that the choices for the above parameter models were chosen to appropriately fit the feedback problems considered in this work. The LQR formulation assumes normally distributed disturbances, naturally aligning with the Gaussian parameter model, while the $\mathcal{H}_\infty$ feedback problem addresses worst-case disturbances, which, from a stochastic perspective, can be approximated by a uniform distribution, representing an agnostic view of uncertainty.

In the following section, we review various feedback subproblems, how they can be solved as special cases of a two-player differential game, and finally how to embed their solutions into the nominal optimal control problem to reduce parameter sensitivities.

III. Feedback Desensitization

The general purpose of feedback control is to design a sequence (or function) of feedback gains that minimizes some performance metric with respect to some dynamic system. In our context, we are dealing with a linear system characterizing the deviation of the state from its nominal trajectory, which is perturbed by the unknown parameters of the system. No matter the feedback design used, a key observation here is that the *linearized* system is affected not only by the choice of the feedback matrix but also (implicitly) by the choice of the nominal trajectory via the matrices A_t , B_t , and D_t ; see Eqs. (7) and (10). In order to increase the levels of achievable desensitization and also alleviate issues associated with a decoupled nominal/feedback design, we wish to shape the nominal trajectory in tandem with choosing a feedback controller. To this end, we modify the nominal objective function (1) to include a penalty related to the feedback costs, that is,

[†]Strictly speaking, this violates the original assumption that p belongs to a compact set $\mathcal{P}$. However, in practice, we can always ensure that $p \in \mathcal{P}$ with very high probability by suitably choosing Σ_p or by taking a truncated Gaussian.

$$\mathcal{J}(\bar{u}, t_f) = \bar{\mathcal{J}}(\bar{u}, t_f) + \varepsilon \tilde{\mathcal{J}}(K(\bar{u}), t_f) \quad (11)$$

where $\bar{\mathcal{J}}$ is as in Eq. (1), and $\tilde{\mathcal{J}}$ is the feedback cost associated with the corresponding feedback design problem (i.e., $\mathcal{H}_2, \mathcal{H}_\infty$), and $\varepsilon > 0$ is a tunable hyperparameter. Note that the feedback gain matrix K depends on the nominal input. Minimizing the cost in Eq. (11) subject to the nominal dynamics (2) and constraints (3) is known as feedback-desensitized optimal control (F-DOC). In Sec. III.A we demonstrate that all the feedback subproblems considered in this work can be seen as special cases of a broader two-player differential game between the controller and nature. This unified perspective provides the natural framework we use in Sec. III.A.2 to subsequently design the feedback cost function as defined in Eq. (11).

A. Non-Zero-Sum Nash Games

To embed feedback desensitization within the nominal OCP, we adopt a game-theoretic approach where the interaction between the control input and worst-case disturbances is modeled as a non-zero-sum differential game. In the sequel, we show that the resulting Nash equilibrium conditions naturally yield optimal feedback laws for various performance and robustness objectives. For simplicity of notation, in this section we drop the tilde from all variables in the following discussion; however, each variable is still understood to be the deviation from its nominal value.

We are interested in finding the Nash equilibrium strategies w_t^* and u_t^* that satisfy

$$\mathcal{J}_1(u_t^*, w_t^*) \leq \mathcal{J}_1(u_t^*, w_t) \quad \forall w_t, \forall t \in [t_0, t_f] \quad (12)$$

$$\mathcal{J}_2(u_t^*, w_t^*) \leq \mathcal{J}_2(u_t, w_t^*) \quad \forall u_t, \forall t \in [t_0, t_f] \quad (13)$$

where

$$\mathcal{J}_1(u, w) = \frac{1}{2} \int_0^{t_f} [\delta^2 w_t^T w_t - z_{2,t}^T z_{2,t}] dt \quad (14)$$

$$\mathcal{J}_2(u, w) = \frac{1}{2} \int_0^{t_f} [z_{1,t}^T z_{1,t} - \tau^2 w_t^T w_t] dt \quad (15)$$

with the dynamics and performance outputs

$$\dot{x}_t = A_t x_t + B_t u_t + D_t w_t, \quad x_0 \text{ given} \quad (16)$$

$$z_{1,t} = \begin{bmatrix} S_{1f}^{1/2} x_f \\ Q_t^{1/2} x_t \\ R_t^{1/2} u_t \end{bmatrix}, \quad z_{2,t} = \begin{bmatrix} S_{2f}^{1/2} x_f \\ Q_t^{1/2} x_t \\ R_t^{1/2} u_t \end{bmatrix} \quad (17)$$

where $S_{1f}, S_{2f} > 0$. In this problem formulation, one of the players (nature) chooses w to minimize cost $\mathcal{J}_1$ in Eq. (14), and the other player (the control designer) chooses u so as to minimize cost $\mathcal{J}_2$ in Eq. (15). The actions of the two players are coupled through the dynamics (16).

