

Feedback-MPPI: Fast Sampling-Based MPC via Rollout Differentiation – Adios low-level controllers

Tommaso Belvedere, Michael Ziegltrum, Giulio Turrisi, Valerio Modugno

Abstract—Model Predictive Path Integral control is a powerful sampling-based approach suitable for complex robotic tasks due to its flexibility in handling nonlinear dynamics and non-convex costs. However, its applicability in real-time, high-frequency robotic control scenarios is limited by computational demands. This paper introduces Feedback-MPPI (F-MPPI), a novel framework that augments standard MPPI by computing local linear feedback gains derived from sensitivity analysis inspired by Riccati-based feedback used in gradient-based MPC. These gains allow for rapid closed-loop corrections around the current state without requiring full re-optimization at each timestep. We demonstrate the effectiveness of F-MPPI through simulations and real-world experiments on two robotic platforms: a quadrupedal robot performing dynamic locomotion on uneven terrain and a quadrotor executing aggressive maneuvers with onboard computation. Results illustrate that incorporating local feedback significantly improves control performance and stability, enabling robust, high-frequency operation suitable for complex robotic systems.

Index Terms—Optimization and Optimal Control; Motion Control; Legged Robots; Model Predictive Control

I. INTRODUCTION

MODEL Predictive Control (MPC) was initially adopted in the process industry, where slow system dynamics of chemical plants and refineries made it feasible to solve optimal control problems online without strict real-time constraints [1]. In recent years, advances in computational hardware have enabled the widespread use of MPC in robotics, where system dynamics are significantly faster than in traditional applications. To meet real-time requirements, these implementations often rely on simplified internal models, such as reduced-order or template dynamics [2]. Despite ongoing efforts to accelerate computation, the optimization time required to solve an MPC problem at each control step remains a significant bottleneck, making the application of MPC to fast-paced and highly nonlinear robotic systems an open and active research topic [3]. These limitations are particularly pronounced in




Fig. 1. Block scheme describing the Feedback-MPPI method. The proposed gain computation (red path) can leverage the same parallelization capabilities of the standard MPPI algorithm. Braces represent the stack of multiple signals.

tasks requiring high-frequency control, such as agile motion generation and dynamic physical interaction.

To address this computational bottleneck, several approaches have been developed. Methods like distributed optimization and transformer-based constraint handling offer improvements but require significant framework modifications or specialized hardware [4]–[6]. Another strategy approximates MPC policies directly, for example, using neural networks, enabling near-instantaneous control after initial training [7].

Gradient-based methods have achieved great efficiency by linearizing dynamics and constraints at each timestep, notably through Real-Time Iteration schemes [8]. To achieve higher control frequencies, some works proposed the use of equivalent feedback gains that locally approximate the MPC solution, enabling more efficient closed-loop control [9]–[12]. For instance, Dantec et al. [9] use sensitivity analysis in a purely model-based context to derive a high-frequency feedback controller for humanoid robots. Their method computes a first-order approximation of the MPC solution using Riccati gains extracted from a Differential Dynamic Programming (DDP) solver, enabling a fast inner-loop controller to operate at 2 kHz without re-solving the full MPC at each cycle. This effectively bridges the gap between slow optimal planning and fast low-level control. Similar ideas have been exploited for quadruped locomotion [10], [11] and for MPC based on Non Linear Programming (NLP) [12]. Building on this idea, Hose et al. [13] extend sensitivity-based approximation to the learning domain. They propose a parameter-adaptive scheme in which a neural network is first trained to imitate the nominal MPC policy, and then locally corrected at runtime using first-order sensitivities of the MPC solution with respect to varying

Manuscript received: June, 13, 2025; Revised September, 10, 2025; Accepted October, 24, 2025. This paper was recommended for publication by Editor A. Kheddar upon evaluation of the Associate Editor and Reviewers' comments. Experiments presented in this paper were carried out thanks to a platform of the Robotex 2.0 French research infrastructure.

T. Belvedere is with CNRS, Univ Rennes, Inria, IRISA, Campus de Beaulieu, Rennes, France. E-mail: tommaso.belvedere@irisa.fr

M. Ziegltrum and V. Modugno are with the Department of Computer Science, University College London, Gower Street, WC1E 6BT, London, UK. E-mail: {michael.ziegltrum.24, v.modugno}@ucl.ac.uk

G. Turrisi is with the Dynamic Legged Systems Laboratory, Istituto Italiano di Tecnologia (IIT), Genova, Italy. E-mail: giulio.turrisi@iit.it

Code and videos available at: <https://feedback-mppi.github.io>

Digital Object Identifier (DOI): 10.1109/LRA.2025.3630871

system parameters. This allows the learned controller to generalize across different system configurations without retraining, making it robust and suitable for deployment on embedded hardware with limited resources. In a similar vein, Tokmak et al. [14] introduce an automated method to approximate nonlinear MPC solutions while preserving closed-loop guarantees. Their approach, ALKIA-X, uses kernel interpolation augmented with local sensitivity information to construct an explicit, non-iterative approximation of the MPC control law.

An alternative to traditional gradient-based MPC methods is the Model Predictive Path Integral (MPPI) control, a stochastic, sampling-based technique capable of handling highly nonlinear, non-convex problems without explicit gradients. MPPI has been effectively applied to aggressive autonomous driving [15], [16], agile quadrotor maneuvers [17], and contact-rich quadruped locomotion [18], [19]. However, MPPI's computational load and noisy actions remain challenges, and can cause oscillations during hovering [17]. Recent advancements such as GPU-accelerated computations [19], low-pass filtering [20], and adaptive sampling using learned priors [21] aim to enhance real-time applicability and stability.

