

Differential Dynamic Programming With Nonlinear Safety Constraints Under System Uncertainties

Gokhan Alcan  and Ville Kyrki 

Abstract—Safe operation of systems such as robots requires them to plan and execute trajectories subject to safety constraints. When those systems are subject to uncertainties in their dynamics, it is challenging to ensure that the constraints are not violated. In this letter, we propose Safe-CDDP, a safe trajectory optimization and control approach for systems under additive uncertainties and nonlinear safety constraints based on constrained differential dynamic programming (DDP). The safety of the robot during its motion is formulated as chance constraints with user-chosen probabilities of constraint satisfaction. The chance constraints are transformed into deterministic ones in DDP formulation by constraint tightening. To avoid over-conservatism during constraint tightening, linear control gains of the feedback policy derived from the constrained DDP are used in the approximation of closed-loop uncertainty propagation in prediction. The proposed algorithm is empirically evaluated on three different robot dynamics with up to 12 degrees of freedom in simulation. The computational feasibility and applicability of the approach are demonstrated with a physical hardware implementation.

Index Terms—Optimization and optimal control, constrained motion planning, planning under uncertainty, robot safety, motion and path planning.

I. INTRODUCTION

IN MANY real-world applications, robots are situated in uncertain environments with stochastic dynamics, where they are required to satisfy particular safety constraints such as collision avoidance or physical limits of their actuators. Within this context, *safety* can be defined as the feasibility and stability of a control policy that achieves the requirements of the desired task while satisfying the safety constraints considering uncertainties affecting the system.

Safety in robotics is an active research area and can be studied as trajectory optimization and control under constraints in a stochastic environment. In this context, model predictive control (MPC) is a useful framework, as it allows the optimization of state and input trajectories based on an objective function under chance constraints [1]. In direct optimization of such a problem, chance constraints are typically transformed to deterministic ones by constraint tightening with precomputed fixed




Fig. 1. Safe-CDDP successfully generates optimal trajectories under system uncertainties for complex and underactuated robots. Iteratively generated safe trajectories (orange) and the real trajectories (green) traveled by 2D car-like robot (Left) and 3D quadrotor robot (Right) in constrained environments.

controller gains [2] or robust constraints are defined using large confidence bounds of uncertainties [3]. In both cases, the true effect of feedback on uncertainty propagation is generally omitted, which leads to excessive conservatism and related loss of performance.

To address the question of how to avoid the loss of performance related to conservative constraint tightening, we propose a novel safe trajectory optimization and control approach, which we call Safe Constrained Differential Dynamic Programming (Safe-CDDP). The proposed method extends standard differential dynamic programming (DDP) to handling nonlinear state and input constraints in the presence of additive system uncertainties. The constraints are handled by modeling the problem as a chance-constrained optimal control problem under uncertainty and using constraint tightening [4] to turn it into a deterministic constrained problem, thus generating an appropriate safety margin for the constraints. To reduce the uncertainty in the prediction horizon, we incorporate the control gains of a locally optimal affine feedback policy derived in the backward pass of DDP, which yields a good approximation for optimized trajectories at convergence and avoids excessive conservatism.

To apply the method in real-time feedback control, we also need to address the issue that the tightened constraints are only virtual in the sense that the system noise may cause them to be violated within the safety margin even when they are tightened appropriately. Thus, the utilization of such a constrained formulation in an optimizer creates the risk of a constraint violation, which the optimization process needs to handle appropriately.

The main contributions of our work are:

- 1) A principled way of reflecting the effect of feedback on uncertainty propagation in prediction using dynamic gains from the locally optimal control derived in constrained DDP formulation to avoid excessive conservatism.

Manuscript received June 21, 2021; accepted December 19, 2021. Date of publication January 7, 2022; date of current version January 20, 2022. This letter was recommended for publication by Associate Editor S. Weiss and Editor C. Gosselin upon evaluation of the reviewers' comments. This work was supported by the Academy of Finland B-REAL Project under Grant 328399. (Corresponding author: Gokhan Alcan.)

The authors are with the Intelligent Robotics Group, Department of Electrical Engineering and Automation (EEA), Aalto University, 02150 Espoo, Finland (e-mail: gokhan.alcan@aalto.fi; ville.kyrki@aalto.fi).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2022.3141192>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2022.3141192

- 2) A safe model predictive control approach (Safe-CDDP) with general nonlinear constraints in the presence of additive system uncertainties.
- 3) Simulation studies for three types of robot platform dynamics demonstrating that Safe-CDDP ensures stochastic constraint satisfaction for complex and under-actuated robots in safety-required tasks (Fig. 1).
- 4) Validation via hardware implementation for a differential drive robot that the proposed method is computationally feasible using a low-power CPU.

Compared to existing formulations to similar problems, the proposed approach simultaneously modes the closed-loop effects of additive uncertainties using dynamic gains instead of pre-computed fix gains, and it is not limited to box constraints but includes general nonlinear state and input constraints.

II. RELATED WORK

There exist different notions of safety in robotics control literature. In safety-critical control, safety has been investigated through the terms of set-invariance and reachability analysis. A safety set is defined as a set of allowed states for a system and the task of controlling the system becomes ensuring the invariance of this safety set [5]. Reachability of such target sets is formulated as an optimal control problem for safety analysis and synthesis of safe controllers [6]. Even though dynamic programming principles similar to the scope of this paper can be utilized to solve reachability-based value functions, those methods suffer from the usual curse of dimensionality, solving them online on embedded hardware is mostly infeasible and they are generally overly conservative.

