\mathcal{L}_1 -Adaptive MPPI Architecture for Robust and Agile Control of Multirotors

Jintasit Pravitra^{1*}, Kasey A. Ackerman^{2*}, Chengyu Cao³, Naira Hovakimyan², and Evangelos A. Theodorou¹

Abstract—This paper presents a multirotor control architecture, where Model Predictive Path Integral Control (MPPI) and \mathcal{L}_1 adaptive control are combined to achieve both fast model predictive trajectory planning and robust trajectory tracking. MPPI provides a framework to solve nonlinear MPC with complex cost functions in real-time. However, it often lacks robustness, especially when the simulated dynamics are different from the true dynamics. We show that the \mathcal{L}_1 adaptive controller robustifies the architecture, allowing the overall system to behave similar to the nominal system simulated with MPPI. The architecture is validated in a simulated multirotor racing environment.

I. INTRODUCTION

As multirotor aircraft continue to be integrated into our daily lives, researchers are challenged by the demand to automate complex tasks such as urban air transport, package delivery, autonomous racing, indoor exploration, or landing on a moving platform. These tasks often involve agile maneuvers and require complex trajectory planning. Despite recent advances, numerous challenges related to the problems of online optimal trajectory planning and replanning remain to be addressed.

Model predictive path integral control (MPPI) [1]–[3] offers a framework to efficiently solve a finite horizon nonlinear optimal control problem without restrictions on the form of the state cost function. The method is sampling based, and it leverages recent advances in GPU programming. Thousands of trajectories are propagated forward in parallel, and the optimal control is obtained by weight-averaging the controls from these trajectories (with each weight corresponding to the respective trajectory cost).

However, much like any simulation-based optimal control algorithm, MPPI suffers a degradation of robustness when the simulated dynamics are different from the true dynamics. The true trajectory may diverge from the planned trajectory because the control sequence optimized for the nominal dynamics is applied to the off-nominal true dynamics. The new initial state (for the purpose of MPC) can be significantly

¹Jintasit Pravitra and Evangelos A. Theodorou School of Aerospace Engineering, Georgia Institute USA Technology, Atlanta, GA 30332, jpravitra, evangelos.theodorou@gatech.edu

²Kasey A. Ackerman and Naira Hovakimyan are with the Mechanical Science and Engineering Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801 kaacker2, nhovakim@illinois.edu

³Chengyu Cao is with the Department of Mechanical Engineering, University of Connecticut, Storrs, CT 06269 chengyu.cao@uconn.edu *equal contribution

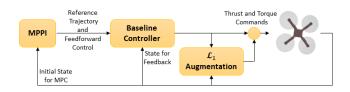


Fig. 1. Schematic of overall system architecture.

different from the previously planned state, thus rendering the warm-started controls to be far from optimal or in the worst case detrimental. A typical strategy to combat this issue is to include a tracking controller on top of the trajectory generator. For example in [4], iterative linear quadratic Gaussian (iLQG) is applied on top of MPPI to steer the states back to the planned trajectory. iLQG, however, requires a smooth cost function and also requires successive linearization of dynamics and cost function. Here we take a different approach and improve robustness through augmentation of an adaptive control element.

Among the various adaptive control methods, \mathcal{L}_1 adaptive control [5], [6] has been widely adopted due to its attractive properties of fast adaption, guaranteed robustness, and predictable transient response. These properties have been verified — consistently with the theory — in multiple manned and unmanned flight tests and simulations, on a diverse array of aircraft, including fixed-wing aircraft [7]-[13], multirotors [14]–[16], helicopters [17], and air-defense missiles [18]. The \mathcal{L}_1 adaptive control law can be formulated either as a standalone controller or as an augmentation of a baseline controller. In the augmentation architecture, the adaptive component can be conveniently disabled without affecting the nominal closed-loop dynamics, facilitating a distinctive comparison between the system behavior with and without the adaptive control component. Since adaptive control is a purely reactive control methodology, there is an opportunity to incorporate the advantages of baseline control algorithms that are capable of planning, such as MPC algorithms.

Previously, in [19], an \mathcal{L}_1 controller was used as an augmentation to a linear MPC, with a SISO linear reference model describing the desired behavior of the vehicle position response. While simpler to implement, a reference model which neglects the underlying vehicle dynamics may result in unrealistic desired behavior. Moreover, for certain systems, it may be advantageous to implement nonlinear reference

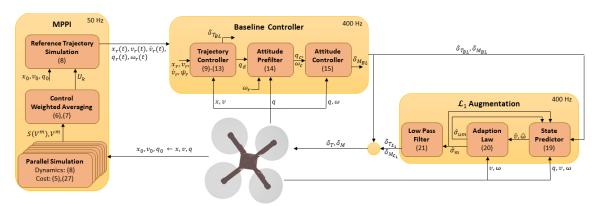


Fig. 2. Detailed control system architecture.

dynamics. By taking the nominal closed-loop dynamics (baseline controller and plant) as the reference dynamics, we explicitly account for the nonlinearity of the vehicle response and compensate the uncertainties directly. To the best of the authors' knowledge, there has never been an architecture where nonlinear MPC with complex cost function is robustified using \mathcal{L}_1 adaptive augmentation, which is the contribution of this paper.