To further address these challenges, this work proposes integrating local linear feedback approximations into sampling-based MPPI controllers, inspired by Riccati-based feedback strategies from gradient-based MPC. The proposed method (Fig. 1) enhances responsiveness without necessitating full re-optimization at each step, unlocking high-frequency feedback control. The main contributions of this paper are:

- The introduction of a novel method for computing local linear feedback gains within MPPI, extending the Riccati-based feedback MPC framework to stochastic sampling-based controllers.
- Demonstrating the effectiveness of the approach in approximating the MPPI action to provide high-frequency feedback, improving accuracy, and action smoothness.
- The approach is validated on quadrupedal robots in simulation and on aerial drones in real-world experiments, both performing dynamic motions.

Our results bridge the gap between sampling-based and gradient-based MPC, demonstrating that local feedback-enhanced sampling-based MPC as a viable strategy for real-time, complex robotic systems.

The remainder of the paper is structured as follows: in Sect. II we describe the proposed methodology, first presenting standard MPPI (Sect. II-A) and then extending it by computing a local approximation of its optimal solution from local variations in the initial state (Sect. II-B). This is first validated by comparing it with a linear quadratic regulator for a toy problem (Sect. II-D) and, in Sect. III, through simulations and experiments on two robotic platforms, a quadruped legged robot and a quadrotor aerial vehicle.

II. METHODOLOGY

A. Background on Model Predictive Path Integral Control

Model Predictive Path Integral (MPPI) control is a sampling-based stochastic optimization approach widely employed to solve optimal control problems characterized by

nonlinear dynamics, complex, non-convex cost functions, and potentially non-differentiable constraints [17], [19]. MPPI addresses optimal control problems defined as follows

$$\begin{aligned} \theta^* = \arg \min_{\theta} \quad & \ell_N(\mathbf{x}_N) + \sum_{i=0}^{N-1} \ell_i(\mathbf{x}_i, \mathbf{u}_i) \\ \text{s.t.} \quad & \mathbf{x}_0 = \hat{\mathbf{x}} \\ & \mathbf{x}_{i+1} = \mathbf{f}(\mathbf{x}_i, \mathbf{u}_i) \\ & \mathbf{u}_i = \pi(\theta, \mathbf{x}_i, t_i), \end{aligned} \quad (1)$$

where θ denotes the decision variables, or parameters, \mathbf{u}_i the control inputs at time t_i , \mathbf{x}_i represents the system state evolving according to the dynamics $\mathbf{f}(\mathbf{x}, \mathbf{u})$, ℓ_i denotes the cost functions defined over the control horizon N and, finally, $\hat{\mathbf{x}}$ indicates the current feedback state. For generality, we consider the control inputs as parametrized by a possibly state dependent function $\pi(\theta, \mathbf{x}, t)$ that maps sampled parameters θ into an input trajectory, possibly also clipping inputs to their upper/lower bounds. This can be used to decrease the dimensionality of the search space and to improve the smoothness of the input trajectory [21]. Possible choices are direct sampling (zero-order) [22], cubic splines [19], or Halton splines [23].

In MPPI, parameter samples are drawn from a Gaussian distribution centered at an initial guess $\bar{\theta}$, which may be obtained from the solution at the previous step, according to $\theta \sim \mathcal{N}(\bar{\theta}, \Sigma)$. Each parameter sample $\theta^k = \bar{\theta} + \Delta\theta^k$ generates a rollout by propagating the system dynamics forward in time, producing a set of candidate trajectories. Each trajectory is associated with a cost, with J^k denoting the total cost associated with the k -th sampled trajectory:

$$J^k := \ell_N(\mathbf{x}_N^k) + \sum_{i=0}^{N-1} \ell_i(\mathbf{x}_i^k, \mathbf{u}_i^k). \quad (2)$$

The optimal parameter update in MPPI is then computed using a weighted average of parameter perturbations:

$$\theta^* = \bar{\theta} + \sum_{k=0}^K \omega^k \Delta\theta^k, \quad (3)$$

with trajectory-specific weights ω^k determined by their costs J^k as follows:

$$\omega^k = \frac{\mu^k}{\sum_{j=1}^K \mu^j}, \quad \text{with } \mu^k = \exp\left(-\frac{1}{\lambda} (J^k - \rho)\right), \quad (4)$$

where $\rho = \min_{k=1, \dots, K} J^k$, and λ is a tuning parameter balancing exploration and exploitation [17], [24]. In practical implementations, MPPI utilizes GPU acceleration for parallelization of trajectory evaluations, significantly reducing computational time and facilitating real-time performance on complex systems such as agile aerial and legged robots [17], [19]. The computed optimal parameters from (3) directly yield the commanded control action via:

$$\mathbf{u}^* = \pi(\theta^*, \hat{\mathbf{x}}, t_0). \quad (5)$$

In standard MPPI frameworks, the resulting optimal action is commanded directly to the system and typically held constant over the sampling interval.

B. Feedback MPPI

A notable drawback of MPC controllers, MPPI included, is that the control frequency is directly tied to the computational rate of the MPPI solution, which is constrained by the platform's processing capabilities. To address this, we introduce a method for computing a first-order approximation of the MPPI action. This approximation enables high-frequency feedback for systems demanding high control bandwidth, eliminating the need for an additional tracking controller.