Robust model predictive control approaches involve direct optimization that enables to employ of nonlinear state and input constraints by considering the uncertainty with a worst-case safety bound. Since robust MPC is computationally complex, tube MPCs [7] are designed as an approximation for robust MPCs but they require to define tube geometry dynamics and in some cases time-consuming offline computation [8]. Moreover, most of the approaches with fixed tube sizes produce overly conservative tubes due to the lack of knowledge regarding state-dependent uncertainty. Lopez *et al.* [7] proposed to optimize both the tube geometry and open-loop trajectory simultaneously to reduce conservativeness. However, they assume that the system has the same number of outputs to be controlled as inputs, which makes the method only feasible for relatively low-dimensional and particular kinds of nonlinear systems [7]. On the other hand, Differential Dynamic Programming (DDP) formulation allows the decomposition of a trajectory optimization problem into smaller ones by limiting the state space to a quadratic trust-region around a current reference solution [9], which results in only local optimality but reduces the processing cost dramatically. Although there is no straightforward way to incorporate nonlinear state and input constraints in the DDP formulation, recent attempts are promising. Next, we provide an overview of previous work related to constrained deterministic and stochastic variants of DDPs.

A. Constrained Deterministic DDPs

Some recent studies target the design of DDP with input constraints [10], [11], [12]. However, for this work methods that incorporate also state constraints are more relevant. Xie

et al. [13] utilized Karush-Kuhn-Tucker (KKT) conditions and modified the DDP formulation in the presence of nonlinear state and input constraints according to the active set of constraints at each time step to keep the trajectories feasible. Aoyama *et al.* [14] utilized the work developed in [13] and extended it through an Augmented Lagrangian method by considering a set of penalty functions. Pavlov *et al.* [15] introduced primal-dual interior-point DDP to handle nonlinear state and input constraints. The method is quite promising since it does not require to identify the active set of constraints, but it needs second-order derivatives of the state constraint. Augmented Lagrangian TRajjectory Optimizer (ALTRO) [16] was also recently developed to fuse the advantageous sides of direct methods and DDPs. It handles general deterministic nonlinear state and input constraints with fast convergence and also enables to start with infeasible initial trajectories.

The aforementioned methods achieve promising performances in handling nonlinear constraints and increase the usability of DDPs in robotics tasks. However, in the presence of system uncertainty, they lack safety provisions, which are the primary issue studied in this work.

B. Stochastic DDPs

Todorov and Li [17] developed input constrained iterative Linear Quadratic Gaussian control for nonlinear stochastic systems involving multiplicative noise with standard deviation proportional to the control signals. Theodorou *et al.* [18] then extended it to stochastic DDPs for state and control multiplicative process noise. In their work, the derived control policy does not depend on the statistical characteristics of the noise when the stochastic dynamics have only additive noise. Pan and Theodorou [19] introduced a data-driven probabilistic unconstrained DDP using Gaussian processes to address the uncertainty for dynamics models. Celik *et al.* [20] proposed to model state and action physical limits as probabilistic chance constraints in a DDP structure to avoid catastrophic greedy updates. The method does not apply to tasks that need to employ an arbitrary nonlinear constraints on states such as obstacle avoidance, since it involves only linear box-constraints on states and actions as their limits. Ozaki *et al.* [21] proposed a modification for stochastic DDP developed in [18] considering the disturbances and the uncertainties as random state perturbations for an unconstrained deterministic dynamical system and employing unscented transform. They then further improved the method by tube stochastic DDP [22] to handle the control constraints but nonlinear constraints on states were not included.

Our approach widens the extent of stochastic DDPs by considering general nonlinear state and input constraints in the presence of additive uncertainties.

III. PROPOSED METHOD

We consider a nonlinear discrete-time dynamical system of the form

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) + \boldsymbol{\omega}_k, \quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^n$ and $\mathbf{u}_k \in \mathbb{R}^m$ are the system state and the control input at time step k , respectively. $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a nonlinear state transition function which is assumed to be twice differentiable. $\boldsymbol{\omega}_k$ is assumed to be spatially uncorrelated independent and identically distributed noise, drawn from a zero

mean Gaussian distribution with a known covariance matrix, $\omega_k \sim \mathcal{N}(\mathbf{0}, \Sigma^\omega)$.

For a given initial state \mathbf{x}_0 , a goal state \mathbf{x}_{goal} , and a time horizon N , the aim is to find an input trajectory $\{\mathbf{u}_0, \dots, \mathbf{u}_{N-1}\}$ that minimizes the expectation of an objective function

$$J(\mathbf{u}_0, \dots, \mathbf{u}_{N-1}) = \ell^f(\mathbf{x}_N) + \sum_{k=0}^{N-1} \ell(\mathbf{x}_k, \mathbf{u}_k), \quad (2)$$

where $\ell^f: \mathbb{R}^n \rightarrow \mathbb{R}$ and $\ell: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ are the final cost and the running cost, respectively. By considering the input limitations and external constraints on the states, chance-constrained MPC for this problem can be formulated as

$$\begin{aligned} \min_{\mathbf{u}_0, \dots, \mathbf{u}_{N-1}} \quad & \mathbb{E} \left(J(\mathbf{u}_0, \dots, \mathbf{u}_{N-1}) \right) \\ \text{subject to} \quad & \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) + \omega_k, \\ & \mathbf{u}_{min} \leq \mathbf{u}_k \leq \mathbf{u}_{max}, \\ & \Pr[g(\mathbf{x}_k, \mathbf{u}_k) \leq \mathbf{0}] > \beta_k, \\ & \mathbf{x}_0 = \bar{\mathbf{x}}, \end{aligned} \quad (3)$$

for all $k = 0, \dots, N-1$. Here we simply add box-constraint on control inputs defined by the boundaries \mathbf{u}_{min} and \mathbf{u}_{max} . \mathbf{g} is a vector of c constraints in the form of differentiable functions representing the deterministic state constraints and $\beta_k = [\beta_{1,k}, \dots, \beta_{c,k}]$ is a vector of minimum satisfaction probabilities for these constraints for time step k .

Future predicted states in MPC formulation will result in stochastic distribution due to the noise ω_k in state transition. This leads to chance constraints on states, therefore the safety of the predicted trajectory is determined by the probability of constraint satisfaction. Typically, sufficiently safe control actions are desired to cope with the effect of the uncertainties. To achieve this, chance constraints can be converted into deterministic ones by tightening the constraints [23], such that for each original constraint function $g_i(\cdot)$ the tightened constraint is denoted by $\tilde{g}_i(\cdot)$. Overly conservative tightening based on the propagated uncertainty such as employing open-loop uncertainty propagation narrows dramatically the admissible ranges of the states and can make the optimization intractable or even unsolvable [24].