II. ARCHITECTURE

The architecture — depicted in Fig. 1 — consists of three components. The first component is MPPI, which acts as a nonlinear MPC controller generating the optimized trajectory and the feedforward control sequence. Due to the recent advances in GPU programming, MPPI can replan at a rate as fast as 50 Hz. The second component is the baseline controller, which provides some robustification by recomputing the attitude command to correct for position and velocity deviation. The attitude command modification is obtained through geometric control for multirotors, and attitude tracking is achieved through a quaternion-based PD controller. The third component is the \mathcal{L}_1 adaptive augmentation. By defining the state predictor to be the same as the MPPI nominal dynamics, the \mathcal{L}_1 augmentation compensates for the mismatch between the nominal dynamics and the true dynamics. Hence, the tuning effort on the actual hardware can be reduced. A detailed architecture with all subcomponents and signal flows can be found in Fig. 2.

III. MULTIROTOR DYNAMICS

In this section, we describe the multirotor dynamics and quantify its (lumped) uncertainties. As depicted in Fig. 3, let the inertial basis vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ be aligned with the North-West-Up directions. Define the rotation matrix from the body frame to the inertial frame to be $\mathbf{R}_B^I \triangleq \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 & \mathbf{b}_3 \end{bmatrix}$. Define the attitude kinematics matrix to be

$$\mathbf{\Omega}(\boldsymbol{\omega}) \triangleq \frac{1}{2} \begin{bmatrix} 0 & -\omega_1 & -\omega_2 & -\omega_3 \\ \omega_1 & 0 & \omega_3 & -\omega_2 \\ \omega_2 & -\omega_3 & 0 & \omega_1 \\ \omega_3 & \omega_2 & -\omega_1 & 0 \end{bmatrix}. \tag{1}$$

The dynamics, with state vector $\mathbf{X} = \{\mathbf{x}, \mathbf{v}, \mathbf{q}, \boldsymbol{\omega}\}$

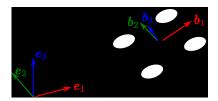


Fig. 3. Inertial and vehicle coordinate systems.

and control vector $\boldsymbol{\delta} = \{\delta_T \quad \boldsymbol{\delta}_M\}$, are

$$\dot{\mathbf{x}} = \mathbf{v}$$
 $\bar{m}\dot{\mathbf{v}} = \mathbf{b}_{3}\bar{T}_{\delta_{T}}\delta_{T} - \bar{m}g\mathbf{e}_{3} + \bar{\zeta}$ (2
 $\dot{\mathbf{q}} = \mathbf{\Omega}(\boldsymbol{\omega})\mathbf{q}$
 $\bar{\mathbf{J}}\dot{\boldsymbol{\omega}} = \bar{\mathbf{M}}_{\delta_{M}}\delta_{M} - \boldsymbol{\omega} \times \bar{\mathbf{J}}\boldsymbol{\omega} + \bar{\boldsymbol{\xi}},$

where $\bar{\zeta}$ and $\bar{\xi}$ are unknown forces and moments which encapsulate exogenous disturbance forces, model uncertainty, and unmodeled dynamics. Such uncertainties may include — but are not limited to — unknown aerodynamics, misalignment of the mass center, and parametric uncertainty. The parameters $\bar{m}, \bar{\mathbf{J}}, \bar{T}_{\delta_T}, \bar{M}_{\delta_M}$ correspond to mass, moment of inertia, thrust control derivative, and moment control derivative matrix respectively. These quantities are unknown in practice and only their nominal values m, \mathbf{J} , T_{δ_T} , \mathbf{M}_{δ_M} are available. The multirotor dynamics can be rewritten with the nominal parameters and the lumped uncertainties (ζ and ξ) as

$$\dot{\mathbf{x}} = \mathbf{v} \qquad \dot{\mathbf{v}} = \frac{T_{\delta_T} \delta_T}{m} \mathbf{b}_3 - g \mathbf{e}_3 + \frac{T_{\delta_T}}{m} \mathbf{R}_B^I \boldsymbol{\zeta} \quad (3)$$

$$\dot{\mathbf{q}} = \boldsymbol{\Omega}(\boldsymbol{\omega}) \mathbf{q} \qquad \dot{\boldsymbol{\omega}} = \mathbf{J}^{-1} \mathbf{M}_{\delta_M} \boldsymbol{\delta}_M + \mathbf{J}^{-1} \mathbf{M}_{\delta_M} \boldsymbol{\xi}$$

with
$$\zeta = \frac{m}{\bar{m}T_{\delta_T}}\mathbf{R}_I^B\bar{\zeta} + \left(\frac{m\bar{T}_{\delta_T}}{\bar{m}T_{\delta_T}} - 1\right)\mathbf{R}_I^B\mathbf{b}_3\delta_T$$
 and $\boldsymbol{\xi} = \mathbf{M}_{\delta_M}^{-1}\mathbf{J}\bar{\mathbf{J}}^{-1}\bar{\boldsymbol{\xi}} + \left(\mathbf{M}_{\boldsymbol{\delta}_M}^{-1}\mathbf{J}\bar{\mathbf{J}}^{-1}\mathbf{M}_{\boldsymbol{\delta}_M} - \mathbb{I}\right)\boldsymbol{\delta}_M - \mathbf{M}_{\boldsymbol{\delta}_M}^{-1}\mathbf{J}\bar{\mathbf{J}}^{-1}\left(\boldsymbol{\omega}\times\bar{\mathbf{J}}\boldsymbol{\omega}\right)$. A similar uncertainty parameterization can be found in [14].

IV. MPPI TRAJECTORY GENERATION

A. Review of MPPI

In this section, we provide a brief review of MPPI. Readers are encouraged to see [20] for the stochastic Hamilton-Jacobi-Bellman derivation, [1]–[3] for the information theoretic derivation, and [21], [22] for the stochastic optimization

derivation. In this paper, we use the information theoretic derivation.