In order to compute a local approximation of the MPPI action, we derive the sensitivity of the optimal inputs \mathbf{u}^* in (5) to variations in the current initial state $\hat{\mathbf{x}}$. This quantity, denoted as $\mathbf{F} \in \mathbb{R}^{n_u \times n_x}$, maps state variations into inputs and can be used as a local feedback control gain to obtain a new action of the kind

$$\mathbf{u} = \mathbf{u}^* + \mathbf{F}(\hat{\mathbf{x}} - \mathbf{x}_{\text{sp}}), \quad (6)$$

where \mathbf{x}_{sp} represents a local setpoint acquired from the MPPI solution¹ which will be tracked by the inner loop.

To derive an expression for the MPPI gain matrix \mathbf{F} , it is sufficient to differentiate the optimal solution obtained through importance sampling (3)–(4). In fact, applying the chain rule to (5) leads to:

$$\mathbf{F} = \frac{\partial \mathbf{u}^*}{\partial \mathbf{x}_0} = \frac{\partial \pi}{\partial \theta} \frac{\partial \theta^*}{\partial \mathbf{x}_0} + \frac{\partial \pi}{\partial \mathbf{x}_0}, \quad (7)$$

where $\frac{\partial \pi}{\partial \theta}$ and $\frac{\partial \pi}{\partial \mathbf{x}_0}$ depend on the specific policy parametrization choice, e.g., direct sampling or cubic splines. Then, one has to differentiate the optimal parameters θ^* . Recalling (3):

$$\frac{\partial \theta^*}{\partial \mathbf{x}_0} = \sum_{k=0}^K \Delta \theta^k \frac{\partial \omega^k}{\partial \mathbf{x}_0}. \quad (8)$$

Then, noting that

$$\frac{\partial \mu^k}{\partial \mathbf{x}_0} = -\frac{\mu^k}{\lambda} \left(\frac{\partial J^k}{\partial \mathbf{x}_0} - \frac{\partial \rho}{\partial \mathbf{x}_0} \right), \quad (9)$$

applying the quotient rule to (4) and simplifying, yields

$$\frac{\partial \omega^k}{\partial \mathbf{x}_0} = \frac{\omega^k}{\lambda} \left(\frac{\partial J^k}{\partial \mathbf{x}_0} - \sum_{j=1}^K \omega^j \frac{\partial J^j}{\partial \mathbf{x}_0} \right). \quad (10)$$

Plugging back to (8) one has

$$\frac{\partial \theta^*}{\partial \mathbf{x}_0} = \sum_{k=0}^K \Delta \theta^k \frac{\omega^k}{\lambda} \left(\frac{\partial J^k}{\partial \mathbf{x}_0} - \sum_{j=1}^K \omega^j \frac{\partial J^j}{\partial \mathbf{x}_0} \right), \quad (11)$$

which can be combined with (7) to obtain the desired MPPI gains, depending on the chosen control parametrization:

$$\mathbf{F} = \sum_{k=0}^K \frac{\partial \pi}{\partial \theta} \Delta \theta^k \frac{\omega^k}{\lambda} \left(\frac{\partial J^k}{\partial \mathbf{x}_0} - \sum_{j=1}^K \omega^j \frac{\partial J^j}{\partial \mathbf{x}_0} \right) + \frac{\partial \pi}{\partial \mathbf{x}_0}. \quad (12)$$

For instance, in the case of direct control sampling, where $\theta = (\mathbf{u}_0, \dots, \mathbf{u}_{N-1})$, with $\frac{\partial \pi}{\partial \theta} = (\mathbf{I}, \mathbf{0}, \dots, \mathbf{0})$ and $\frac{\partial \pi}{\partial \mathbf{x}_0} = \mathbf{0}$:

$$\mathbf{F} = \sum_{k=0}^K \Delta \mathbf{u}_0^k \frac{\omega^k}{\lambda} \left(\frac{\partial J^k}{\partial \mathbf{x}_0} - \sum_{j=1}^K \omega^j \frac{\partial J^j}{\partial \mathbf{x}_0} \right). \quad (13)$$

¹In contrast to several works that simply set this setpoint as the initial state \mathbf{x}_0 at which the solution was computed, we draw inspiration from [25] and continuously update \mathbf{x}_{sp} by linearly interpolating between \mathbf{x}_0 and \mathbf{x}_1 .

Similar expressions can be derived also in the case of different input parametrization.

Remark 1. One of the advantages of sampling-based MPC over gradient-based MPC is the ability to optimize over discontinuous optimal control problems. Since our gain computation relies on the gradients of the rollout costs, we do require local differentiability of the dynamics and costs in (1). This is, however, not a limitation: in fact, most robotic applications are described by at least piecewise smooth functions. Indeed, such gains will only approximate the local behavior of the system away from discontinuities. This is similar to the fact that linear MPC approximations do not track active set changes when dealing with inequality constraints [12]. Still, while computing gains requires local differentiability of the rollouts, the ability of the solver to explore the solution space and optimize over discontinuous non-convex problems is not affected.

Remark 2. In many MPPI implementations, including ours, generic inequality constraints can be managed using indicator functions by adding a large constant cost to trajectories that violate the constraint. Since this cost term is (locally) constant, it does not influence the gain computation, meaning the feedback system remains unaware of such constraints. A straightforward way to overcome this limitation is to incorporate constraints through smooth barrier functions, as demonstrated in Feedback-MPC [10]. On the other hand, since input constraints are treated by directly clipping the sampled trajectory inside the rollouts, MPPI gains will naturally account for such constraints if active.