In order to reflect the effect of feedback on uncertainty propagation in prediction and avoid over-conservatism, we propose the following constrained differential dynamic programming-based model predictive control structure. An initial trajectory is optimized in successive backward and forward passes, while the deterministic constraints are updated periodically in constraint tightening. In each backward pass (described in detail in Section III-A), the derivatives of the action-value function ($Q_k(x_k, u_k)$) and optimum value function ($V_k(x_k)$) are calculated backward starting from the final step in the horizon (N) using the current estimated (or initialized) trajectories and pre-defined objective function as in (2). In forward passes (Section III-B), the nominal state-control trajectory is updated in forward starting from the initial state by using the derivatives obtained in the backward pass and performing local optimizations considering the constraints. In constraint tightening (Section III-C), chance constraints in (3) are converted into deterministic ones by tightening the constraints using the derived locally optimal control from the backward pass in the uncertainty propagation

through prediction. These steps will be detailed in the next sections together with the entire algorithm.

A. Backward Pass

The backward pass that follows the formulation in [13] with the assumption of deterministic constraints briefly explained here for completeness. Using standard DDP formulation, action-value function $Q_k(\mathbf{x}_k, \mathbf{u}_k)$ can be approximated as a quadratic function as

$$\begin{aligned} Q(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}) &\approx Q + Q_{\mathbf{x}}^\top \delta\mathbf{x} + Q_{\mathbf{u}}^\top \delta\mathbf{u} \\ &\quad + \frac{1}{2}(\delta\mathbf{x}^\top Q_{\mathbf{x}\mathbf{x}} \delta\mathbf{x} + \delta\mathbf{u}^\top Q_{\mathbf{u}\mathbf{u}} \delta\mathbf{u}) \\ &\quad + \delta\mathbf{x}^\top Q_{\mathbf{x}\mathbf{u}} \delta\mathbf{u} \end{aligned} \quad (4)$$

where $\delta\mathbf{x}$ and $\delta\mathbf{u}$ are the deviations about the nominal action-state pair (\mathbf{x}, \mathbf{u}) . The derivatives of Q function are then

$$\begin{aligned} Q_{\mathbf{x}} &= \ell_{\mathbf{x}} + f_{\mathbf{x}}^\top V'_{\mathbf{x}}, \\ Q_{\mathbf{u}} &= \ell_{\mathbf{u}} + f_{\mathbf{u}}^\top V'_{\mathbf{x}}, \\ Q_{\mathbf{x}\mathbf{x}} &= \ell_{\mathbf{x}\mathbf{x}} + f_{\mathbf{x}}^\top V'_{\mathbf{x}\mathbf{x}} f_{\mathbf{x}} + (V'_{\mathbf{x}} f_{\mathbf{x}\mathbf{x}}), \\ Q_{\mathbf{u}\mathbf{u}} &= \ell_{\mathbf{u}\mathbf{u}} + f_{\mathbf{u}}^\top V'_{\mathbf{x}\mathbf{x}} f_{\mathbf{u}} + (V'_{\mathbf{x}} f_{\mathbf{u}\mathbf{u}}), \\ Q_{\mathbf{x}\mathbf{u}} &= \ell_{\mathbf{x}\mathbf{u}} + f_{\mathbf{x}}^\top V'_{\mathbf{x}\mathbf{x}} f_{\mathbf{u}} + (V'_{\mathbf{x}} f_{\mathbf{x}\mathbf{u}}). \end{aligned} \quad (5)$$

where V is the value function (see [13], [14] for details). In order to simplify the notations, we dropped the time step indices, used prime to indicate the next time step and used subscripts for the derivatives. The terms in parentheses in (5) (second derivatives of the dynamics) provide better local fidelity by capturing the nonlinear effects of the system. Calculation of those tensors can be costly especially for systems with complex dynamics and high-dimensional states such as humanoids, and therefore they can be discarded to obtain faster convergence by sacrificing fidelity, where in that case DDP refers to iterative LQR. The method developed in this study is applicable to both cases.

In unconstrained situations, locally optimal control deviations can be computed by minimizing (4) with respect to $\delta\mathbf{u}$ resulting in

$$\delta\mathbf{u}^* = -Q_{\mathbf{u}\mathbf{u}}^{-1}(Q_{\mathbf{u}\mathbf{x}}\delta\mathbf{x} + Q_{\mathbf{u}}) \triangleq \bar{\mathbf{K}}\delta\mathbf{x} + \bar{\mathbf{d}} \quad (6)$$

where $\bar{\mathbf{K}}$ is the linear feedback gain and $\bar{\mathbf{d}}$ is the affine term.

Under state and/or input constraints, the optimal control problem turns into

$$\begin{aligned} \min_{\delta\mathbf{u}} \quad & Q(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}), \\ \text{subject to} \quad & \tilde{\mathbf{g}}(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}) \leq \mathbf{0}. \end{aligned} \quad (7)$$

where the constraint vector $\tilde{\mathbf{g}}$ includes the deterministic state constraints obtained in the constraint tightening step including input limitations. In order to incorporate the constraints in the estimation of action-value function (4), constraints' first order approximations are obtained as

$$\begin{aligned} \tilde{\mathbf{g}}(\mathbf{x} + \delta\mathbf{x}, \mathbf{u} + \delta\mathbf{u}) &\approx \tilde{\mathbf{g}}(\mathbf{x}, \mathbf{u}) \\ &\quad + \underbrace{\tilde{\mathbf{g}}_{\mathbf{u}}(\mathbf{x}, \mathbf{u}) \delta\mathbf{u}}_{\triangleq \mathbf{C}} + \underbrace{\tilde{\mathbf{g}}_{\mathbf{x}}(\mathbf{x}, \mathbf{u}) \delta\mathbf{x}}_{\triangleq \mathbf{D}} \end{aligned} \quad (8)$$

Xie *et al.* [13] proposed to check the set of active constraints during the iterations to make sure that the analytical approximation of the value function around the nominal trajectories still yields a good approximation. To achieve this, all active constraints are included in a set which contains all equality constraints that are always active and inequality constraints for which the value is greater than $-\epsilon$ ($\tilde{g}_i > -\epsilon$) where ϵ is a small constant for numerical stability.