Consider the discrete-time dynamics $\mathbf{s}_{t+1} = \mathbf{F}(\mathbf{s}_t, \mathbf{v}_t)$, where \mathbf{v}_t is a Gaussian distributed random control input $\mathbf{v}_t \sim \mathcal{N}(\mathbf{u}_t, \mathbf{\Sigma})$. We are interested in finding the mean control sequence $\mathbf{U} = \left\{\mathbf{u}_0, \mathbf{u}_1, \dots \mathbf{u}_{T-1}\right\}$ that minimizes the cost function

$$J(\mathbf{U}) = \mathbb{E}\left[\phi(\mathbf{s}_T) + \sum_{t=0}^{T-1} \left[Q(\mathbf{s}_t) + \lambda \mathbf{u}_t^\mathsf{T} \mathbf{\Sigma}^{-1} \mathbf{u}_t\right]\right]. \quad (4)$$

Note that the state running cost, $Q(\mathbf{s}_t)$, and the terminal cost, $\phi(\mathbf{s}_T)$, can be arbitrarily complex. For each control sequence realization $\mathbf{V}^m = \{\mathbf{v}_0^m, \mathbf{v}_1^m, \dots \mathbf{v}_{T-1}^m\}$, define $S(\mathbf{V}^m)$ to be the state-dependent portion of the cost

$$S(\mathbf{V}^m) = \phi(\mathbf{s}_T^m) + \sum_{t=0}^{T-1} Q(\mathbf{s}_t^m). \tag{5}$$

Start from the initial condition s_0 . Given the control sequence from the previous iteration U_{k-1} , simulate (thousands of) trajectories in parallel, each with a different control sequence realization V^m . The costs are collected for each rollout and are mapped to the trajectory weights:

$$w(\mathbf{V}^m) = \exp\left(-\frac{1}{\lambda}\left(S(\mathbf{V}^m) - \sum_{t=0}^{T-1} \mathbf{u}_{t,k-1}^\mathsf{T} \mathbf{\Sigma}^{-1} \mathbf{v}_t^m - \rho\right)\right)$$
(6)

The term ρ is included in order to prevent arithmetic underflow. The value of ρ is typically set to the minimum cost among all sampled trajectories. This term does not affect the solution because of the normalization (as will be seen in the next step). The optimal control at each time step can be approximated as

$$\mathbf{u}_{t,k} = \mathbf{u}_{t,k-1} + \frac{1}{\sum_{m=1}^{M} w(\mathbf{V}^m)} \sum_{m=1}^{M} \left[w(\mathbf{V}^m) \epsilon_t^m \right], \quad (7)$$

where $\epsilon_t^m = \mathbf{v}_t^m - \mathbf{u}_{t,k-1}$ and M is the number of samples.

B. Multirotor Application

To apply MPPI to a multirotor trajectory generation problem, we propagate the kinematic equations and the disturbance-free translational dynamics. The MPPI controller treats the angular velocity as a command to be tracked by the lower-level controller, where uncertainty and disturbances in the rotational motion are compensated by the adaptive control law. The sampled control inputs are thrust and angular velocity: $\mathbf{v}_t = [\delta_T \ u_{\omega_1} \ u_{\omega_2} \ u_{\omega_3}]^\mathsf{T}$.

However, injecting discontinuous angular velocity is not a realistic representation of the low-level controller. To represent this effect, the sampled angular velocity is low-pass filtered before it enters the attitude kinematics. In summary, we define the MPPI states to be $s = \{\mathbf{x}, \mathbf{v}, \mathbf{q}, \boldsymbol{\omega}\}$ and propagate the trajectories using the following equations:

$$\dot{\mathbf{x}} = \mathbf{v} \qquad \dot{\mathbf{v}} = \frac{T_{\delta_T} \delta_T}{m} \mathbf{b}_3 - g \mathbf{e}_3 \qquad (8)$$

$$\dot{\mathbf{q}} = \mathbf{\Omega} \Big([\omega_1 \ \omega_2 \ \omega_3]^\mathsf{T} \Big) \mathbf{q} \qquad \dot{\omega}_i = \frac{u_{\omega_i} - \omega_i}{\tau_{\omega_i}}, \ i = 1, 2, 3$$

After the trajectories are propagated and the optimal control sequence is obtained through (6) and (7), the optimal control sequence is simulated through (8) to obtain the reference trajectory $\{\mathbf{x}_r(t), \mathbf{v}_r(t), \mathbf{q}_r(t), \boldsymbol{\omega}_r(t)\}$. Additionally, the reference acceleration sequence $\dot{\mathbf{v}}_r(t)$ is also saved to be used in the baseline controller.

V. BASELINE CONTROLLER

We provide the procedure to obtain the baseline control input $\delta_{T_{BL}}$ and $\delta_{M_{BL}}$. The procedure follows the geometric control methods as in [23], [24], and [25].