C. Implementation

Notably, the proposed gain computation (12) still makes use of parallel computation to achieve real-time performance even with a large number of samples. In fact, the cost gradients $\frac{\partial J^k}{\partial \mathbf{x}_0}$ can be evaluated separately for each rollout and later combined with the appropriate trajectory weights ω^k . The whole Feedback-MPPI algorithm has been implemented in a standalone Python library that we released as open source. This library makes use of JAX to run JIT compiled code on CUDA-compatible GPUs to perform parallel rollouts. Moreover, we make use of the Automatic Differentiation capabilities provided by JAX to compute the cost gradients in (12) in parallel together with the rollouts, improving the computational efficiency of the method.

D. Illustrative example

To assess the efficacy of the proposed method in determining the optimal gains for the MPPI solution, we first illustrate its behavior through a quantitative example involving an inverted pendulum system. These gains are compared to those obtained by applying an infinite horizon discrete-time Linear Quadratic Regulator (LQR) to the linearized system around the upward equilibrium. Although a formal proof of the relationship between MPPI and Riccati gains is not available at the time of writing, it is reasonable to assume that the methods should yield similar gains under similar circumstances.

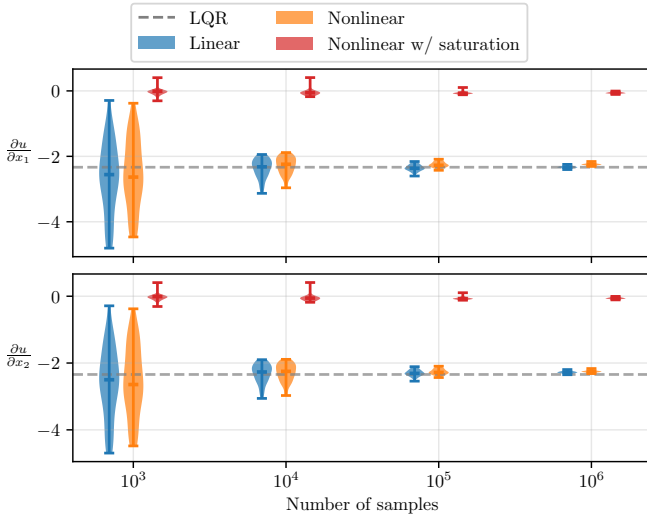


Fig. 2. Comparison between MPPI gains obtained for the inverted pendulum system under different conditions (linear/nonlinear rollout, input saturation) for an increasing number of samples, and LQR gains computed around the upward equilibrium (dashed gray line).

In this scenario, the algebraic Riccati equation was solved using $\mathbf{Q} = \mathbf{I}$ and $\mathbf{R} = 1$, resulting in the optimal Riccati gain matrix \mathbf{F}_{LQR} and the solution matrix \mathbf{S} . Next, we formulated an equivalent MPPI problem, aligning the running cost with that of the LQR and setting the terminal cost as $\ell_N = \mathbf{x}^T \mathbf{S} \mathbf{x}$ to solve the infinite horizon problem. Three cases are considered for MPPI: first, using the same linearized dynamics as the LQR, then with the nonlinear inverted pendulum dynamics and, finally, including input bounds to saturate the optimal solution. The MPPI was then solved multiple times from an initial condition $\mathbf{x}_0 = (\pi/8, 0.0)$, with gains computed using (12). This procedure was repeated while varying the number of samples K . The results in Fig. 2 demonstrate that in the unconstrained cases (Linear and Nonlinear) MPPI gains $\mathbf{F} = \left(\frac{\partial u}{\partial x_1}, \frac{\partial u}{\partial x_2} \right)$ are on average similar to LQR optimal gains for both linear and nonlinear rollouts, with a variance shrinking significantly as the number of samples increases. As the gain computation relies on the linearization around the rollout trajectories, minor differences can be ascribed to the different state trajectories obtained with the linearized model. In the input constrained case, MPPI gains instead converge to zero: this effect, which correctly predicts the fact that no additional feedback action can be added to the control, stems from the fact that the partial derivative $\frac{\partial \pi}{\partial \theta}$ collapses to zero when inputs are clipped.

III. APPLICATIONS

Riccati-like gains have consistently demonstrated their effectiveness in various applications through numerous studies and practical implementations [9]–[11], [25], [26]. Given this well-established track record, our objective in this section is not to revalidate their utility. Instead, we aim to illustrate how these gains can be successfully integrated within the framework of sampling-based MPC running at low frequency to dynamically control two robotic platforms, and how the ad-

TABLE I
MPPI SETTINGS FOR THE QUADRUPEL EXPERIMENTS

	Symbol	Value	
Solver	N	10 steps	
	δ_t	0.02 s	
	K	5000 samples	
	Σ	$9\mathbf{I}$	
	λ	1	
Weights	\mathbf{Q}_r	Running	Terminal
	\mathbf{Q}_v	diag(0, 0, 1500)	diag(0, 0, 1500)
	\mathbf{Q}_ϕ	diag(200, 200, 200)	diag(200, 200, 200)
	\mathbf{Q}_ω	diag(1500, 1500, 0)	diag(1500, 1500, 0)
	\mathbf{R}	diag(20, 20, 50)	diag(20, 20, 50)
		$10^{-4}\mathbf{I}$	—

ditional computational effort balances with performance gains even when compared to MPPI running at higher frequencies.

A. Quadruped motion control

We first test the proposed method on a simulated quadruped robot based on Aliengo², a 24kg electric quadruped developed by Unitree. The simulations are performed in MuJoCo, running on a laptop with an Intel i7-13700H CPU and an Nvidia 4050 6GB GPU. All relevant settings are reported in Tab. I.