The optimization problem in (7) is then converted to

$$\begin{aligned} \min_{\delta \mathbf{u}} \quad & Q_{\mathbf{u}}^\top \delta \mathbf{u} + \frac{1}{2} (\delta \mathbf{u}^\top Q_{\mathbf{uu}} \delta \mathbf{u}) + \delta \mathbf{x}^\top Q_{\mathbf{xu}} \delta \mathbf{u}, \\ \text{subject to} \quad & \mathbf{C} \delta \mathbf{u} = \mathbf{D} \delta \mathbf{x}. \end{aligned} \quad (9)$$

An analytical solution to this problem through KKT conditions [25] can be found by solving the pair of equations

$$\begin{aligned} Q_{\mathbf{uu}} \delta \mathbf{u} + Q_{\mathbf{ux}} \delta \mathbf{x} + Q_{\mathbf{u}} + \mathbf{C}^\top \lambda &= \mathbf{0} \\ \mathbf{C} \delta \mathbf{u} - \mathbf{D} \delta \mathbf{x} &= \mathbf{0} \end{aligned} \quad (10)$$

where λ is a vector of Lagrangian multipliers. By solving (10) with pruned matrices $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{D}}$ that ensures nonexistence of negative λ values (check [13] for details), locally optimal input deviations can be again expressed as a function of state deviations as follows:

$$\delta \mathbf{u}^* = \mathbf{K} \delta \mathbf{x} + \mathbf{d} \quad (11)$$

where the control parameters \mathbf{K} and \mathbf{d} are now calculated as

$$\begin{aligned} \mathbf{K} &= -Q_{\mathbf{uu}}^{-1} Q_{\mathbf{ux}} \\ &\quad - Q_{\mathbf{uu}}^{-1} \tilde{\mathbf{C}}^\top (\tilde{\mathbf{C}} Q_{\mathbf{uu}}^{-1} \tilde{\mathbf{C}}^\top)^{-1} \tilde{\mathbf{D}} \\ &\quad + Q_{\mathbf{uu}}^{-1} \tilde{\mathbf{C}}^\top (\tilde{\mathbf{C}} Q_{\mathbf{uu}}^{-1} \tilde{\mathbf{C}}^\top)^{-1} \tilde{\mathbf{C}} Q_{\mathbf{uu}}^{-1} Q_{\mathbf{ux}}, \\ \mathbf{d} &= -Q_{\mathbf{uu}}^{-1} \left(Q_{\mathbf{u}} - \tilde{\mathbf{C}}^\top (\tilde{\mathbf{C}} Q_{\mathbf{uu}}^{-1} \tilde{\mathbf{C}}^\top)^{-1} \tilde{\mathbf{C}} Q_{\mathbf{uu}}^{-1} Q_{\mathbf{u}} \right). \end{aligned} \quad (12)$$

It is observed that $\delta \mathbf{u}^*$ at time step k requires the knowledge of first and second derivatives of the value function for next time step ($k+1$) as shown in (5). By plugging the optimal control found in (11) into the approximated action-value function (4), the value function can be approximated as

$$V \approx \frac{1}{2} \delta \mathbf{x}^\top V_{\mathbf{xx}} \delta \mathbf{x} + V_{\mathbf{x}}^\top \delta \mathbf{x} + \mathbf{c} \quad (13)$$

where \mathbf{c} is a constant term and the derivative terms can be calculated as

$$\begin{aligned} V_{\mathbf{x}} &= Q_{\mathbf{x}} + \mathbf{K} Q_{\mathbf{u}} + \mathbf{K}^\top Q_{\mathbf{uu}} \mathbf{d} + Q_{\mathbf{ux}}^\top \mathbf{d}, \\ V_{\mathbf{xx}} &= Q_{\mathbf{xx}} + \mathbf{K}^\top Q_{\mathbf{uu}} \mathbf{K} + \mathbf{K}^\top Q_{\mathbf{ux}} + Q_{\mathbf{ux}}^\top \mathbf{K}. \end{aligned} \quad (14)$$

$V_{\mathbf{x}}$ and $V_{\mathbf{xx}}$ at final step in the horizon (N) can be calculated as the first and second derivatives of the final cost function (ℓ^f), respectively. In this way, derivatives of the Q function (5) for each step in the predicted trajectory can be consecutively calculated backward starting from the final step.

B. Forward Pass

As the successive step of a backward pass to update the trajectories, forward pass implemented in [13] is briefly explained here for the sake of completeness.

To enforce that the constraints are also satisfied in the forward pass, we formulate the forward pass as the solution of a QP

problem

$$\begin{aligned} \min_{\delta \mathbf{u}} \quad & Q_{\mathbf{u}}^\top \delta \mathbf{u} + \frac{1}{2} (\delta \mathbf{u}^\top Q_{\mathbf{uu}} \delta \mathbf{u}) + \delta \mathbf{x}^\top Q_{\mathbf{xu}} \delta \mathbf{u}, \\ \text{subject to} \quad & \tilde{g}(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u}) \leq 0, \\ & \mathbf{u}_{\min} \leq \mathbf{u} + \delta \mathbf{u} \leq \mathbf{u}_{\max}. \end{aligned} \quad (15)$$

The solution of the problem will then represent the optimal deviation in control input $\delta \mathbf{u}$ while respecting the constraints.

The state is then updated by forward integration using the optimal input deviation $\delta \mathbf{u}^*$ as

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k + \delta \mathbf{u}^*) \quad (16)$$

C. Constraint Tightening

This section explains how the chance constraints formulated in (3) are transformed into the deterministic ones which are required for backward and forward passes. In order to obtain the constraints in a deterministic form, we tighten the constraints using the propagated uncertainty through prediction, which can be decreased efficiently once a feedback policy is employed in the optimization.