A. Baseline Trajectory Tracking Controller

The desired specific force includes the feedforward term $\dot{\mathbf{v}}_r$ from MPPI and feedback terms to correct for position and velocity deviation:

$$\mathbf{f}_d = \sup_{\mathbf{a}_{max}} \left[\dot{\mathbf{v}}_r + \mathbf{K}_P(\mathbf{x}_r - \mathbf{x}) + \mathbf{K}_D(\mathbf{v}_r - \mathbf{v}) \right] + g\mathbf{e}_3, \quad (9)$$

where $\operatorname{sat}[\cdot]$ is a function that clamps the acceleration to be between $[-\mathbf{a}_{max}, \mathbf{a}_{max}]$. The baseline throttle command is proportional to the norm of the specific force:

$$\delta_{T_{BL}} = (m/T_{\delta_T})||\mathbf{f}_d||. \tag{10}$$

Define the desired rotation matrix to be $\mathbf{R}_D^I \triangleq [\mathbf{d}_1 \ \mathbf{d}_2 \ \mathbf{d}_3]$. For a vehicle where all rotors are aligned in a single plane, \mathbf{f}_d must be aligned with \mathbf{d}_3 :

$$\mathbf{d}_3 = \mathbf{f}_d / ||\mathbf{f}_d||. \tag{11}$$

We obtain the reference heading ψ_r from \mathbf{q}_r and constrain the local level projection of the vehicle's nose to be

$$\mathbf{l}_1 = [\cos \psi_r \quad \sin \psi_r \quad 0]^\mathsf{T}. \tag{12}$$

The other two components of the \mathbf{R}_D^I can be obtained using the following relationships:

$$\mathbf{d}_2 = (\mathbf{d}_3 \times \mathbf{l}_1) / ||\mathbf{d}_3 \times \mathbf{l}_1|| \qquad \mathbf{d}_1 = \mathbf{d}_2 \times \mathbf{d}_3. \tag{13}$$

At this point, the desired attitude \mathbf{R}_D^I is completely defined.

B. Baseline Attitude Controller

Define \mathbf{q}_d to be the quaternion representation of \mathbf{R}_D^I , and define \mathbf{R}_R^I to be the rotation matrix representation of \mathbf{q}_r . Since MPPI and the baseline controller run at different rates, the calculation (9) - (13) may result in a discontinuous attitude command. The following prefilter is recommended in order to smooth the command

$$\dot{\mathbf{q}}_c = \mathbf{\Omega}(\omega_c)\mathbf{q}_c \tag{14}$$

$$\dot{\boldsymbol{\omega}}_{c} = \mathbf{K}_{\omega f} \big[\underset{\boldsymbol{\omega}_{max}}{\operatorname{sat}} [\mathbf{K}_{\omega f}^{-1} \mathbf{K}_{qf} \widetilde{\mathbf{Q}} (\mathbf{q}_{d}, \mathbf{q}_{c})] + (\mathbf{R}_{I}^{B} \mathbf{R}_{R}^{I} \boldsymbol{\omega}_{r} - \boldsymbol{\omega}_{c}) \big],$$

where $\widetilde{\mathbf{Q}}(\cdot,\cdot)$ is the quaternion error function [26]. The filtered quaternion \mathbf{q}_c and angular velocity $\boldsymbol{\omega}_c$ become the commands of the baseline PD attitude controller:

$$\delta_{M_{BL}} = \mathbf{K}_{\omega} \left[\underset{\boldsymbol{\omega}_{max}}{\text{sat}} \left[\mathbf{K}_{\omega}^{-1} \mathbf{K}_{q} \widetilde{\mathbf{Q}}(\mathbf{q}_{c}, \mathbf{q}) \right] + (\boldsymbol{\omega}_{c} - \boldsymbol{\omega}) \right]$$
 (15)

Since the sole purpose of the prefilter is to avoid feeding discontinuous attitude commands to the baseline controller, the prefilter gains \mathbf{K}_{qf} and $\mathbf{K}_{\omega f}$ should be greater than the baseline gains \mathbf{K}_q and \mathbf{K}_{ω} .

VI. \mathcal{L}_1 Adaptive Augmentation

The \mathcal{L}_1 augmentation implements a nonlinear reference model [27], with estimation of both matched and unmatched uncertainties [28]. The piecewise constant adaptive law formulation has been chosen over the projection operator formulation for its numerical robustness. The unmatched uncertainty arises in the translational dynamics, since a multirotor (with rotors aligned in a single plane) can only provide linear acceleration along the body z-axis. Because the uncertainties ζ , ξ , appear purely in the vehicle dynamics, the kinematic part of (3) can be omitted and only the dynamics are considered:

$$\dot{\mathbf{v}} = -g\mathbf{e}_{3} + \frac{T_{\delta_{T}}}{m}\mathbf{b}_{3}(\delta_{T_{BL}} + \delta_{T_{\mathcal{L}_{1}}} + \zeta_{m}) + \frac{T_{\delta_{T}}}{m}\begin{bmatrix}\mathbf{b}_{1} & \mathbf{b}_{2}\end{bmatrix}\boldsymbol{\zeta}_{um}$$

$$\dot{\boldsymbol{\omega}} = \mathbf{J}^{-1}\mathbf{M}_{\delta_{\mathbf{M}}}(\boldsymbol{\delta}_{\mathbf{M}_{\mathbf{BL}}} + \boldsymbol{\delta}_{\mathbf{M}_{\mathcal{L}_{1}}} + \boldsymbol{\xi}_{\mathbf{m}}),$$
(16)

where substitutions have been made for

$$oldsymbol{\zeta} = egin{bmatrix} oldsymbol{\zeta}_{um} \ oldsymbol{\zeta}_{m} \end{bmatrix}, \quad oldsymbol{\delta}_{T} = oldsymbol{\delta}_{T_{BL}} + oldsymbol{\delta}_{T_{\mathcal{L}_{1}}}, \quad oldsymbol{\delta}_{M} = oldsymbol{\delta}_{M_{BL}} + oldsymbol{\delta}_{M_{\mathcal{L}_{1}}}.$$

