Unifying Quadrotor Motion Planning and Control by Chaining Different Fidelity Models

Rudolf Reiter[1], Chao Qin[2], Leonard Bauersfeld[1] and Davide Scaramuzza[1]

***Abstract*—Many aerial tasks involving quadrotors demand both instant reactivity and long-horizon planning. High-fidelity models enable accurate control but are too slow for long horizons; low-fidelity planners scale but degrade closed-loop performance. We present UNIQUE, a unified MPC that cascades models of different fidelity within a single optimization: a short-horizon, high-fidelity model for accurate control, and a long-horizon, low-fidelity model for planning. We align costs across horizons, derive feasibility-preserving thrust and body-rate constraints for the point-mass model, and introduce transition constraints that match the different states, thrust-induced acceleration, and jerk–body-rate relations. To prevent local minima emerging from nonsmooth clutter, we propose a 3D progressive smoothing schedule that morphs norm-based obstacles along the horizon. In addition, we deploy parallel randomly initialized MPC solvers to discover lower-cost local minima on the long, low-fidelity horizon. In simulation and real flights, under equal computational budgets, UNIQUE improves closed-loop position or velocity tracking by up to 75% compared with standard MPC and hierarchical planner–tracker baselines. Ablations and Pareto analyses confirm robust gains across horizon variations, constraint approximations, and smoothing schedules.**

SUPPLEMENTARY MATERIAL

A summarizing video is available at: https://youtu.be/D7WKPK6ZQs0

I. INTRODUCTION

Fast and safe autonomous quadrotor flight hinges on two competing requirements: *reactivity* to disturbances and dynamic interactions, and *planning* to reason over long horizons for obstacle avoidance, efficiency, and mission objectives. Optimizing a high-fidelity quadrotor model over such horizons is typically intractable in real time. Yet, plans synthesized with low-fidelity models alone are suboptimal and sometimes incompatible with the true vehicle dynamics. This discrepancy is the root cause of the long-standing separation between planning and control for quadrotors [1]–[3].

Hierarchical architectures compute a long-horizon plan using a simplified model (e.g., point-mass/polynomial [1], [4]) and track it with an MPC. This decoupling induces (i) performance limitations imposed by the coarse plan, (ii) redundant optimization in the initial horizon, and (iii) conflicting inner/outer loops that may create infeasibilities or unstable behavior. Pure MPC alternatives retain the high-fidelity model but shorten the horizon to remain real-time feasible, relying on terminal costs/sets that are difficult to obtain for complex scenes involving moving obstacles or nonlinear constraints and are prone to local minima in nonsmooth environments.