Nonlinear transition dynamics $f(\cdot)$ can be approximated to first order by

$$\mathbf{x}_{k+1} \approx \mathbf{x}_k + f_{\mathbf{x}} \delta \mathbf{x} + f_{\mathbf{u}} \delta \mathbf{u} \quad (17)$$

Employing the locally optimal control from (11) results in

$$\mathbf{x}_{k+1} \approx \mathbf{x}_k + f_{\mathbf{x}} \delta \mathbf{x} + f_{\mathbf{u}} \mathbf{K} \delta \mathbf{x} + \mathbf{d} = \mathbf{x}_k + (f_{\mathbf{x}} + f_{\mathbf{u}} \mathbf{K}) \delta \mathbf{x} + \mathbf{d} \quad (18)$$

Closed-loop uncertainty propagation through prediction for the system in (1) then can be approximated as:

$$\Sigma_{k+1}^{\mathbf{x}} = (f_{\mathbf{x}} + f_{\mathbf{u}} \mathbf{K}_k)^\top \Sigma_k^{\mathbf{x}} (f_{\mathbf{x}} + f_{\mathbf{u}} \mathbf{K}_k) + \Sigma_{k+1}^{\omega} \quad (19)$$

Since constrained DDP performs a local optimization at each step through prediction, general nonlinear constraint can be locally interpreted as a half-space constraint for that particular step. By utilizing the relationship between probabilistic and robust invariant sets presented in [4] for linear time-invariant systems, the effect of the propagated uncertainty could be approximately taken into account in constraint definitions for each step as follows:

$$\tilde{g}(\mathbf{x}, \mathbf{u}, \Sigma^{\mathbf{x}}) \triangleq g(\mathbf{x}, \mathbf{u}) + \phi^{-1}(\beta) \sqrt{\tilde{g}_x^\top \Sigma^{\mathbf{x}} \tilde{g}_x} \quad (20)$$

where ϕ^{-1} is the quantile function of the standard normal distribution. Fig. 2 shows that the additional part of the right-hand side in (20) serves for altering the boundaries of the obstacles and generates safety margins based on the propagated uncertainty and the dynamics of the robot. Moreover, the β value adjusts the amount of uncertainty to be considered, which proportionally increases the safety limits.

Although different β values can be assigned for different constraints at different time steps, we did not find it beneficial for safety-required tasks and selected all the β values the same for simplicity. Once all the chance-constraints are reformulated, the MPC formulation of (3) can be written in a deterministic




Fig. 2. Changes in the boundaries of the obstacles after tightening the constraints with different β values.

form as:

$$\begin{aligned}
 & \min_{\mathbf{u}_0, \dots, \mathbf{u}_{N-1}} \ell^f(\mathbf{x}_N) + \sum_{k=0}^{N-1} \ell(\mathbf{x}_k, \mathbf{u}_k) \\
 & \text{subject to } \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \\
 & \quad \mathbf{u}_{min} \leq \mathbf{u}_k \leq \mathbf{u}_{max}, \\
 & \quad \Sigma_{k+1}^x = h(\mathbf{x}_k, \mathbf{u}_k, \mathbf{K}_k, \Sigma_k^x) \text{ Eq.(19)} \\
 & \quad \tilde{\mathbf{g}}(\mathbf{x}_k, \mathbf{u}_k, \Sigma_k^x) \leq \mathbf{0}, \\
 & \quad \mathbf{x}_0 = \bar{\mathbf{x}},
 \end{aligned} \tag{21}$$

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

In order to employ a feedback policy in the optimization, Hewing *et al.* [2] restricted the policy to be a linear state feedback using pre-computed or fixed linear gains. Here, we inherently have a locally optimal affine feedback policy (11) thanks to the structure of DDP. Moreover, that policy with the controller gains derived in the backward pass (12) indeed yields a good approximation for overall nonlinear constrained DDP optimization at convergence. Consequently, the method neither needs to employ some precomputed fixed gains nor needs to solve again a finite LQR problem around the current nominal trajectory.

D. Algorithm in a Nutshell

Here we describe the Safe-CDDP algorithm as a whole with some practicalities to employ it in an MPC framework. Even though the initialization of such trajectory optimization algorithms is still an open problem, applying an unconstrained DDP for a relatively straightforward temporary goal state is indeed a convenient practice to initialize our method with suboptimal state-input trajectories. After initialization, our method assumes that there exists a sufficiently large safe region around the real goal state (\mathbf{x}_{goal}) and implements a model predictive control procedure with decreasing horizon until the robot gets close to the goal state (Algorithm 1). Depending on the state dimensions,

Algorithm 1: MPC using Safe-CDDP

```

1 Init:  $N, \mathbf{x}_0, \mathbf{x}_{goal}, \mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{u}_0, \dots, \mathbf{u}_{N-1}$ 
2 while robot is not close to  $\mathbf{x}_{goal}$  do
3   for iter = 0, ...,  $n_1$  do
4     function Backward_Pass
5        $V_x(\mathbf{x}_N) \leftarrow \ell_x^f(\mathbf{x}_N)$ 
6        $V_{xx}(\mathbf{x}_N) \leftarrow \ell_{xx}^f(\mathbf{x}_N)$ 
7       for  $k = N-1, \dots, 0$  do
8         Calculate derivatives of  $Q$  (Eq. 5)
9         Calculate derivatives of  $V$  (Eq. 14)
10        Save  $Q_{uu}, Q_{xu}, Q_u, K$ 
11     function Forward_Pass
12        $\mathbf{x} \leftarrow \mathbf{x}_0$ 
13       for  $k = 0, \dots, N-1$  do
14          $\delta \mathbf{x} \leftarrow \mathbf{x} - \mathbf{x}_k$ 
15         Solve QP for  $\delta \mathbf{u}$  (Eq. 15)
16         if feasible solution then
17            $\mathbf{x}_{k+1} \leftarrow f(\mathbf{x}, \mathbf{u}_k + \delta \mathbf{u})$ 
18            $\mathbf{x} \leftarrow \mathbf{x}_{k+1}$ 
19         else
20           Regularization
21           break
22     if iter %  $n_2 == 0$  then
23       function Constraint_Tightening
24         Calculate closed-loop uncertainty
25         propagation (Eq. 19)
26         Update constraints (Eq. 20)
27   Apply  $\mathbf{u}_0$  to robot
28    $\mathbf{x}_0 \leftarrow$  measurement
29    $\mathbf{x}_1, \dots, \mathbf{x}_{N-1} \leftarrow \mathbf{x}_2, \dots, \mathbf{x}_N$ 
30    $\mathbf{u}_0, \dots, \mathbf{u}_{N-2} \leftarrow \mathbf{u}_1, \dots, \mathbf{u}_{N-1}$ 
31    $N \leftarrow N-1$ 