The dynamical model of the system adopts the simplified Single Rigid Body Dynamics model formulation [19]. This model focuses on the quadruped’s core translational and rotational movements, omitting the swinging legs’ dynamics given their typically small mass in comparison with the trunk of the robot. The robot’s dynamics is centered around the CoM frame, and is described using two reference frames: an inertial frame \mathcal{W} , and a body-aligned frame \mathcal{C} at the Center of Mass (CoM). The state vector is

$$\mathbf{x} = (\mathbf{r}_c, \mathbf{v}_c, \phi, \omega) \in \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}^3,$$

with position $\mathbf{r}_c = (x, y, z)$ and velocity $\mathbf{v} = (v_x, v_y, v_z)$ expressed in the inertial frame, the robot body orientation $\phi = (\phi, \theta, \psi)$ where ϕ, θ, ψ , are the roll, pitch, and yaw respectively, and the angular velocity $\omega = (\omega_x, \omega_y, \omega_z)$ expressed in body frame. Finally, the equations of motion read as

$$\begin{aligned} \dot{\mathbf{r}}_c &= \mathbf{v}_c \\ \dot{\mathbf{v}}_c &= \frac{1}{m} \sum_{i=1}^4 \delta_i \mathbf{f}_i + \mathbf{g} \\ \dot{\phi} &= \mathbf{E}'^{-1}(\phi) \omega \\ \dot{\omega} &= -\mathbf{I}_c^{-1} (\omega \times \mathbf{I}_c) \omega + \sum_{i=1}^4 \delta_i \mathbf{I}_c^{-1} \mathbf{p}_i \times \mathbf{f}_i, \end{aligned}$$

with m the robot mass, \mathbf{g} the gravitational acceleration and $\mathbf{I}_c \in \mathbb{R}^{3 \times 3}$ the constant inertia tensor centered at the robot’s CoM; \mathbf{E}'^{-1} is a mapping from the robot’s angular velocity to Euler rates; $\mathbf{p}_i \in \mathbb{R}^3$ is the displacement vector between the CoM position \mathbf{r}_c and the i -th robot’s foot. Binary variables $\delta_i = \{0, 1\}$, extracted from a precomputed periodic gait sequence, indicate whether an end-effector makes contact with the environment and can produce Ground Reaction Forces \mathbf{f}_i (GRF), which we optimize in our controller formulation.

²Aliengo quadruped robot: <https://www.unitree.com/products/aliengo>




Fig. 3. Top: Tracking results for the quadruped simulation of F-MPPI against MPPI running at three different control frequencies; Bottom: A snapshot of the quadruped robot tracking a desired velocity (orange arrow) on a randomly generated uneven terrain, while being perturbed by an external disturbance (red arrow). Green arrows represent the GRFs, while the blue arrow depicts the actual robot's velocity.

1) *Simulation results:* In the simulation, we ask the robot to track a set of linear velocity references (between 0, 0.5 m/s) over a randomly generated rough terrain while being subject to random disturbances (between ± 5 Nm). For this task, the cost function is designed to track a desired velocity reference (linear and angular) while maintaining a desired posture (height, roll, and pitch of the robot). This is achieved by incorporating a weighted quadratic state error plus a regularization term for the GRF (gravity compensation)

$$\ell_{\text{track}} = \|\mathbf{x} - \mathbf{x}_{\text{track}}\|_{\mathbf{Q}}^2 + \|\mathbf{u} - \mathbf{u}_{\text{reg}}\|_{\mathbf{R}}^2. \quad (14)$$

The control inputs (GRFs) are sampled using linear splines and are clipped to enforce friction cone constraints for non-slipping conditions [19].

We run F-MPPI at 50Hz, and compare its performance to standard MPPI at three different control frequencies, 50Hz, 80Hz, and 100Hz. We selected 50Hz as base frequency to be consistent with the real quadrotor experiments in Sect.III-B1, where such frequency is limited by the onboard hardware, and to simulate the use of a similar, less powerful, embedded system. Then, we progressively increased the MPPI frequency to reflect the fact that the increased computational cost of F-MPPI relatively decreases the maximum control frequency

TABLE II
MPPI SETTINGS FOR THE QUADROTOR EXPERIMENTS

	Symbol	Value	
Solver	N	15 steps	
	δ_t	0.05 s	
	K	800 samples	
	Σ	25I	
	λ	1	
Weights		Running	Terminal
	\mathbf{Q}_r	diag(100, 100, 125)	diag(100, 100, 125)
	\mathbf{Q}_v	diag(0.5, 0.5, 2.5)	diag(10, 10, 25)
	\mathbf{Q}_q	diag(0.5, 0.5, 50)	diag(10, 10, 100)
	\mathbf{Q}_ω	diag(0.25, 0.25, 25)	diag(0.5, 0.5, 50)
	\mathbf{R}	$10^{-2}\mathbf{I}$	—
	Q_{obs}	10^6	10^6

for a given computational budget. In all cases, we sample 5000 parameters, bringing the computational complexity of F-MPPI to 3 ms per run, while that of MPPI to 2 ms. Our goal is to show that utilizing the MPPI gains brings benefits to the final performance of the system, even when accounting for a decrease in the control frequency due to the additional computations. In all cases, a low-level controller that runs at 500Hz maps the GRFs to joint torques via a simple Jacobian mapping (see [19]). When F-MPPI is used, this mapping consistently relies on GRFs updated at the same rate (500Hz), leveraging the gains calculated at a lower frequency by our algorithm.

The final results, which report the Mean Absolute Errors over 50 trials, are shown in Fig. 3. F-MPPI, computed at 50 Hz, remains competitive even against an MPPI running two times faster (100 Hz), and outperforms MPPI running at a similar computational budget (80 Hz).