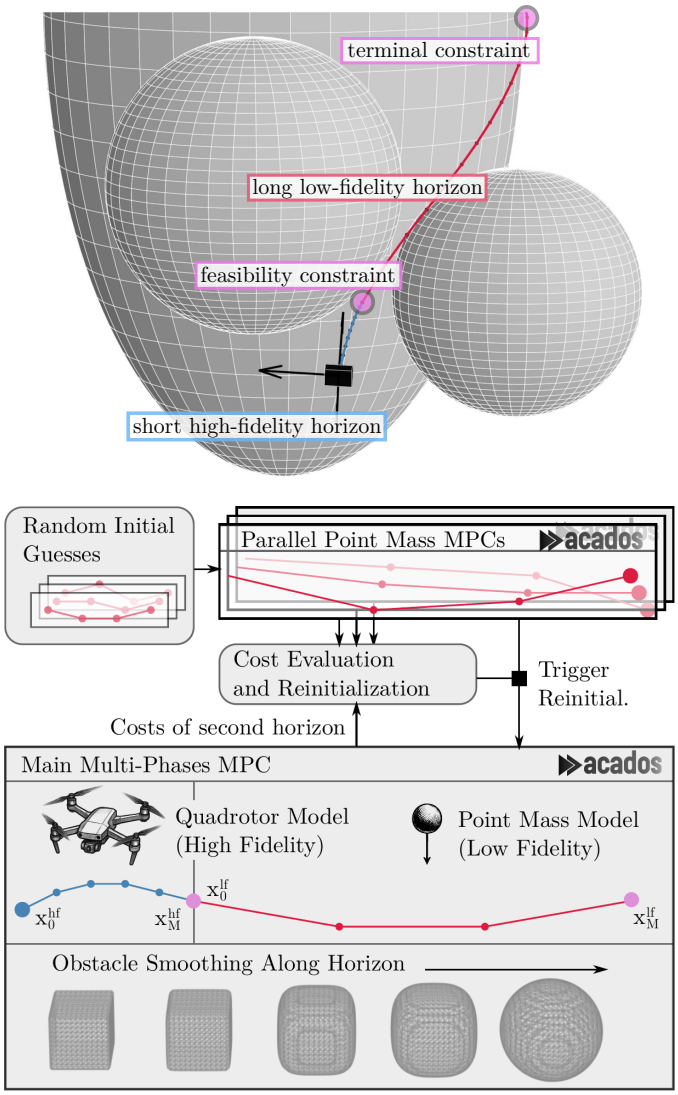

Fig. 1. Conceptual sketch of UNIQUE. We employ a multi-phase MPC with two horizons: one high-fidelity horizon, which is necessary for control, and a long horizon with a lower-fidelity point mass model, which enables long-horizon planning. Local minima of the gradient-based optimizer are avoided by (i) smoothing obstacle shapes along the horizon, and (ii) by running parallel MPCs solely for the computationally comparably cheap low-fidelity model on the second horizon. Whenever the cost on the low-fidelity planning horizon of such a parallel MPC is lower, the second horizon of the multi-phase MPC is reinitialized by the corresponding optimization variables.

We *unify* planning and control inside a single MPC by cascading models of different fidelity along the prediction horizon, as previously shown for quadrupeds [5]. As sketched in Fig. 1, a short high-fidelity phase governs the near-horizon,

[1]Authors are with the Robotics and Perception Group, University of Zurich, Switzerland (https://rpg.ifi.uzh.ch). This work was supported by the European Union's Horizon Europe Research and Innovation Programme under grant agreement No. 101120732 (AUTOASSESS) and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

[2] Institute for Aerospace Studies, University of Toronto, Canada

while a long low-fidelity phase provides foresight for global task performance. The two phases are coupled by transition constraints that align (i) position and velocity, (ii) thrust-induced acceleration, and (iii) jerk/body-rate relations. Cost terms are aligned across horizons, allowing the controller to optimize a consistent objective end-to-end. To improve numerical robustness over long horizons with sharp geometry, we adapt progressive smoothing [6] of norm-based obstacle models for 3D shapes, morphing cube-like sets to ellipsoids with increasing prediction time, which preserves recursive feasibility and enables real time iterations (RTIs) [31] to converge.

In addition to progressive smoothing, we propose an architecture of multiple parallel MPCs to escape local minima. The parallel MPCs solely optimize the computationally comparably cheap low-fidelity trajectory from different random initial guesses, all starting from the final state of the first high-fidelity horizon. Once the cost of one parallel MPC is lower than the cost along the second horizon of the multi-phase MPC, the decision variables are initialized with the lower-cost solution of the corresponding parallel MPC, cf. Fig. 1. Finally, our contributions are as follows:

- Unified two-phase MPC. A single optimization that couples a high-fidelity quadrotor model with a long-horizon point-mass model via equality constraints on position/velocity, thrust–acceleration, and jerk–body-rate mappings, with cost alignment across horizons.
- Feasibility-preserving low-fidelity constraints. Derivation and tractable inner approximations of thrust and body-rate safe sets for the point-mass model that ensure compatibility with high-fidelity actuator and rate limits.
- Progressive 3D obstacle smoothing. A horizon-dependent norm schedule for 3D convex obstacles that improves convergence in nonsmooth scenes while retaining safety, enabling reliable RTI with very long horizons.
- Parallel MPC framework. By optimizing different randomly initialized point-mass model MPCs along the second horizon, local minima are avoided in obstacle-rich environments.
- Comprehensive evaluation. Simulations and real flights showing substantial closed-loop gains over standard MPC and hierarchical baselines under equal compute budgets, including Pareto analyses across horizon variations and feasibility approximations, and demonstrations of long-horizon planning with sub-$5\,\mathrm{ms}$ iteration times.

The UNIQUE framework removes the planning–control performance gap without sacrificing real-time capability. In obstacle-rich tasks, we observe up to 75% lower closed-loop cost compared to standard MPC and hierarchical designs, while reducing online computation compared to long-horizon, high-fidelity MPC.

A. Outline

Following the related works Section II, the Section III formalizes the nominal high-fidelity MPC. Section IV introduces the point-mass planning model and derives feasibility sets. Section V presents the unified formulation, including cost/constraint alignment, transition constraints, progressive smoothing, the parallel MPC framework, and details numerical approximations for real-time efficiency. Sections VI and VII report simulation and real-world evaluations against standard MPC and hierarchical baselines, followed by ablations. We conclude with limitations and future directions in Section VIII.

B. Notation

For the quaternion rotation of a vector $x \in \mathbb{R}^3$ and a quaternion $q \in \mathbb{R}^4$, we use $q \odot x$ and for the quaternion multiplication with a second quaternion $p \in \mathbb{R}^4$, we use $q \cdot p$. A power p of a vector $x \in \mathbb{R}^n$ is taken element wise, with $x^p = [x_1^p, \dots, x_n^p]^\top$. When using different models along an MPC horizon we refer to the trajectory parts as *phases*.