In a more general form, the dynamics can be written as

$$\dot{z} = \mathbf{f} \left(\mathbf{R}_{B}^{I} \right) + \mathbf{g} \left(\mathbf{R}_{B}^{I} \right) \left(\mathbf{u}_{\mathcal{L}_{1}} + \boldsymbol{\sigma}_{m} \right) + \mathbf{g}^{\perp} \left(\mathbf{R}_{B}^{I} \right) \boldsymbol{\sigma}_{um}, (18)$$

where

$$oldsymbol{z} = egin{bmatrix} \mathbf{v}^\mathsf{T} & oldsymbol{\omega}_{\mathcal{L}_1} = egin{bmatrix} \delta_{T_{\mathcal{L}_1}} & \delta_{\mathbf{M}_{\mathcal{L}_1}}^\mathsf{T} \end{bmatrix}^\mathsf{T} \ oldsymbol{\sigma}_m = egin{bmatrix} \zeta_m & oldsymbol{\xi}^\mathsf{T} \end{bmatrix}^\mathsf{T} & oldsymbol{\sigma}_{um} = oldsymbol{\zeta}_{um} \end{cases}$$

$$\begin{split} \mathbf{f}(\mathbf{R}_B^I) &= \begin{bmatrix} -g\mathbf{e}_3 + m^{-1}\mathbf{b}_3T_{\delta_T}\delta_{T_{BL}} \\ \mathbf{J}^{-1}\mathbf{M}_{\delta_\mathbf{M}}\delta_{\mathbf{M}_{\mathbf{BL}}} \end{bmatrix} \\ \mathbf{g}(\mathbf{R}_B^I) &= \begin{bmatrix} m^{-1}\mathbf{b}_3T_{\delta_T} & \mathbf{0}_{3\times 3} \\ \mathbf{0}_{3\times 1} & \mathbf{J}^{-1}\mathbf{M}_{\delta_\mathbf{M}} \end{bmatrix} \\ \mathbf{g}^{\perp}(\mathbf{R}_B^I) &= \begin{bmatrix} m^{-1}\mathbf{b}_1T_{\delta_T} & m^{-1}\mathbf{b}_2T_{\delta_T} \\ \mathbf{0}_{3\times 1} & \mathbf{0}_{3\times 1} \end{bmatrix} \end{split}$$

Then the \mathcal{L}_1 state predictor is defined as

$$\dot{\hat{\mathbf{z}}} = \mathbf{f} + \mathbf{g} \left(\mathbf{u}_{f_{*}} + \hat{\boldsymbol{\sigma}}_{m} \right) + \mathbf{g}^{\perp} \hat{\boldsymbol{\sigma}}_{um} + \mathbf{A}_{s} \tilde{\mathbf{z}}, \tag{19}$$

where $\tilde{z} = \hat{z} - z$ and \mathbf{A}_s is a Hurwitz matrix. For simplicity, we assume $\hat{z}(0) = z(0)$ and note that a non-zero state predictor initialization error results in an exponentially decaying term [29]. Define $\mathbf{\Phi} \triangleq \mathbf{A}_s^{-1} \left(\exp(\mathbf{A}_s T_s) - \mathbb{I} \right)$. The piecewise-constant adaption law is

$$\begin{bmatrix} \hat{\boldsymbol{\sigma}}_{m}(iT_{s}) \\ \hat{\boldsymbol{\sigma}}_{um}(iT_{s}) \end{bmatrix} = - \begin{bmatrix} \mathbb{I}_{4\times4} & \mathbf{0}_{4\times2} \\ \mathbf{0}_{2\times4} & \mathbb{I}_{2\times2} \end{bmatrix} \mathbf{G}(iT_{s})^{-1} \mathbf{\Phi}^{-1} \boldsymbol{\mu}(iT_{s})$$
(20)

where T_s is the time step and where $\mathbf{G}(iT_s) = [\mathbf{g}(\mathbf{R}_B^I) \ \mathbf{g}^{\perp}(\mathbf{R}_B^I)]$ and $\boldsymbol{\mu}(iT_s) = \exp(\mathbf{A}_s T_s) \tilde{\boldsymbol{z}}(iT_s)$ are

evaluated at the i-th time step. The \mathcal{L}_1 control law compensates only for the matched components of the uncertainty within the bandwidth of the strictly proper stable filter $\mathbf{C}(s)$:

$$\mathbf{u}_{\mathcal{L}_1} = -\mathbf{C}(s)\hat{\boldsymbol{\sigma}}_m,\tag{21}$$

which serves to decouple the fast adaptation from the system robustness [5]. The effects of the unmatched component are estimated to improve the performance of the state predictor but, for this application, it is not necessary to compensate for the unmatched components directly; they can be indirectly canceled by the baseline control law. Proofs regarding stability and bounds on states and controls can be found in [27].