B. Quadrotor motion control

The second robotic platform consists of a quadrotor robot based on MikroKopter and running on the Telekyb3 framework for the hardware interface and state estimation. An Unscented Kalman Filter implemented in Telekyb3 is used to fuse the IMU measurements and the motion capture position feedback to obtain an estimate of the state $\hat{\mathbf{x}}$. First, the proposed method has been tested in a Software-In-The-Loop simulation using a Gazebo environment where the robot firmware, state estimation, and motor dynamics are simulated to mimic real-world experiments. The simulation runs on a laptop with an Intel i7-13700H CPU and an Nvidia A1000 6GB GPU. In the experiments, the F-MPPI controller is fully executed *onboard* on an Nvidia Jetson Orin NX 16GB. The same implementation and MPPI settings used in the simulation have been employed in the experiments, with only minor differences in the hardware interface. All relevant settings are reported in Tab. II.

Following the scheme depicted in Fig. 1, the F-MPPI algorithm operates at a frequency of 50 Hz, while the low-level control module, based on the MPPI gain matrix, updates the control action at a frequency of 200 Hz.

We adopt the standard rigid-body model of a quadrotor whose center of mass coincides with its geometric center [27]. The state vector is

$$\mathbf{x} = (\mathbf{r}, \mathbf{v}, \mathbf{q}, \boldsymbol{\omega}) \in \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{S}^3 \times \mathbb{R}^3,$$




Fig. 4. Evolution of the quadrotor position in the x - z plane, stroboscopic view of the quadrotor motion and commanded rotor speeds during the experiment. In blue is the proposed framework, while in orange is the standard MPPI controller running at the maximum frequency allowed by the hardware platform. Thanks to the proposed method, the resulting control action is smoother and the motion more precise, with a reduction of 64.6% and 28.9% in RMSE for the x and z components, respectively. See the supplementary material for a clip of the experiment.

with position $\mathbf{r} = (x, y, z)$ and velocity $\mathbf{v} = (v_x, v_y, v_z)$ expressed in the inertial frame, the unit quaternion $\mathbf{q} = (q_w, q_x, q_y, q_z)$ describing the body attitude and the angular velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)$ expressed in body frame, and \mathbb{S}^3 representing the unit 3-sphere.

With the rotor speeds as inputs $\mathbf{u} = (w_1, w_2, w_3, w_4)$, the total thrust f and control torque $\boldsymbol{\tau}$ generated in body frame are obtained through the allocation matrix $(f, \boldsymbol{\tau}) = \mathbf{T}(l, k_f, k_m)\mathbf{u}^2$, where l denotes the quadrotor arm length and k_f, k_m the propeller aerodynamic coefficients [27]. The equations of motion therefore read

$$\begin{aligned} \dot{\mathbf{r}} &= \mathbf{v}, & \dot{\mathbf{q}} &= \frac{1}{2} \mathbf{q} \otimes \begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix}, \\ \dot{\mathbf{v}} &= \frac{1}{m} \mathbf{R}(\mathbf{q}) \begin{bmatrix} 0 \\ 0 \\ f \end{bmatrix} + \mathbf{g}, & \dot{\boldsymbol{\omega}} &= \mathbf{I}_c^{-1} (\boldsymbol{\tau} - \boldsymbol{\omega} \times \mathbf{I}_c \boldsymbol{\omega}), \end{aligned}$$

where $\mathbf{R}(\mathbf{q}) \in \text{SO}(3)$ is the rotation matrix associated with \mathbf{q} , m is the total mass, \mathbf{g} the gravity force vector, and \mathbf{I}_c the inertia tensor centered at the robot CoM.

Both the simulation and experiments use the same base cost function designed to reach a position goal while trading off control effort. The cost function term for achieving the goal incorporates a weighted quadratic state error:

$$\ell_{\text{goal}} = \|\mathbf{r} - \mathbf{r}_{\text{goal}}\|_{\mathbf{Q}_r}^2 + \|\mathbf{v}\|_{\mathbf{Q}_v}^2 + \|\mathbf{q} - \mathbf{q}_0\|_{\mathbf{Q}_q}^2 + \|\boldsymbol{\omega}\|_{\mathbf{Q}_\omega}^2, \quad (15)$$

where $\mathbf{q}_0 = (1, 0, 0, 0)$. Inputs are sampled using cubic splines with $N_{\text{knot}} = 5$ knot points. The sampled trajectory is then clipped to enforce input limits. To minimize control effort and regularize the solution, the input is kept close to the constant hovering input \mathbf{u}_h with a cost:

$$\ell_{\mathbf{u}} = \|\mathbf{u} - \mathbf{u}_h\|_{\mathbf{R}}^2. \quad (16)$$

1) *Experimental results:* The effectiveness of the F-MPPI method is now demonstrated in a real-world scenario where computational resources are also limited. Specifically, we ex-

cute standard MPPI on the onboard computer at its maximum feasible frequency of 66.7 Hz. During these tests, the quadrotor is tasked with reaching a goal displaced of (1, 0.5) m in the x - z plane. The overall system is tuned for best performance with the standard MPPI running at 66.7 Hz, while the frequency is lowered to 50 Hz for F-MPPI to meet real-time constraints without modifying the problem.

The experimental results, illustrated in Fig. 4, highlight that the proposed F-MPPI method not only reaches the target more quickly but also with superior precision compared to standard MPPI. To quantitatively evaluate performance, we calculate the Root Mean Square Error (RMSE) for both the x and z components after the initial transient phase. The F-MPPI method achieves RMSE values of 0.017 m and 0.027 m for the x and z components, respectively, while the standard MPPI yields higher RMSEs of 0.048 m and 0.038 m.