We assign the same state and control weights Q_t, R_t to both players for simplicity, but this can easily be generalized to different weight matrices. The hyperparameters $\delta, \tau > 0$ in Eqs. (14) and (15) define the type of problem that is being solved (i.e., $\mathcal{H}_2$ or $\mathcal{H}_\infty$ synthesis). The inequality (12) represents the fact that the *worst-case* disturbance, w^* , achieves the lowest total cost with respect to $\mathcal{J}_1$. Achieving the lowest cost in $\mathcal{J}_1$ implies that there is the *largest* energy gain from disturbance input to output, since the second term in the integral in Eq. (14) becomes maximized while the first term becomes minimized. Similarly, inequality (13) represents the fact that the optimal controller, u^* , achieves the lowest total cost with respect to $\mathcal{J}_2$. Intuitively, this implies (with $\tau = 0$) that the input regulates the state with minimum energy, subject to the *worst*-possible disturbance w^* .

The following theorem, based on the work in [21], provides the equilibrium solution to this Nash game.

Theorem III.1: For the system described by Eqs. (16) and (17), there exist Nash equilibrium strategies u_t^*, w_t^* such that the inequality (12) given by the cost (14) is satisfied for all w_t and the inequality (13) given by the cost (15) is satisfied for all u_t if and only if the coupled Riccati matrix differential equations

$$\begin{aligned} -\dot{S}_{1,t} &= S_{1,t} A_t + A_t^T S_{1,t} - Q_t \\ &\quad - [S_{1,t} \quad S_{2,t}] \begin{bmatrix} \delta^{-2} D_t D_t^T & B_t R_t^{-1} B_t^T \\ B_t R_t^{-1} B_t^T & B_t R_t^{-1} B_t^T \end{bmatrix} \begin{bmatrix} S_{1,t} \\ S_{2,t} \end{bmatrix}, \quad (18) \\ S_{1,t_f} &= S_{1f} \end{aligned}$$

$$\begin{aligned} -\dot{S}_{2,t} &= S_{2,t} A_t + A_t^T S_{2,t} + Q_t \\ &\quad - [S_{1,t} \quad S_{2,t}] \begin{bmatrix} \tau^2 \delta^{-4} D_t D_t^T & \delta^{-2} D_t D_t^T \\ \delta^{-2} D_t D_t^T & B_t R_t^{-1} B_t^T \end{bmatrix} \begin{bmatrix} S_{1,t} \\ S_{2,t} \end{bmatrix}, \quad (19) \\ S_{2,t_f} &= S_{2f} \end{aligned}$$

have solutions $S_{1,t} \leq 0$ and $S_{2,t} \geq 0$ on $[t_0, t_f]$. If the solutions exist, the Nash equilibrium strategies are uniquely specified by

$$w_t^* = \delta^{-2} D_t^T S_{1,t} x_t \quad (20)$$

$$u_t^* = -R_t^{-1} B_t^T S_{2,t} x_t \quad (21)$$

We consider the following three cases by varying the values of τ and δ in Eqs. (14) and (15).

1. Linear Quadratic Regulator

The standard LQR problem is obtained by setting $-S_{1,f} = S_{2,f} = Q_f$, $\tau = 0$, and $\delta \rightarrow \infty$, in which case $S_{1,t} = -P_t$ and $S_{2,t} = P_t$, where P_t satisfies the usual LQ Riccati equation:

$$-\dot{P}_t = P_t A_t + A_t^T P_t - P_t B_t R_t^{-1} B_t^T P_t + Q_t, \quad P_{t_f} = Q_f \quad (22)$$

The corresponding optimal values for the control and the disturbance are $u_t^* = -R_t^{-1} B_t^T P_t x_t$ and $w_t^* \equiv 0$, respectively. Note that in this case, the problem turns into just a *single*-player game, i.e., a standard optimal control problem.

2. $\mathcal{H}_\infty$ Synthesis

To formulate the $\mathcal{H}_\infty$ control problem, we set $-S_{1,f} = S_{2,f} = Q_f$ and $\tau = \delta$, which yields $P_{\infty,t} = -S_{1,t} = S_{2,t}$, where $P_{\infty,t}$ satisfies

$$\begin{aligned} -\dot{P}_{\infty,t} &= A_t^T P_{\infty,t} + P_{\infty,t} A_t - P_{\infty,t} (B_t R_t^{-1} B_t^T - \delta^{-2} D_t D_t^T) P_{\infty,t} + Q_t, \\ P_{\infty,t_f} &= Q_f \end{aligned} \quad (23)$$

The optimal strategies in this case become $u_t^* = -R^{-1} B^T P_{\infty,t} x_t$, and $w_t^* = \delta^{-2} D_t^T P_{\infty,t} x_t$. Comparing the costs (14 and 15), we see that $\mathcal{J}^* = -\mathcal{J}_1^* = \mathcal{J}_2^*$ in the case $\tau = \delta$, so the $\mathcal{H}_\infty$ solution corresponds to a zero-sum Nash game.

3. Mixed $\mathcal{H}_2/\mathcal{H}_\infty$ Synthesis

The mixed $\mathcal{H}_2/\mathcal{H}_\infty$ problem is retrieved by setting $\tau = 0$ for some $\delta < \infty$. This gives $P_{1,t} = S_{1,t}$ and $P_{2,t} = S_{2,t}$ with S_1, S_2 given by Eqs. (18) and (19). The corresponding equilibrium solutions are given in Eqs. (20) and (21). Note that, in this case, the two optimal costs $\mathcal{J}_1^*$ and $\mathcal{J}_2^*$ will *not* be equal to each other. As a result, this is a non-zero-sum game. The reader is referred to [21] for a rigorous proof of how this solution minimizes the $\mathcal{H}_2$ norm while satisfying the $\mathcal{H}_\infty$ constraint.