Remark 1. Define the (invertible) elementary column operator

$$\mathbf{E} = \begin{bmatrix} \mathbf{0}_{4 \times 2} & \mathbb{I}_{4 \times 4} \\ \mathbb{I}_{2 \times 2} & \mathbf{0}_{2 \times 4} \end{bmatrix}$$
 (22)

and let $\mathbf{H} = \mathbf{GE}$. Then

$$\mathbf{H} = \begin{bmatrix} m^{-1} \mathbf{R}_B^I T_{\delta_T} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{J}^{-1} \mathbf{M}_{\delta_{\mathbf{M}}} \end{bmatrix}, \tag{23}$$

and we can directly obtain \mathbf{G}^{-1} as

$$\mathbf{G}^{-1} = \mathbf{E}\mathbf{H}^{-1} = \begin{bmatrix} m\mathbf{b}_{3}^{\mathsf{T}}T_{\delta_{T}}^{-1} & \mathbf{0}_{1\times3} \\ \mathbf{0}_{3\times3} & \mathbf{M}_{\delta_{M}}^{-1}\mathbf{J} \\ m\mathbf{b}_{1}^{\mathsf{T}}T_{\delta_{T}}^{-1} & \mathbf{0}_{1\times3} \\ m\mathbf{b}_{2}^{\mathsf{T}}T_{\delta_{T}}^{-1} & \mathbf{0}_{1\times3} \end{bmatrix}$$
(24)

Since T_{δ_T} and $\mathbf{M}_{\delta_{\mathbf{M}}}$ are known and the rotation matrix \mathbf{R}_B^I is obtained from sensor/estimator feedback, the nullspace component \mathbf{g}^{\perp} can be directly obtained as well. Therefore, no matrix inversion operations are needed to compute either $\mathbf{g}^{\perp}(t)$ or $\mathbf{G}^{-1}(t)$.

Since the \mathcal{L}_1 adaptive augmentation is intended to be implemented in a discrete-time environment, we formulate the discrete version of the adaptation law as

$$\begin{bmatrix} \hat{\boldsymbol{\sigma}}_{m,k} \\ \hat{\boldsymbol{\sigma}}_{um,k} \end{bmatrix} = - \begin{bmatrix} \mathbb{I}_{4\times4} & \mathbf{0}_{4\times2} \\ \mathbf{0}_{2\times4} & \mathbb{I}_{2\times2} \end{bmatrix} \mathbf{G}_k^{-1} \mathbf{\Phi}^{-1} \boldsymbol{\mu}_k.$$
 (25)

Here, $(\cdot)_k$ is equivalent to $(\cdot)(kT_s)$. The discrete version of the control law can be written as

$$\mathbf{x}_{f,k+1} = \mathbf{A}_f \mathbf{x}_{f,k} + \mathbf{B}_f \hat{\boldsymbol{\sigma}}_{m,k} \qquad \mathbf{u}_{\mathcal{L}_1,k} = -\mathbf{C}_f \mathbf{x}_{f,k},$$

where $\{\mathbf{A}_f, \mathbf{B}_f, \mathbf{C}_f\}$ are the discrete-time state space matrices defining the low-pass filter $\mathbf{C}(s)$, and \mathbf{x}_f is the filter state. The discrete-time state predictor is propagated via

$$\hat{\boldsymbol{z}}_{k+1} = \hat{\boldsymbol{z}}_k + \left[\mathbf{f}_k + \mathbf{g}_k \left(\mathbf{u}_{\mathcal{L}_1,k} + \hat{\boldsymbol{\sigma}}_{m,k} \right) + \mathbf{g}_k^{\perp} \hat{\boldsymbol{\sigma}}_{um,k} + \mathbf{A}_s \tilde{\boldsymbol{z}}_k \right] T_s,$$
(26)

where $\tilde{\boldsymbol{z}}_k = \hat{\boldsymbol{z}}_k - \boldsymbol{z}_k$.

VII. RESULTS

The control architecture is evaluated in FlightGoggles, a photorealistic quadrotor racing simulation environment based on Unity and ROS [30]. The objective is to fly through the designated gates in the correct sequence and finish the course as fast as possible. We assume that the gate locations are known a priori and that the state estimation is perfect. A screenshot of FlightGoggles is shown in Fig. 4.

¹For a system $\dot{z} = \mathbf{f}(z,t) + \mathbf{g}(z,t)\mathbf{u} + \boldsymbol{\sigma}(t)$, the uncertainty, $\boldsymbol{\sigma}(t)$, can be decomposed into matched and unmatched components, respectively $\boldsymbol{\sigma}_m(t) \in \operatorname{span}(\mathbf{g}(z,t))$ and $\boldsymbol{\sigma}_{um}(t) \in \operatorname{span}(\operatorname{Null}(\mathbf{g}(z,t)))$.



Fig. 4. FlightGoggles multirotor racing simulation environment.

A. Cost Function and Parameters

The MPPI cost function is designed to aggregate the following objectives:

- maintain the (3-dimensional) air corridor's centerline;
- point the front of the multirotor toward the next gate;
- heavily penalize trajectories that violate the air corridor;
- maintain the commanded speed;
- give a bonus to trajectories that pass through gates.

The expression of the cost function is

$$Q(\mathbf{s}) = 450M(\mathbf{x}) + 250|\text{mod}(\psi_{cmd} - \psi)| + 10000\mathbb{1}_{out}(\mathbf{x}) + 150|V_{cmd} - \sqrt{v_x^2 + v_y^2}| - 150\sum \mathbb{1}_{gate}(\mathbf{x}),$$
(27)

where $M(\mathbf{x})$ is the air corridor cost map look-up as shown in Fig. 5. We select the planning horizon to be 1.5 s. Since the units of δ_T and δ_M are arbitrary, we can conveniently select $T_{\delta_T}=1.0$ and $\mathbf{M}_{\delta_M}=\mathbb{I}_{\mathbf{3}\times\mathbf{3}}$ so that δ_T is in the unit of Newtons and δ_M is in the unit of Newton-meters. This can be done without loss of generality. The derivation of the lumped uncertainties is still valid with the scaled values of \bar{T}_{δ_T} and $\bar{\mathbf{M}}_{\delta_M}$. The rest of the parameters are shown in Table I.