An additional advantage of F-MPPI is the significantly smoother control commands it generates, as shown in Fig. 4 (bottom right). This smoothness can have practical benefits, such as reducing actuator wear and lowering power consumption, thus enhancing the overall efficiency and longevity of the system.

2) *Computation time analysis:* In order to study the impact of the proposed method on the computation time, we have collected data by running several simulations of the quadrotor on the onboard Jetson computer with and without the MPPI gains calculation. Results obtained with different control horizons N and number of samples K are reported in Fig. 5. The difference between F-MPPI and standard MPPI can be ascribed to the gain computation, which increases overall runtime by roughly 40%–70%. The overhead grows with the number of samples and the prediction horizon; it reaches the high end of this range when $K = 1200$, where the rollouts outnumber the available GPU cores.

It should be recognized that this approach has a higher




Fig. 5. Computation times for the proposed method (F-MPPI) and standard MPPI in the quadrotor regulation problem, analyzed across varying control horizons and numbers of parallel samples. The analysis was conducted on an Nvidia Jetson Orin NX 16GB with 1024 GPU cores.

computational burden compared to methodologies developed for gradient-based MPC that rely on Riccati recursion [9] or NLP sensitivities [26]. Ultimately, the performance gains will be determined by the specific system characteristics that dictate whether the additional computational cost is worthwhile. Nevertheless, we demonstrated that the proposed method is a viable option for sampling-based MPC. As evidenced by the experimental results, even running F-MPPI at the low frequency of 50 Hz can yield significant performance improvements compared to standard MPPI thanks to the high-frequency feedback provided by the MPPI gains implemented in the inner-loop.

C. Application to hybrid systems and non-convex constraints

As the proposed method relies on differentiating the rollout trajectories, MPPI gains only provide a first-order approximation of the optimal solution (cfr. Remark 1). However, we demonstrate that this limitation does not hinder the method’s applicability to discontinuous problems arising from hybrid dynamical systems or the presence of inequality constraints.

1D Hopper. The method was applied to a 1D Hopper robot whose hybrid dynamics is simulated through Mujoco MJX. This robot can extend its leg to push on the ground and propel upwards. However, due to the leg extension being limited, it has to resort to a bouncing motion to reach higher positions. In this setting, we have implemented the F-MPPI method running at 50 Hz with gains closing the low-level loop at 500 Hz. The robot starts from the desired height of 0.5 m and repeatedly bounces around the same height. To mimic a real-world scenario, the velocity of the robot is




Fig. 6. Comparison of the height evolutions of the 1D hopper robot under the two different control laws. The results were obtained by directly performing sampling and differentiation through the dynamics simulator.

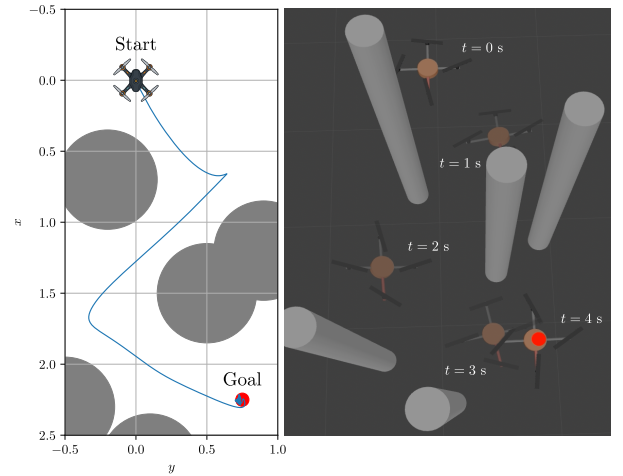


Fig. 7. Path executed by the robot when navigating among obstacles. Top-down view with obstacles’ radii augmented by the quadrotor dimension (left). Stroboscopic view of the motion in the Gazebo simulator (right). See the supplementary material for a clip of the simulation.

perturbed with Gaussian noise. Figure 6 shows results with and without the use of MPPI gains. Without the additional correction provided by the inner loop, disturbances make the robot enter a local minimum, remaining stuck to the ground. On the other hand, the proposed method is able to complete the motion, demonstrating how MPPI gains are able to locally regulate the system across the two hybrid modes.

Quadrotor navigation among obstacles. Despite generic constraints not being explicitly accounted for in the gain computation (cfr. Remark 2), the method can effortlessly incorporate them and handle non-convex optimization problems. To illustrate this, we conducted simulations involving navigating through a series of stationary circular obstacles to reach the designated goal position. Following [17], the

obstacle avoidance constraint is enforced by incorporating a barrier term into the cost function. This significantly penalizes trajectories that collide with the environment. The barrier term can be added to the cost as $\ell_{\text{obs}} = Q_{\text{obs}} \mathbf{1}_{\mathbf{x} \in \chi_{\text{obs}}}$, where χ_{obs} represents the space occupied by obstacles.

Figure 7 depicts the path executed in the navigation scenario. The robot successfully navigates to the goal in approximately 4 s, avoiding the obstacles or getting stuck in local minima, before hovering at the goal position.

IV. CONCLUSION AND FUTURE WORK

In this letter, we introduced a way to efficiently compute an approximation of the sampling-based MPPI predictive control method in the form of linear feedback gains. This has allowed to extend the typical paradigm used in DDP-like and gradient-based methods for producing high-frequency MPC actions through the use of Riccati gains. The effectiveness of the method has been demonstrated on two robotic platforms, a legged robot and an aerial vehicle, through simulations and experiments. Results show that, similarly to related works, introducing such high-frequency MPPI approximation provides a benefit over plain MPPI, even when accounting for the higher computational cost. This development paves the way to several possible advancements in sampling-based MPC where such gains could be embedded, such as sensitivity-based robust MPC [26], safe reinforcement learning [28], and data augmentation for sample efficient learning [29]. Furthermore, future research will delve into computational aspects to streamline the gains calculation, for instance, optimizing the backpropagation process or further leveraging parallelization. Additionally, techniques to incorporate inequality constraints into the gains calculation will be investigated.