In the following section, we use the solution to the above differential game (and its limiting cases) to desensitize the nominal OCP. Specifically, we will take the expected value of the optimal controller

cost $\mathcal{J}_2^* = \mathcal{J}_2(u_i^*, w_i^*)$ in Eq. (15) with respect to the distribution of parameter deviations $\tilde{p} \sim \mathcal{N}(0, \Sigma_p)$ as the desensitized cost $\tilde{\mathcal{J}}$ in Eq. (11).

B. Construction of Desensitized Costs

In the following, we derive the desensitized costs for the various feedback subproblems presented in Sec. III.A.2 and show their dependence on the sensitivity components and feedback gains. The general idea is to quantify the expected feedback cost due to parametric uncertainties and integrate this cost into the nominal optimization problem. It should be reminded that the desensitized cost function and the associated feedback control gain are computed after the original nonlinear system is linearized about the state and control nominal trajectory.

1. Open-Loop DOC

As a comparison with the proposed method, we will use the desensitized cost $\tilde{\mathcal{J}}$ arising from purely open-loop control [24], which results from the F-DOC cost by setting all the feedback gains to zero. To this end, the only variables that are included in the feedback cost are the state deviations. Taking the standard LQR cost (equivalently, Eq. (15) with $u_i \equiv 0$) and removing the control deviations yields

$$\tilde{\mathcal{J}}_{\text{OL}}(\bar{u}_t, t_f) = \mathbb{E} \left[\frac{1}{2} \tilde{x}_f^T Q_f \tilde{x}_f + \frac{1}{2} \int_{t_0}^{t_f} \tilde{x}_t^T Q_t \tilde{x}_t dt \right] \quad (24)$$

Using the first-order approximation of the state variations in Eq. (5) and taking the expectation with respect to the parameter deviations gives

$$\begin{aligned} \tilde{\mathcal{J}}_{\text{OL}}(\bar{u}_t, t_f) &= \frac{1}{2} \mathbb{E}[\tilde{p}^T S_f^T Q_f S_f \tilde{p}] + \frac{1}{2} \int_{t_0}^{t_f} \mathbb{E}[\tilde{p}^T S_t^T Q_t S_t \tilde{p}] dt \\ &= \frac{1}{2} \text{tr}[\mathbb{E}[Q_f S_f \tilde{p} \tilde{p}^T S_f^T]] + \frac{1}{2} \int_{t_0}^{t_f} \text{tr}[\mathbb{E}[Q_t S_t \tilde{p} \tilde{p}^T S_t^T]] dt \\ &= \frac{1}{2} \text{tr}(Q_f S_f \Sigma_p S_f^T) + \frac{1}{2} \int_{t_0}^{t_f} \text{tr}(Q_t S_t \Sigma_p S_t^T) dt \end{aligned} \quad (25)$$

where in the second equality we cyclically arrange the terms under the trace operator, and in the last equality, we use the statistical properties of the uncertain parameter. DOC solutions are more robust than just optimizing the nominal trajectory, as we are incorporating the sensitivity of the states with respect to the parameters directly into the optimization problem through the additional term $\tilde{\mathcal{J}}_{\text{OL}}$ in Eq. (25).

2. $\mathcal{H}_2$ -Feedback DOC

Taking the expectation of the optimal LQR feedback cost in Eq. (15) yields

$$\tilde{\mathcal{J}}_{\mathcal{H}_2}(\bar{u}_t, t_f) = \mathbb{E} \left[\frac{1}{2} \tilde{x}_f^T Q_f \tilde{x}_f + \frac{1}{2} \int_{t_0}^{t_f} \tilde{x}_t^T (Q_t + K_t^T R_t K_t) S_t \Sigma_p S_t^T \tilde{x}_t dt \right] \quad (26)$$

Following a similar logic to open-loop DOC, we obtain

$$\tilde{\mathcal{J}}_{\mathcal{H}_2}(\bar{u}_t, t_f) = \text{tr}(Q_f S_f \Sigma_p S_f^T) + \int_{t_0}^{t_f} \text{tr}[(Q_t + K_t^T R_t K_t) S_t \Sigma_p S_t^T] dt \quad (27)$$

where $K_t = -R_t^{-1} B_t^T P_t$, and P_t satisfies the Riccati differential equation (22). We expect solutions from LQR-FDOC to be more robust than open-loop DOC, since we are not only incorporating the sensitivities of the state with respect to the parameters but also incorporating the solution to the LQR problem as a means of desensitizing the nominal trajectory.