```

horizon length (N) and the number of constraints, iteration number (n_1) required to optimize the trajectory should be chosen beforehand. Instead of a fixed iteration number for trajectory optimization, one can iterate forward and backward passes until a predetermined convergence is achieved. However, similar to the selection of horizon length for most of the trajectory planning algorithms, the sufficient number of iterations (n_1) that saves computational resources is also a hyper-parameter that can be tuned by trial and error. In this study, we choose $n_1 = 10$ and keep it fixed for simplicity.

Solving (21) in a DDP framework revealed a chicken-egg problem, where the closed-loop uncertainty propagation requires the control gains, and the control gains are calculated in the backward pass of DDP constrained by reformulated chance-constraints using uncertainty propagation. To overcome this dilemma, we initially neglect the noise in transition dynamics at the beginning of backward and forward passes and treat the constraints as deterministic by removing the constraint satisfaction probability. Afterward, the deterministic constraints are updated via constraint tightening.

Note that in this framework, the action-value function is approximated up to second order for trajectory optimization,

and the first-order approximation of nonlinear transition dynamics is employed in constraint tightening through uncertainty propagation. Frequent updates in closed-loop uncertainty propagation lead to employing control gains that do not yield a good approximation for action-value function, and therefore it causes oscillations in predicted trajectories and prevents the optimizer from converging. Therefore, we suggest a slower update schedule ($n_2 = n_1/2 = 5$ as a rule of thumb) where iterations get settled for better control gains, then the constraints are updated and the estimated trajectory is optimized for updated constraints again before applying the first input to the robot as depicted in Algorithm 1. The hyper-parameters n_1 and n_2 can be tuned by trial and error depending on the nonlinearity of both system dynamics and constraints. The less complexity allows less iteration and frequent constraint updates.

Finally, it should be noted that the regularization scheme (line 19 in Algorithm 1) and the choice of optimization parameters such as step size and trust region in search space affect the convergence of the algorithm. In our implementation, we used a diminishing trust-region for infeasible solutions and ensured the numerical stability of matrix inversions by adding a regularization term (Please refer to [13], [14], [26] for regularization details).

Due to the stochastic setting, the constraints are satisfied with the desired probability β as shown in (3). The quadratic programming solver is able to detect that the solution is infeasible such that all constraints cannot be satisfied, thus, the approach is able to detect infeasibility. However, it is possible that the tightened constraints are not satisfied even when the original constraints are, due to noise in the system. If the aforementioned regularization scheme is not enough to retain feasibility, the tightened constraints can be relaxed as long as the safety requirement of the robotic application allows.

IV. EXPERIMENTS

In this section, we present repetitive simulation results for statistical evaluation and demonstrate hardware implementation to assess the proposed method in the scope of computational feasibility and applicability to real robots.

A. Dynamical Systems in Simulation

In this section, we statistically evaluate our Safe-CDDP algorithm on three different robot dynamics in simulation: (i) 2D point robot (4 states, 2 inputs) [13], (ii) 2D car-like robot (4 states, 2 inputs) [13], and (iii) 3D quadrotor robot (12 states, 4 inputs) [29], [30]. Prediction horizons (N) for the robots are selected as 100, 120 and 50, respectively. The defined task for each robot dynamics is reaching a goal position while avoiding stationary obstacles. The characteristics of process noise (Σ_k^x) is assumed to be known and the positions of the obstacles and the robot itself are assumed to be externally provided without noise. We also assume that there exist safe regions around the initial and goal positions. Once the robot arrived in the safe region around the goal, the task is considered successfully achieved. Another control algorithm could be subsequently triggered to maintain the robot in that region with accurate positioning, which is not a part of this study.

Stationary obstacles in the environment are formulated as circular constraints as defined in [13] and [14]. For the 3D quadrotor, it is also possible to use spherical constraints but we preferred

TABLE I
CONSTRAINT VIOLATION RESULTS

	CDDP [13]	Safe-CDDP		
2D Point Robot		$\beta=0.90$	$\beta=0.95$	$\beta=0.99$
Num. of violated episodes	64	14	8	0
Avg. in violated episodes	1.98	1.43	1.38	0
Total avg. of violations	1.27	0.20	0.11	0
Car-like Robot		$\beta=0.90$	$\beta=0.95$	$\beta=0.99$
Num. of violated episodes	38	2	0	0
Avg. in violated episodes	1.63	1.0	0	0
Total avg. of violations	0.62	0.02	0	0
3D Quadrotor Robot		$\beta=0.90$	$\beta=0.95$	$\beta=0.99$
Num. of violated episodes	66	34	27	15
Avg. in violated episodes	5.7	4.26	5.19	4.07
Total avg. of violations	3.76	1.45	1.4	0.61

to choose cylindrical constraints to test the performance in a more challenging scenario. In addition to obstacle avoidance, i.e., nonlinear state constraints, the limits on control inputs are also employed. Only for the quadrotor robot, Safe-CDDP is implemented in iLQR form for simplicity.

Once the task configuration is fixed for a robot, it was performed repeatedly 100 times for each of 4 different β values {0.5, 0.90, 0.95, 0.99}. Note that the case with $\beta = 0.5$ corresponds to classical CDDP without safety precautions, since $\phi^{-1}(0.5) = 0$.