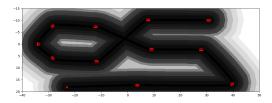


Fig. 5. Top view of the air corridor cost map.

B. Results

In the presence of disturbance, \mathcal{L}_1 -MPPI is expected to outperform the standalone MPPI. We verify this hypothesis

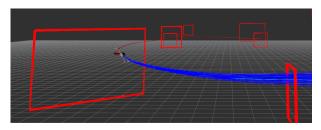


Fig. 6. Ground station visualization of simulation.

TABLE I SIMULATION PARAMETERS

Nominal Plant				
m	1.0	kg		
J	$diag(4.9, 4.9, 4.9) \times 10^{-3}$	$kg m^2$		
MPPI				

MPPI				
Σ	diag(1.5 , 0.4 , 0.4 , 0.4)	N, rps		
M	7,200			
λ	1.4			
$\tau_{\omega_1}, \ \tau_{\omega_2}, \ \tau_{\omega_3}$	0.25	S		

Baseline Controller		
\mathbf{K}_{P}	diag(6.0, 6.0, 6.0)	
\mathbf{K}_D	diag(4.0, 4.0, 4.0)	
\mathbf{K}_q	diag(1.0, 1.0, 1.0)	
\mathbf{K}_{ω}	diag(0.15, 0.15, 0.15)	
\mathbf{a}_{max}	$15.0 \ m/s^2$	
ω_{max}	$2.0 \ rps$	

\mathcal{L}_{1} Controller				
T_s	2.5	ms		
C(s)	$\frac{15}{(s+15)}\mathbb{I}_{4\times4}$			
\mathbf{A}_{s}	$-5.0\mathbb{I}_{6\times6}$			

by modifying the vehicle dynamics simulation in the following ways:

- 1) no modification other than the inclusion of aerodynamic drag (since drag is not modeled in (8), some drag compensation is expected in \mathcal{L}_1 -MPPI drag is also included in all other test cases);
- 2) mass increase by 50%;
- 3) moment of inertia increase by 100% in all axes;
- 4) constant nose-up pitching moment disturbance of $0.1\ Nm$ (equivalent to center of gravity offset);
- 5) reduction in motor thrust by 40% (reduction in both \bar{T}_{δ_T} and $\bar{\mathbf{M}}_{\delta_M}$).

Figure 6 shows a visualization for one of the runs. For each case, we run the race 15 times. The average lap time is shown in Fig 7. The bar graph is shown only if all 15 runs are successfully completed. (The vehicle neither crashes nor diverges.) The fact that not all cases are successful is also shown in Table II.

Figure 7 and Table II show that \mathcal{L}_1 augmentation successfully reduces lap time. For the cases of extreme disturbances, \mathcal{L}_1 augmentation successfully "saves" the vehicle that would have crashed if only the baseline controller was employed.



Fig. 7. Average lap time for each simulation case.

VIII. CONCLUSION

In this paper, we proposed a multirotor control architecture where MPPI acts as a nonlinear MPC controller, while model uncertainty was compensated through an \mathcal{L}_1 adaptive controller. It was shown in the quadrotor racing scenario that inclusion of \mathcal{L}_1 augmentation improves performance and, in

 $\label{eq:TABLE II} \textbf{Success and failure for each case}$

Case	\mathcal{L}_1 off	\mathcal{L}_1 on
1)	✓	✓
2)	X	✓
3)	✓	✓
4)	X	✓
5)	X	✓

the presence of large uncertainties, was imperative to the success of the mission. Potential future work includes flight testing, detailed analysis of robustness, and incorporation of the control limits.

ACKNOWLEDGMENT

The authors would like to thank the members of the Georgia Tech AlphaPilot team who initiated the implementation of MPPI with FlightGoggles, including Grady Williams, Sam Yi Ting, Jason Gibson, Manan Gandhi, Keuntaek Lee, Bogdan Vlahov, Camilo Duarte, Akash Patel, and Matthew Houghton. The authors would also like to thank Aditya Gahlawat from UIUC for the fruitful discussion.

This work is supported by NASA and NSF CPS awards #1932529 and #1932288.

REFERENCES

- [1] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [2] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1714–1721.
- [3] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [4] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. M. Rehg, and E. Theodorou, "Robust sampling based model predictive control with sparse objective information." in *Robotics: Science and Systems*, 2018.
- [5] N. Hovakimyan and C. Cao, L₁ Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation. SIAM, 2010.
- [6] N. Hovakimyan, C. Cao, E. Kharisov, E. Xargay, and I. M. Gregory, "L₁ adaptive control for safety-critical systems," *IEEE Control Systems Magazine*, vol. 31, no. 5, pp. 54–104, 2011.
- [7] K. A. Ackerman, E. Xargay, R. Choe, N. Hovakimyan, M. C. Cotting, R. B. Jeffrey, M. P. Blackstun, T. P. Fulkerson, T. R. Lau, and S. S. Stephens, "L₁ stability augmentation system for Calspan's variable-stability Learjet," in AIAA Guidance, Navigation, and Control Conference, 2016, AIAA-2016-0631.
- [8] M. C. Cotting, R. B. Jeffrey, M. P. Blackstun, T. P. Fulkerson, T. R. Lau, S. S. Stephens, K. A. Ackerman, E. Xargay, R. Choe, and N. Hovakimyan, "'Can I get \(\mathcal{L}_1 \) on?!' Providing consistent handling qualities on Calspan's variable-stability Learjet," in AIAA Guidance, Navigation, and Control Conference, 2016, AIAA-2016-0632.
- [9] J. Puig-Navarro, K. A. Ackerman, N. Hovakimyan, M. C. Cotting, D. J. Duke, M. J. Carrera, N. C. McCaskey, D. Esposito, J. M. Peterson, and J. R. Tellefsen, "An L₁ adaptive stability augmentation system designed for MIL-HDBK-1797 level 1 flying qualities," in AIAA SciTech Forum, 2019, AIAA-2019-1083.
- [10] K. A. Ackerman, J. Puig-Navarro, N. Hovakimyan, M. C. Cotting, D. J. Duke, M. J. Carrera, N. C. McCaskey, D. Esposito, J. M. Peterson, and J. R. Tellefsen, "Recovery of desired flying characteristics with an L₁ adaptive control law: Flight test results on Calspan's VSS Learjet," in AIAA SciTech Forum, 2019, AIAA-2019-1084.