3. $\mathcal{H}_\infty$ -Feedback DOC

Plugging the feedback gains for the two players in Eq. (15) (with $\tau = \delta$) and taking the expectation yields

$$\begin{aligned} \tilde{\mathcal{J}}_{\mathcal{H}_\infty}(\bar{u}_t, t_f) &= \text{tr}(Q_f S_f \Sigma_p S_f^T) \\ &+ \int_{t_0}^{t_f} \text{tr}[(Q_t + K_t^T R_t K_t - \delta^2 L_t^T L_t) S_t \Sigma_p S_t^T] dt \end{aligned} \quad (28)$$

where $K_t = -R_t^{-1} B_t^T P_{\infty,t}$, $L_t = \delta^{-2} D_t^T P_{\infty,t}$, and $P_{\infty,t}$ satisfies the Riccati differential equation (23). We expect solutions from $\mathcal{H}_\infty$ F-DOC to be more robust than LQR F-DOC since the associated feedback problem solved is more robust, at the cost of optimality with respect to the nominal objective function.

4. Mixed $\mathcal{H}_2/\mathcal{H}_\infty$ -Feedback DOC

Lastly, we incorporate both performance and robustness into our augmented OCP by weighing the $\mathcal{H}_\infty$ cost from Eq. (14) with the $\mathcal{H}_2$ cost from Eq. (15) via

$$\tilde{\mathcal{J}}_{\mathcal{H}_2/\mathcal{H}_\infty}(\bar{u}_t, t_f) = \tilde{\mathcal{J}}_{\mathcal{H}_\infty}(\bar{u}_t, t_f) + \beta \tilde{\mathcal{J}}_{\mathcal{H}_2}(\bar{u}_t, t_f) \quad (29)$$

where $\beta > 0$ is a weight parameter that weighs the importance of $\mathcal{H}_\infty$ with $\mathcal{H}_2$. Note that the case $\beta = 0$ corresponds to pure $\mathcal{H}_\infty$ feedback, while the case $\beta \rightarrow \infty$ corresponds to pure LQR feedback.

C. Summary of F-DOC Architecture

An overview of the F-DOC workflow is shown in Algorithm 1. First, the nonlinear system is linearized about the nominal parameter values and (at the moment, unknown) nominal trajectory and input. Then, an associated feedback problem (LQR, $\mathcal{H}_\infty$, or $\mathcal{H}_2/\mathcal{H}_\infty$) is solved for the resulting linearized system, which yields the associated Riccati matrices P_t and the corresponding feedback gains K_t . The sensitivity dynamics are then propagated using the feedback gains and linearized dynamics. Lastly, using Eq. (5), the desensitized cost is constructed by taking the expectation of the optimal feedback cost from Eq. (15) with respect to the distribution of parameter deviations $\tilde{p} \sim \mathcal{N}(0, \Sigma_p)$. This resulting objective is then optimized (subject to the nominal system dynamics and constraints, sensitivity dynamics, and Riccati dynamics) to obtain the robust nominal trajectory and control input, as well as the optimal feedback gains.

Algorithm 1: F-DOC Architecture

Input: Cost tradeoff $\varepsilon > 0$, nominal parameter $\bar{p}$, nonlinear system f , feedback problem (τ, δ) , initial guess $(\bar{u}_t, \bar{x}_t)$

Output: Optimal nominal control and state trajectories $(\bar{u}_t^*, \bar{x}_t^*)$, final time t_f^*

- 1: Compute dynamics linearizations A_t, B_t, D_t from Eqs. (7) and (10)
- 2: Compute optimal gains K_t from Eq. (21)
- 3: Encode sensitivity dynamics constraints from Eq. (9)
- 4: Compute desensitized cost $\tilde{\mathcal{J}}$ for chosen feedback problem
- 5: Minimize augmented cost (11) with respect to nominal dynamics (16) and nominal constraints (3) and (4)

IV. Numerical Examples

We consider two examples illustrating the benefits of the proposed control architecture to generate trajectories and feedback controls that are insensitive to system parameter variations. For each example below, the F-DOC problem is solved using GPOPS-II [27], which is an off-the-shelf, general-purpose NLP solver. All simulations were run on a 32 GB Intel i7-10750H at 2.60 GHz desktop computer.

A. Zermelo's Problem

We first start with a simple two-dimensional problem to compare each DOC method. The state follows the dynamics

$$\dot{x} = \cos u + p_1 y \quad (30a)$$

$$\dot{y} = \sin u + p_2 \quad (30b)$$

where the parameters obey the statistics $p \sim \mathcal{N}(\bar{p}, \Sigma_p)$, where $\bar{p} = [10, -0.5]^T$ and $\Sigma_p = \text{diag}(\sigma_1^2, \sigma_2^2)$, where the standard deviation of p_i ($i = 1, 2$) is computed as $\sigma_i = (\eta_i/3)\bar{p}_i$, with $\eta = [0.1, 0.05]^T$ the 3σ error from the mean. Equations (30) are often used to describe the problem of steering a ship in uncertain sea currents given by p_2 in the y -direction, and p_1 in the x -direction. The input is simply the steering angle u . The nominal objective is to *maximize* the range of the ship in the x -direction in the presence of the uncertain currents, that is,