B. Simulation Results

In order to assess the safety of the obtained trajectories, we utilized 3 metrics regarding constraint violations (Table I). Once the robot finishes the task with either success or failure, we call it one episode and repeat the same task with a fixed configuration for 100 episodes. “Violated episode” refers to an episode in which the robot violates a constraint, i.e., collides with an obstacle, at least once. Average in violated episodes is the ratio of the total number of collisions and the number of violated episodes. Similarly, the total average of violations is the ratio of the total number of collisions and the total number of episodes, i.e., 100 in our case.

Table I depicts that Safe-CDDP algorithm provides safety in terms of obstacle collision due to uncertainties. By selecting $\beta=0.99$, Safe-CDDP performed well for point robot and car-like robot without any constraint violation, where at least one constraint violation occurred in 64 and 38 percent of the tasks for these robots when the safety precaution was removed ($\beta=0.5$). Selecting $\beta=0.95$ was even sufficient for the car-like robot to obtain safe trajectories with zero collision. This implies that increasing β values, i.e., increasing the confidence bound on the approximated uncertainty propagation leads optimizer to be more conservative and force it to find trajectories away from the obstacles if admissible.

It should be noted that the performance of Safe-CDDP highly depends on the admissible region of states possible for the task definition. For the 3D quadrotor robot, we deliberately select a goal position above the cylinder to assess the capabilities of the method. Fig. 3 presents all the traveled trajectories by quadrotor using Safe-CDDP with different β values. Once β is set to 0.5, i.e., in the case of classical CDDP, the majority (66%) of the tasks were finished with violations. In Safe-CDDP case with $\beta=0.99$, collision avoidance was not satisfied completely as in the cases of




Fig. 3. Comparison of traveled trajectories of a 3D quadrotor in constrained environment. (a): Classical CDDP without safety precaution. Safe-CDDP with β values of (b): 0.90, (c): 0.95 and (d): 0.99. Trajectories in blue represent constraint respected ones and the ones in red for constraint violations.




Fig. 4. Comparison of CDDP, constraint tightening with fixed gains, and Safe-CDDP for trajectory optimization of a car-like robot in terms of (a) traveled trajectories (b) uncertainty propagation in position, and (c) Frobenius norm of gains.

point robot and car-like robot, but it is dramatically improved by diminishing from %66 to %15. This indicates that Safe-CDDP can achieve remarkable performances in decreasing the number of constraint violations for complex and under actuated systems in the presence of system uncertainties.

We further compared our method with the one that uses fixed gains in constraint tightening, similar to the closest collocation method, called Cautious MPC [2]. An ancillary linear controller was employed in constraint tightening with $\beta = 0.99$ for all prediction steps. Even though the fixed controller gains reduced the uncertainty propagation in position up to 40 prediction steps, it started to increase exponentially afterward (Fig. 4(b)). Constraint tightening based on this propagated uncertainty resulted in a more conservative trajectory (Fig. 4(a)). On the other hand, Safe-CDDP avoids excessive conservatism with the same confidence bounds ($\beta = 0.99$) by increasing the norm of gains through the prediction horizon appropriately (Fig. 4(c)).

C. Hardware Implementation

As an experimental setup, we employed Turtlebot3 Waffle Pi that is a differential drive type mobile robot. In order to test Safe-CDDP on Turtlebot, we placed two circular objects with different sizes in the robot's environment as stationary obstacles. Gmapping was used to map the environment and the fusion of Adaptive Monte Carlo Localization (AMCL) and odometry were employed to retrieve the states of the robot (x, y, θ) with an estimation uncertainty. Circular obstacles with the radius of 0.15 and 0.11 m were detected at ($x:0.85, y:0$) and ($x:0.5, y:0.85$) locations, respectively. Turtlebot was initially located in

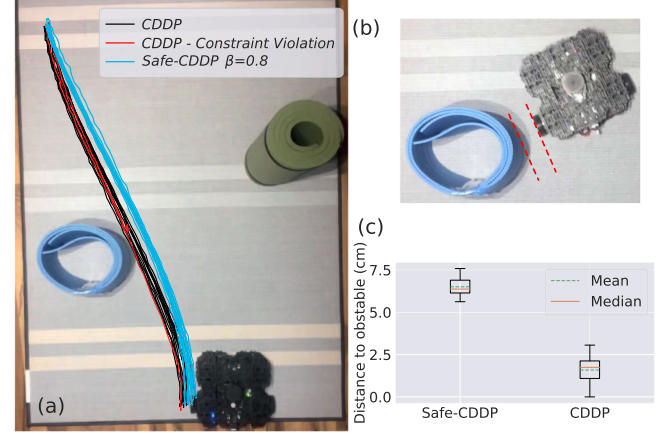


Fig. 5. (a) Comparison of traveled trajectories optimized by CDDP vs Safe-CDDP (b) Safety is determined by the distance between the wheel of Turtlebot and closest obstacle (c) Distributions of distance to closest obstacle.

($x:0, y:0, \theta:0$) and reaching to goal state ($x:1.4, y:0.6, \theta:0$) via Safe-CDDP was performed with the prediction horizon $N = 90$. Standard deviations of position noise retrieved by AMCL were approximately 0.001.

Safe-CDDP algorithm was implemented as a ROS node in C++ using OSQP [27] and Eigen [28] libraries. The execution of Safe-CDDP node was tested on both Turtlebot's CPU (Raspberry Pi-4 Model-B with 8 GB RAM) and a workstation laptop (Intel i7-6820HQ CPU, 2.70 GHz, 8 cores, 16 GB RAM). By setting $\beta = 0.8$, Turtlebot was able to plan and successfully travel a safe trajectory to the goal position avoiding the obstacles with reasonable safety margins, whereas CDDP results in the trajectories that pass very close to obstacles and sometimes collide with one of the obstacles (Fig. 5).

Our proposed approach does not require any additional QPs to be solved compared to CDDP. Only the added computation is in the form of uncertainty propagation and constraint tightening, which corresponds to less than %2 of CDDP execution. The worst-case execution times of a single iteration including backward and forward pass with the trajectory length $N = 90$ are approximately 63 ms and 22 ms for Raspberry Pi and workstation, respectively. Then the execution times decrease dramatically during traveling due to a decrease in active constraints and prediction horizon. This ensures the applicability of the proposed method in real robots with a control frequency faster than 10 hz.