ACKNOWLEDGMENT

The authors express their gratitude to Odichimma Ezeji for her assistance in implementing obstacle avoidance in the quadrotor simulation.

REFERENCES

- [1] S. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.
- [2] S. Katayama, M. Murooka, and Y. Tazaki, "Model predictive control of legged and humanoid robots: models and algorithms," *Advanced Robotics*, vol. 37, no. 5, pp. 298–315, 2023.
- [3] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. D. Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2024.
- [4] L. Amatiucci, G. Turrissi, A. Bratta, V. Barasuol, and C. Semini, "Accelerating model predictive control for legged robots through distributed optimization," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 1–7.
- [5] E. Adabag, M. Atal, W. Gerard, and B. Plancher, "Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 9787–9794.
- [6] K. Nguyen, S. Schoedel, A. Alavilli, B. Plancher, and Z. Manchester, "TinyMPC: Model-predictive control on resource-constrained micro-controllers," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 1–7.
- [7] M. Hertneck, J. Köhler, S. Trimpe, and F. Allgöwer, "Learning an approximate model predictive controller with guarantees," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 543–548, 2018.
- [8] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM Journal on Control and Optimization*, vol. 43, no. 5, pp. 1714–1736, 2005.
- [9] E. Dantec, M. Taïx, and N. Mansard, "First order approximation of model predictive control solutions for high frequency feedback," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, p. 4448–4455, 2022.
- [10] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4730–4737.
- [11] H. Li and P. M. Wensing, "Cafe-mpc: A cascaded-fidelity model predictive control framework with tuning-free whole-body control," *IEEE Transactions on Robotics*, vol. 41, pp. 837–856, 2025.
- [12] V. M. Zavala and L. T. Biegler, "The advanced-step nmmpc controller: Optimality, stability and robustness," *Automatica*, vol. 45, no. 1, p. 86–93, 2009.
- [13] H. Hose, A. Gräfe, and S. Trimpe, "Parameter-adaptive approximate mpc: Tuning neural-network controllers without retraining," in *Proc. 6th Conf. on Learning for Dynamics and Control (LADC)*, vol. 242, 2024, pp. 349–360.
- [14] A. Tokmak, C. Fiedler, M. N. Zeilinger, S. Trimpe, and J. Köhler, "Automatic nonlinear mpc approximation with closed-loop guarantees," *IEEE Trans. on Automatic Control*, 2025.
- [15] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [16] 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.
- [17] M. Minařík, R. Pěnička, V. Vonásek, and M. Saska, "Model predictive path integral control for agile unmanned aerial vehicles," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 13 144–13 151.
- [18] J. Carius, R. Ranftl, F. Farshidian, and M. Hutter, "Constrained stochastic optimal control with learned importance sampling: A path integral approach," *The International Journal of Robotics Research*, vol. 41, no. 2, pp. 189–209, 2022.
- [19] G. Turrissi, V. Modugno, L. Amatiucci, D. Kanoulas, and C. Semini, "On the benefits of gpu sample-based stochastic predictive controllers for legged locomotion," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 13 757–13 764.
- [20] P. Kicki, "Low-pass sampling in model predictive path integral control," *arXiv preprint arXiv:2503.11717*, 2025.
- [21] T. Howell, N. Gileadi, S. Tunyasuvunakool, K. Zakka, T. Erez, and Y. Tassa, "Predictive Sampling: Real-time Behaviour Synthesis with MuJoCo," 2022, arXiv:2212.00541.
- [22] 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.
- [23] C. Pezzato, C. Salmi, E. Trevisan, M. Spahn, J. Alonso-Mora, and C. Hernández Corbato, "Sampling-based model predictive control leveraging parallelizable physics simulations," *IEEE Robotics and Automation Letters*, vol. 10, no. 3, pp. 2750–2757, 2025.
- [24] G. Rizzi, J. J. Chung, A. Gawel, L. Ott, M. Tognon, and R. Siegwart, "Robust Sampling-Based Control of Mobile Manipulators for Interaction With Articulated Objects," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1929–1946, 2023.
- [25] R. Subburaman and O. Stasse, "Delay robust model predictive control for whole-body torque control of humanoids," in *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, 2024, p. 600–606.
- [26] T. Belvedere, M. Cognetti, G. Oriolo, and P. Robuffo Giordano, "Sensitivity-aware model predictive control for robots with parametric uncertainty," *IEEE Transactions on Robotics*, vol. 41, pp. 3039–3058, 2025.
- [27] G. Antonelli, E. Cataldi, F. Arrichiello, P. Robuffo Giordano, S. Chiaverini, and A. Franchi, "Adaptive Trajectory Tracking for Quadrotor MAVs in Presence of Parameter Uncertainties and External Disturbances," *IEEE Trans. on Control Systems Technology*, vol. 26, no. 1, pp. 248–254, 2018.
- [28] M. Zanon and S. Gros, "Safe reinforcement learning using robust mpc," *IEEE Trans. on Automatic Control*, vol. 66, no. 8, p. 3638–3652, 2021.
- [29] A. Tagliabue and J. P. How, "Efficient deep learning of robust policies from mpc using imitation and tube-guided data augmentation," *IEEE Transactions on Robotics*, vol. 40, p. 4301–4321, 2024.