$$\bar{J}(\bar{u}) = -\bar{x}_f \quad (31)$$

where the final time $t_f = 1$ is fixed. Additionally, the boundary conditions are $x_0 = y_0 = 0$ and $y_f = 0$. The parameters for each of the three subproblems for F-DOC are shown in Table 1. To put each method on equal footing, we compare the *closed-loop* Monte Carlo trajectories for 100 samples of the parameters, as shown in Fig. 1. The baseline solution is the one that optimizes the nominal cost (no feedback cost) subject to the nominal dynamics and constraints. For the baseline and open-loop DOC cases, we use a decoupled control architecture, which computes the solution of an LQR problem after the nominal trajectory and control have been found. As mentioned, this is not needed for the F-DOC methods since we obtain the optimal feedback gains as a byproduct of the nominal optimization process. Table 2 displays the terminal downrange distance $\bar{x}_f$ and the volume V_f of the terminal covariance ellipses. It is clear that among the F-DOC methods, the LQR feedback has the largest downrange distance since it achieves the best performance (in the sense of maximizing the cost function); however, it has the largest terminal covariance as well. The $\mathcal{H}_\infty$ feedback solution has the *smallest* downrange distance but has the smallest covariance, and hence it is the most robust of the three methods. Lastly, the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ feedback solution is in-between the two extremes, and it is up to the designer to choose how to balance performance and robustness. This point is further illustrated in Figs. 2b and 2c, which show the sensitivity functions and nominal trajectories, respectively. We see graphically that each method (from left to right) progressively lowers the final downstream value (Fig. 2c) while decreasing sensitivity to the parameters by achieving a lower terminal sensitivity S_f (Fig. 2b).

To further assess the effectiveness of each feedback strategy, we examine the terminal points and compare their dispersions. Shown in Fig. 3 is a comparison of the nominal (baseline) solution, open-loop DOC, as well as the three F-DOC methods. In all cases, we again see that coupled F-DOC feedback strategies achieve lower terminal covariances, with even the worst-case (LQR) giving a decrease in the covariance due to the uncertain parameters as compared with decoupled open-loop DOC. Figure 3 also visually confirms the results obtained in Table 2.

B. Longitudinal Hypersonic Entry Vehicle

We now address a problem of more practical interest, namely, that of hypersonic entry in an uncertain atmosphere. The longitudinal dynamics of a (point-mass) vehicle in unpowered, atmospheric flight around a spherical, nonrotating planet are described by the system of differential equations




Fig. 1 Monte Carlo samples of uncertain parameters.

Table 2 Nominal terminal downrange and terminal volumes of the 3σ ellipses

Parameter	Baseline	Open loop	LQR	$\mathcal{H}_2/\mathcal{H}_\infty$	$\mathcal{H}_\infty$
$\bar{x}_f$, m	2.1415	2.1321	2.1142	1.9452	1.8608
V_f , cm ²	185.17	116.36	83.11	33.34	20.16

$$\dot{r} = v \sin \gamma \quad (32a)$$

$$\dot{\phi} = \frac{v}{r} \cos \gamma \quad (32b)$$

$$\dot{v} = -\frac{D}{m} - \frac{\mu \sin \gamma}{r^2} \quad (32c)$$

$$\dot{\gamma} = \frac{L}{mv} - \frac{\mu \cos \gamma}{r^2 v} + \frac{v \cos \gamma}{r} \quad (32d)$$

where r is the radial distance from the planet to the center of the vehicle, ϕ is the longitude, v is the planet-relative velocity, and γ is the planet-relative flight path angle. The planet's gravitational parameter is denoted by $\mu = 3.986 \times 10^{14} \text{ m}^3/\text{s}^2$, the mass by m , and the lift and drag are given by

$$L = \frac{1}{2} \rho v^2 S C_L(\alpha), \quad D = \frac{1}{2} \rho v^2 S C_D(\alpha) \quad (33)$$

in terms of the atmospheric density ρ , the reference area S , and the lift and drag coefficients C_L and C_D , which are both functions of the angle of attack α . We assume that the atmospheric density follows the exponential model

$$\rho(r) = \rho_0 \exp\left(-\frac{r - r_0}{H_s}\right) \quad (34)$$

where $\rho_0 = 1.225 \text{ kg/m}^3$ is density at the surface of the planet, $r_0 = 6,371 \text{ km}$ is the radius of the planet, and $H_s \sim \mathcal{N}(\bar{H}_s, \sigma_{H_s}^2)$ is the nominal scale height with statistics $\bar{H}_s = 7,254.24$ and

Table 1 Parameter values for open loop, LQR feedback, mixed $\mathcal{H}_2/\mathcal{H}_\infty$, and $\mathcal{H}_\infty$ feedback DOC

System	LQG		$\mathcal{H}_\infty$		Mixed $\mathcal{H}_2/\mathcal{H}_\infty$	
$\bar{p}$ [10 - 0.5]	Q	I_2	δ	5	β	10
η [0.1 0.05]	Q_f	blkdiag(1, 10)	Q_∞	I_2	Q_2	I_2
ϵ 100	R	1	Q_{f_∞}	blkdiag(1, 10)	Q_{f_2}	blkdiag(1, 10)
T 1	—	—	R_∞	1	R_2	1




Fig. 2 Comparison of nominal trajectories (c) and nominal control (a), as well as sensitivity dynamics (b) among the various feedback desensitization strategies.




Fig. 3 Comparison of terminal splash points among all methods from 100 MC runs.

$\sigma_{H_s} = 0.033\bar{H}_s$, which is set so that the 3σ variation is equal to 10% of the nominal value.