V. CONCLUSION

We presented a novel safe trajectory optimization approach for nonlinear systems with nonlinear state and input constraints under additive system uncertainties. The proposed method reduces over conservatism in the approximation of uncertainty propagation through prediction. It also enables to impose safety precautions on general nonlinear constraints. Computational feasibility and applicability of the proposed approach were shown in hardware implementation, and the effectiveness of the method was validated on three different robot dynamics in simulation.

The proposed approach approximates the future prediction uncertainty by propagating it considering the implicit gain of the DDP feedback controller. It is likely that potential approximation errors in this prediction can be compensated conservatively by increasing the safety factor. However, an exact solution would be useful to limit the need for this, but whether it is possible to find such an exact solution remains an interesting open issue.

Recently, model-based reinforcement learning (RL) has received lots of attention, and integration of optimal control with safety guarantees to systems with learned dynamics seems to offer great possibilities. The proposed approach would be applicable to data-driven models of system dynamics models when their prediction uncertainty is modeled as additive Gaussian noise, which is the case for example for Gaussian Process models. Thus, we believe that the proposed method can be integrated as a part of a RL system that would be able to provide safety guarantees also for its exploration.

REFERENCES

- [1] T. A. N. Heirung, J. A. Paulson, J. O'Leary, and A. Mesbah, "Stochastic model predictive control-how does it work?" *Comput. Chem. Eng.*, vol. 114, pp. 158–170, 2018.
- [2] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious model predictive control using gaussian process regression," *IEEE Trans. Control Syst. Technol.*, vol. 28 no. 6, pp. 2736–2743, Nov. 2020.
- [3] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, "Robust constrained learning-based NMPC enabling reliable mobile robot path tracking," *Int. J. Robot. Res.*, vol. 35 no. 13, pp. 1547–1563, 2016.
- [4] M. Lorenzen, F. Dabbene, R. Tempo, and F. Allgöwer, "Constraint-tightening and stability in stochastic model predictive control," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3165–3177, Jul. 2017.
- [5] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames, "Towards a framework for realizable safety critical control through active set invariance," in *Proc. IEEE Int. Conf. Cyber-Phys. Syst.*, 2018, pp. 98–106.
- [6] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier-value functions for safety-critical control," 2021, *arXiv:2104.02808*.
- [7] B. T. Lopez, J. J. E. Slotine, and J. P. How, "Dynamic tube MPC for nonlinear systems," in *Proc. Amer. Control Conf.*, 2019, pp. 1655–1662.
- [8] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *Int. J. Robot. Res.*, vol. 36, no. 8, pp. 947–982, 2017.
- [9] G. Lantoine and R. P. Russell, "A hybrid differential dynamic programming algorithm for constrained optimal control problems part 1: Theory," *J. Optim. Theory Appl.*, vol. 154, no. 2, pp. 382–417, 2012.
- [10] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 1168–1175.
- [11] J. M. S. J. Sola and C. M. A. Santamaria-Navarro, "Squash-box feasibility driven differential dynamic programming," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2020, pp. 7637–7644.
- [12] N. Doshi, F. R. Hogan, and A. Rodriguez, "Hybrid differential dynamic programming for planar manipulation primitives," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 6759–6765.
- [13] Z. Xie, C. K. Liu, and K. Hauser, "Differential dynamic programming with nonlinear constraints," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 695–702.
- [14] Y. Aoyama, G. Boutselis, A. Patel, and E. A. Theodorou, "Constrained differential dynamic programming revisited," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 9738–9744.
- [15] A. Pavlov, I. Shames, and C. Manzie, "Interior point differential dynamic programming," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 6, pp. 2720–2727, Nov. 2021.
- [16] T. A. Howell, B. E. Jackson, and Z. Manchester, "ALTRO: A fast solver for constrained trajectory optimization," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2019, pp. 7674–7679.
- [17] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proc. Amer. Control Conf.*, 2005, pp. 300–306.
- [18] E. Theodorou, Y. Tassa, and E. Todorov, "Stochastic differential dynamic programming," in *Proc. Amer. Control Conf.*, 2010, pp. 1125–1132.
- [19] Y. Pan and E. Theodorou, "Probabilistic differential dynamic programming," in *Proc. IEEE Adv. Neural Inf. Process. Syst.*, 2014, pp. 1907–1915.
- [20] O. Celik, H. Abdulsamad, and J. Peters, "Chance-constrained trajectory optimization for non-linear systems with unknown stochastic dynamics," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2019, pp. 6828–6833.
- [21] N. Ozaki, S. Campagnola, R. Funase, and C. H. Yam, "Stochastic differential dynamic programming with unscented transform for low-thrust trajectory design," *J. Guid., Control, Dyn.*, vol. 41, no. 2, pp. 377–387, 2018.
- [22] N. Ozaki, S. Campagnola, and R. Funase, "Tube stochastic optimal control for nonlinear constrained trajectory optimization problems," *J. Guid., Control, Dyn.*, vol. 43, no. 4, pp. 645–655, 2020.
- [23] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
- [24] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*, Springer, pp. 207–226, 1999.
- [25] J. Nocedal and S. Wright, *Numerical Optimization*. Springer Science & Business Media, 2nd ed. 2006.
- [26] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2012, pp. 4906–4913.
- [27] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *Math. Program. Comput.*, vol. 12, no. 4, pp. 637–672, 2020.
- [28] G. Guennebaud and B. Jacob, "Eigen," 2010. [Online]. Available: <http://eigen.tuxfamily.org>
- [29] G. Alcan and M. Unel, "Robust hovering control of a quadrotor using acceleration feedback," in *Proc. IEEE Int. Conf. Unmanned Aircr. Syst.*, 2017, pp. 1455–1462.
- [30] H. Zaki, G. Alcan, and M. Unel, "Robust trajectory control of an unmanned aerial vehicle using acceleration feedback," *Int. J. Mechatronics Manuf. Syst.*, vol. 12, no. 3–4, pp. 298–317, 2019.