- [11] E. Xargay, N. Hovakimyan, V. Dobrokhodov, I. Kaminer, C. Cao, and I. Gregory, "L₁ adaptive control in flight," *Advances in Intelligent and Autonomous Aerospace Systems*, vol. 241, pp. 129–172, 2012.
- [12] T. Leman, E. Xargay, G. Dullerud, N. Hovakimyan, and T. Wendel, "£1 adaptive control augmentation system for the X-48B aircraft," in AIAA Guidance, Navigation and Control Conference, 2009, AIAA-2009-5619.
- [13] R. Choe, O. Stroosma, E. Xargay, H. Damveld, N. Hovakimyan, J. Mulder, and H. Damveld, "A handling qualities assessment of a business jet augmented with an L₁ adaptive controller," in AIAA Guidance, Navigation and Control Conference, 2011, AIAA-2011-6610
- [14] S. Mallikarjunan, B. Nesbitt, E. Kharisov, E. Xargay, N. Hovakimyan, and C. Cao, "L₁ adaptive controller for attitude control of multirotors," in AIAA Guidance, Navigation, and Control Conference, 2012, AIAA-2012-4831.
- [15] S. Mallikarjunan, Z. Zuo, B. Nesbitt, and N. Hovakimyan, "L₁ back-stepping for robust trajectory tracking," *IFAC Proceedings Volumes*, vol. 46, no. 30, pp. 1–6, 2013.
- [16] J. Wang, F. Holzapfel, E. Xargay, and N. Hovakimyan, "Non-cascaded dynamic inversion design for quadrotor position control with L₁ augmentation," in *Proc. of the CEAS Specialist Conf. on Guidance,* Navigation & Control, Delft, Netherlands, 2013.
- [17] M. Bichlmeier, "Certifiable \mathcal{L}_1 adaptive control for helicopters," Ph.D. dissertation, Technische Universität München, 2016.
- [18] F. Peter, F. Holzapfel, E. Xargay, and N. Hovakimyan, " \mathcal{L}_1 adaptive augmentation of a missile autopilot," in AIAA Guidance, Navigation, and Control Conference, 2012, AIAA-2012-4832.
- [19] K. Pereida and A. P. Schoellig, "Adaptive model predictive control for high-accuracy trajectory tracking in changing conditions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7831–7837.
- [20] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal* of Guidance, Control, and Dynamics, vol. 40, no. 2, pp. 344–357, 2017
- [21] Z. Wang, G. Williams, and E. A. Theodorou, "Information theoretic model predictive control on jump diffusion processes," in 2019 American Control Conference (ACC). IEEE, 2019, pp. 1663–1670.
- [22] G. I. Boutselis, Z. Wang, and E. A. Theodorou, "Constrained sampling-based trajectory optimization using stochastic approximation," arXiv preprint arXiv:1911.04621, 2019.
- [23] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *IEEE International Conference on Robotics* and Automation, 2011, pp. 2520–2525.
- [24] M. Faessler, F. Fontana, C. Forster, and D. Scaramuzza, "Automatic re-initialization and failure recovery for aggressive flight with a monocular vision-based quadrotor," in *IEEE International Conference* on Robotics and Automation (ICRA), 2015, pp. 1722–1729.
- [25] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor UAV on SE(3)," in 49th IEEE conference on decision and control (CDC), 2010, pp. 5420–5425.
- [26] E. N. Johnson, "Limited authority adaptive flight control," Ph.D. dissertation, School of Aerospace Engineering, Georgia Institute of Technology Atlanta, GA, 2000.
- [27] X. Wang and N. Hovakimyan, "L₁ adaptive controller for nonlinear time-varying reference systems," Systems & Control Letters, vol. 61, no. 4, pp. 455–463, 2012.
- [28] E. Xargay, N. Hovakimyan, and C. Cao, "£1 adaptive controller for multi-input multi-output systems in the presence of nonlinear unmatched uncertainties," in *Proceedings of the 2010 American Control* Conference. IEEE, 2010, pp. 874–879.
- [29] C. Cao and N. Hovakimyan, "Effect of non-zero initialization error on the performance bounds in L₁ adaptive control architecture," in AIAA Guidance, Navigation and Control Conference and Exhibit, 2007, AIAA-2007-6645.
- [30] W. Guerra, E. Tal, V. Murali, G. Ryou, and S. Karaman, "Flightgoggles: Photorealistic sensor simulation for perception-driven robotics using photogrammetry and virtual reality," arXiv preprint arXiv:1905.11377, 2019.