There are three pathwise constraints on the vehicle, namely, the heating rate Λ , normal load n , and dynamic pressure q , given by

$$\Lambda = k_\lambda \sqrt{\rho} v^3 \leq \Lambda_{\max} \quad (35a)$$

$$n = \frac{1}{m} \sqrt{L^2 + D^2} \leq n_{\max} \quad (35b)$$

$$q = \frac{1}{2} \rho v^2 \leq q_{\max} \quad (35c)$$

where k_λ is a heating constant, and the thresholds are determined by the vehicle characteristics. Lastly, define the state as $x := [r, \phi, v, \gamma]^T$

and the input as $u = \alpha$. We assume that there is some low-level control that can accurately execute the computed optimal control command α_k^* at time t_k . The nominal objective is to maximize the terminal longitude, that is, $\bar{\mathcal{J}} = -\bar{\phi}_f$, with the final time t_f a free variable. We consider a vehicle model from [28] with the application of reentry. To this end, the initial and terminal conditions and the vehicle parameters are given in Tables 3 and 4.

From Table 3, the initial speed is around Mach 8 and the final speed should be in the range $[0.8, 1.8]$ Mach. For the augmented cost, we use the state weight matrices $Q_i \equiv 0$, $t \in [0, t_f]$, and $Q_f = \text{diag}(\bar{Q}_r/r_0^2, \bar{Q}_\phi, \bar{Q}_v/r_0 g_0, \bar{Q}_\gamma)$, where the terminal weight is expressed in dimensional units, and $\bar{Q}_r = \bar{Q}_v = 100$ and $\bar{Q}_\phi = \bar{Q}_\gamma = 1$. To improve numerical performance, the problem is first

Table 3 Boundary conditions

State	Initial value	Final value	Units
h	45,720	1,524	m
ϕ	0	—	deg
v	2,743.2	[274.32, 609.60]	m/s
γ	-1	[-6, 0]	deg
t	0	—	s

Table 4 Space Shuttle characteristics

Parameter	Value	Units
S	242.38	m ²
m	92,079.34	kg
$\Lambda_{\max}$	4×10^5	W/m ²
$q_{\max}$	5.07×10^4	N/m ²
$n_{\max}$	$5g_0$	m/s ²
k_λ	9.4369×10^{-5}	kg ^{1/2} /m ^{3/2}

nondimensionalized with lengths and time being scaled by r_0 (radius of the Earth) and $\sqrt{r_0/g_0}$ (where g_0 is the Earth gravitational acceleration at r_0), respectively. For the DOC methods, we use $\varepsilon = 100$; for $\mathcal{H}_\infty$, we use $\delta = 200$. Lastly, the control weight is $R = 1000$.

As in the example in Sec. IV.A, we randomly sample 100 values of the uncertain scale height (as shown in Fig. 4c) and use them for

each Monte Carlo run. The optimal nominal control is shown in Fig. 4b; the trajectories in Fig. 4a; the sensitivity dynamics in Fig. 5a; and the dynamic pressure, load factor, and heating rate histories in Fig. 5b. From Fig. 4a, all solutions feature some form of phugoid oscillations during the descent of the vehicle, which is standard in high L/D vehicles [29]. The desensitized solutions, however, oscillate more than that of the baseline, and there is a notable divergence




Fig. 4 Comparison of optimal trajectories (a) and control (b) between baseline (black, dash-dot), DOC (black, dot), F-DOC-LQG (blue), and F-DOC- $\mathcal{H}_\infty$ (red), as well as parameter samples used for Monte Carlo runs (c).




Fig. 5 Sensitivity function components and heating rate, dynamic pressure, and load factor time histories for baseline (black, dash-dot), DOC (black, dot), F-DOC-LQG (blue), and F-DOC- $\mathcal{H}_\infty$ (red) methods.

in the trajectories after 100 s elapsed time. Additionally, the flight path angle starts oscillating near the end of the horizon for the desensitized solutions; this occurs during the period of fastest descent, implying that the vehicle is trying to pitch up and down to maintain speed in the uncertain atmosphere. These oscillations are also evident in the control, as shown in Fig. 4b, as a means of desensitizing once the atmosphere gets thicker upon descent.

Similarly, Fig. 5b quantitatively shows how robust the various methods are to uncertainties; the red and blue curves, corresponding to the LQR and $\mathcal{H}_\infty$ feedback methods, respectively, have the smallest final sensitivities at the end of the horizon. Even the open-loop DOC control yields a sensitivity near zero, with the exception of the sensitivity in longitude for each case. The longitude is not affected as

much by the uncertainty in the scale height, and as a result, it does not need to be as regulated when compared to the other three state variables. Lastly, from Fig. 5b, all open-loop nominal trajectories satisfy the pathwise constraints with no saturations.

To better see the differences between DOC and F-DOC for this problem, we directly compare the decoupled (DOC+LQR feedback) and coupled (F-DOC) splash points and also compare the open-loop versus closed-loop performance among the F-DOC methods to see how adding a feedback control affects desensitization. Figure 6 shows both the samples and distributions of the terminal state error from the nominal solution, and Fig. 7 shows the splash points between each pair of the final state variables. The decoupled method uses the solution to the DOC problem and adds an LQR controller after the




Fig. 6 Decoupled DOC with LQR (red) and coupled FDOC-LQG (black) terminal trajectory behaviors.




Fig. 7 Decoupled DOC with LQR feedback (red) and coupled LQR-F-DOC (black) terminal splash points.

fact by solving the associated Riccati equation using the nominal DOC solution; the feedback method uses the solution to the F-DOC LQR problem and uses the same gains from the optimization.

Figure 6 shows that even though the terminal trajectories are similar, the coupled F-DOC method achieves less dispersion at the final time. The final longitude errors are roughly the same; however, those errors are small to begin with, whereas the height above ground level (AGL), speed, and flight path errors are more substantial, and those we see more of a reduction. Furthermore, in Fig. 6a, the terminal distribution for F-DOC is sharper-peaked around zero compared to that of DOC with LQR. Interestingly, the flight path angle is skewed for the DOC method, giving a nonzero error for all Monte Carlo (MC) runs, whereas F-DOC is centered around zero with a higher probability.

Lastly, Fig. 7 shows that the F-DOC method has a smaller covariance in the terminal splash points compared to that of DOC. There is a 25% reduction in the terminal altitude error, from a 3σ value of 800 to 600 m, and a 33% reduction in the terminal velocity error, from a 3σ value of 12 to 8 m/s. Clearly, both methods are successful in lowering the sensitivity to the uncertain scale height, but the F-DOC is significantly better, with the added benefit of optimizing the gains and nominal control simultaneously.

V. Conclusions

In this paper, we have developed a coupled nominal/feedback design framework to simultaneously optimize nominal controllers that are robust to parametric uncertainties in the system's model as well as the associated feedback controller. This is done by solving the associated feedback problem via the solution to an equivalent differential game and by augmenting the nominal cost with the associated feedback game cost. The F-DOC strategy uses the solution to the feedback problem as a means of shaping the nominal trajectory, which is more robust than just nominal trajectory optimization and open-loop trajectory desensitization alone. The proposed framework has been applied to the Zermelo problem and the hypersonic entry vehicle problem, and it was shown to result in the most robust trajectories among all methods. Extensions of this work may include analyzing other risk measures or cost regularizers that will aid in the desensitization of trajectories.

Acknowledgments

Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

References

- [1] Matni, N., Ames, A. D., and Doyle, J. C., "A Quantitative Framework for Layered Multirate Control: Toward a Theory of Control Architecture," *IEEE Control Systems Magazine*, Vol. 44, No. 3, 2024, pp. 52–94.
<https://doi.org/10.1109/MCS.2024.3382388>
- [2] Todorov, E., and Li, W., "A Generalized Iterative LQG Method for Locally-Optimal Feedback Control of Constrained Nonlinear Stochastic Systems," *American Control Conference*, Inst. of Electrical and Electronics Engineers, New York, June 2005, pp. 300–306.
<https://doi.org/10.1109/ACC.2005.1469949>
- [3] Lu, P., "Entry Guidance and Trajectory Control for Reusable Launch Vehicle," *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 1, 1997, pp. 143–149.
<https://doi.org/10.2514/2.4008>
- [4] Winsor, C., and Roy, R., "The Application of Specific Optimal Control to the Design of Desensitized Model Following Control Systems," *IEEE Transactions on Automatic Control*, Vol. 15, No. 3, 1970, pp. 326–333.
<https://doi.org/10.1109/TAC.1970.1099453>
- [5] Subbayan, R., Sarma, V. V. S., and Vaithilingam, M. C., "An Approach for Sensitivity-Reduced Design of Linear Regulators," *International Journal of Systems Science*, Vol. 9, No. 1, 1978, pp. 65–74.
<https://doi.org/10.1080/002071727808941679>
- [6] Tang, D. L., "The Design of a Desensitized Controller—A Sensitivity Function Approach," *American Control Conference*, Inst. of Electrical and Electronics Engineers, New York, May 1990, pp. 3103–3111.
<https://doi.org/10.23919/ACC.1990.4791294>
- [7] Cheli, F., Resta, F., and Sabbioni, E., "Observer Design for Active Suspension Control for Passenger Cars," *Proceedings of the ASME Conference on Engineering System Design and Analysis*, American Soc. of Mechanical Engineers, New York, 2006, pp. 1–9.
<https://doi.org/10.1115/ESDA2006-95564>
- [8] Makkapati, V. R., Dor, M., and Tsiotras, P., "Trajectory Desensitization in Optimal Control Problems," *IEEE Conference on Decision and Control*, Inst. of Electrical and Electronics Engineers, New York, Dec. 2018, pp. 2478–2483.
<https://doi.org/10.1109/CDC.2018.8619577>
- [9] Zames, G., "Feedback and Optimal Sensitivity: Model Reference Transformations, Multiplicative Seminorms, and Approximate Inverses," *IEEE Transactions on Automatic Control*, Vol. 26, No. 2, 1981, pp. 301–320.
<https://doi.org/10.1109/TAC.1981.1102603>
- [10] Chang, B.-C., and Pearson, J., "Optimal Disturbance Reduction in Linear Multivariable Systems," *IEEE Transactions on Automatic Control*, Vol. 29, No. 10, 1984, pp. 880–887.
<https://doi.org/10.1109/TAC.1984.1103409>
- [11] Seywald, H., and Kumar, R. R., "Desensitized Optimal Trajectories," *6th Annual AAS/AIAA Space Flight Mechanics Meeting*, AIAA Paper 1996-0107, 1996.
- [12] Seywald, H., "Desensitized Optimal Trajectories with Control Constraints," *13th AAS/AIAA Space Flight Mechanics Meeting*, AIAA Paper 2003-147, 2003.
- [13] Seywald, K., and Seywald, H., "Desensitized Optimal Control," *AIAA SciTech Forum*, AIAA Paper 2019-0651, Jan. 2019.
<https://doi.org/10.2514/6.2019-0651>
- [14] Malyuta, D., Yu, Y., Elango, P., and Açıkmeşe, B., "Advances in Trajectory Optimization for Space Vehicle Control," *Annual Reviews in Control*, Vol. 52, Nov. 2021, pp. 282–315.
<https://doi.org/10.1016/j.arcontrol.2021.04.013>
- [15] Hespanha, J. P., *Linear Systems Theory*, Princeton Press, Princeton, NJ, 2018.
- [16] Zhou, K., Doyle, J., and Glover, K., *Robust and Optimal Control*, Vol. 40, Prentice Hall, Upper Saddle River, NJ, 1996.
- [17] Lewis, F., Vrabie, D., and Syrmos, V., *Optimal Control of Continuous-Time Systems*, Wiley, Hoboken, NJ, 2012, pp. 110–176, Chap. 3.
<https://doi.org/10.1002/9781118122631.ch3>
- [18] Gahinet, P., and Apkarian, P., "A Linear Matrix Inequality Approach to $\mathcal{H}_\infty$ Control," *International Journal of Robust and Nonlinear Control*, Vol. 4, No. 4, 1994, pp. 421–448.
<https://doi.org/10.1002/rnc.4590040403>
- [19] Wang, H.-S., Yung, C.-F., and Chang, F.-R., "Bounded Real Lemma and $\mathcal{H}_\infty$ Control for Descriptor Systems," *American Control Conference*, Inst. of Electrical and Electronics Engineers, New York, June 1997, pp. 2115–2119.
<https://doi.org/10.1109/ACC.1997.611064>
- [20] Bernstein, D., and Haddad, W., "LQG Control with an $\mathcal{H}_\infty$ Performance Bound: A Riccati Equation Approach," *IEEE Transactions on Automatic Control*, Vol. 34, No. 3, 1989, pp. 293–305.
<https://doi.org/10.1109/9.16419>
- [21] Limebeer, D., Anderson, B., and Hendel, B., "A Nash Game Approach to Mixed $\mathcal{H}_2/\mathcal{H}_\infty$ Control," *IEEE Transactions on Automatic Control*, Vol. 39, No. 1, 1994, pp. 69–82.
<https://doi.org/10.1109/9.273340>
- [22] Foley, M., and Schmitendorf, W. E., "On a Class of Nonzero-Sum, Linear-Quadratic Differential Games," *Journal of Optimization Theory and Applications*, Vol. 7, No. 5, 1971, pp. 357–377.
<https://doi.org/10.1007/BF00933999>
- [23] Limebeer, D. J. N., Anderson, B. D. O., Khargonekar, P. P., and Green, M., "A Game Theoretic Approach to $\mathcal{H}_\infty$ Control for Time-Varying Systems," *Siam Journal on Control and Optimization*, Vol. 30, No. 2, 1992, pp. 262–283.
<https://doi.org/10.1137/0330017>

- [24] Makkapati, V. R., Ridderhof, J., Tsiotras, P., Hart, J., and van Bloemen Waanders, B., "Desensitized Trajectory Optimization for Hypersonic Vehicles," *IEEE Aerospace Conference*, Inst. of Electrical and Electronics Engineers, New York, March 2021, pp. 1–10.
<https://doi.org/10.1109/AERO50100.2021.9438511>
- [25] Cesari, L., *Optimization - Theory and Applications*, Springer-Verlag, New York, 1983.
<https://doi.org/10.1007/978-1-4613-8165-5>
- [26] Khalil, H. K., *Nonlinear Systems*, 3rd ed., Prentice Hall, Upper Saddle River, NJ, 2002, Chap. 3.
- [27] Patterson, M. A., and Rao, A. V., "GPOPS-II: A MATLAB Software for Solving Multiple-Phase Optimal Control Problems Using hp-Adaptive Gaussian Quadrature Collocation Methods and Sparse Nonlinear Programming," *ACM Transactions Mathematical Software*, Vol. 41, No. 1, 2014, pp. 1–37.
<https://doi.org/10.1145/2558904>
- [28] Betts, J. T., *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, Soc. for Industrial and Applied Mathematics, Philadelphia, PA, 2010.
<https://doi.org/10.1137/1.9780898718577>
- [29] Lu, P., Forbes, S., and Baldwin, M., "Gliding Guidance of High L/D Hypersonic Vehicles," *AIAA Guidance, Navigation, and Control (GNC) Conference*, AIAA Paper 2013-4648, 2013.
<https://doi.org/10.2514/6.2013-4648>

D. Scheeres
Associate Editor