II. RELATED WORK

Classical hierarchical architectures. In traditional software stacks, a high-level planner solves a simplified problem, often in discrete or low-dimensional state/control spaces, and provides a reference path or trajectory to a lower-level controller [4]. Graph search [7]–[9] and sampling-based methods such as rapidly exploring random trees (RRTs) [10] are common choices. For quadrotors, the authors in [1] employ a point-mass model for planning, while Model Predictive Contouring Control (MPCC) [2] and its extensions [3] apply the decoupled approach to racing by optimizing along a precomputed trajectory. While these planners are computationally efficient and capture high-level objectives, the reduced dimensionality leads to suboptimal trajectories, and the hierarchical decomposition introduces coupling issues between the planner and the controller.

Point-mass motion planning. Approximating the quadrotor as a low-fidelity integrator chain [1] and allowing small deviations from its true dynamics [11], [12] can significantly simplify the motion planning problem. In particular, Hehn and D'Andrea [4] proposed a third-order point-mass model for online planning under actuation constraints, while higher-order (snap) models [13] further improve trajectory smoothness and accuracy at the expense of increased computational complexity. Moreover, exploiting the differential flatness of simple quadrotor models enables long-horizon motion planning for complex maneuvers [14], [15] without significantly compromising model fidelity.

However, these approaches operate as open-loop motion planners and must be complemented by a feedback tracking controller. One key limitation arises from their inherently smooth lower-fidelity plans, unable to account for more complicated aerodynamic effects, necessitating conservatism. In addition, they often require the current jerk state as an initial condition, which is difficult to obtain in practice since it depends on the rate of change of thrust [13].

Nevertheless, these simplified models form the foundation of the long-horizon component in our unified MPC framework, where we address the limitations as mentioned above by introducing a multi-phase structure preceding the point-mass phase, enforcing feasibility-preserving constraints, and ensuring consistency with high-fidelity quadrotor dynamics.

Long-horizon MPC extensions. Several approaches aim to embed long-horizon reasoning directly into MPC. One line of work uses terminal costs learned from data. For instance, reinforcement learning is used to approximate value functions [16], [17] or supervised learning for convex quadratic surrogates [18]. If a policy is available, long-term behavior can be approximated by closed-loop costing [19], [20], including recent neural-policy rollouts [21]. However, these methods rely on pretrained policies or value functions, which may be environment-specific. Other research has studied recursive feasibility under multiple phases [22], developed multi-phase MPC formulations within modern solvers such as `acados` [23], and provided foundational analysis of multi-stage problems [24]. Li et al. [25] introduced a dual-model MPC for legged and aerial robots, although it was only tested in offline simulations and lacked recursive feasibility, rendering the formulation unsafe. Later, the authors extended the multi-model formulation for quadrupeds [5], enabling feasibility. Our work emphasizes real-time feasibility and numerical stability and demonstrates the approach in closed-loop quadrotor experiments.

III. Nominal Controller

In the following, we describe the nominal nonlinear MPC that serves as a basis for UNIQUE. First, the high-fidelity model is introduced, followed by the essential system constraints, obstacle constraints, and the MPC formulation.

A. High-Fidelity Control Model

The high-fidelity quadrotor is modeled following [26] as a 6-degree-of-freedom rigid body with mass m and a diagonal moment of inertia matrix $J = \mathrm{diag}(\begin{bmatrix} J_x & J_y & J_z \end{bmatrix})$.¡

The model includes the positions $p \in \mathbb{R}^3$, velocities in the world coordinate frame $v \in \mathbb{R}^3$, orientation formulated as quaternions $q \in S^3 \subset \mathbb{R}^4$, body rates $\omega \in \mathbb{R}^3$, single rotor thrusts $f_i^\top = \begin{bmatrix} 0 & 0 & f_{z,i} \end{bmatrix}$ with $i \in \{1,2,3,4\}$ and the collective z-components of the single rotor thrusts $f_z^\top = \begin{bmatrix} f_{z,1} & f_{z,2} & f_{z,3} & f_{z,4} \end{bmatrix} \in \mathbb{R}^4$. Thrusts are generated by the rotors with the rotor speed Ω_i and the geometric location $r_{\mathrm{p},i} \in \mathbb{R}^3$ of the i-th rotor in the body frame via $f_{z,i} = \Omega_i^2 c_l$ using the thrust coefficient $c_\mathrm{l} \in \mathbb{R}$.

The total thrust is the sum of all single rotor thrusts f_i, with $T(f_z) = \sum_{i=1}^4 f_i$, and controlled by the corresponding derivatives $u^{\mathrm{hf}} := \dot{f}_z \in \mathbb{R}^4$. When now computing the total force $f : \mathbb{R}^3 \times \mathbb{R}^4 \to \mathbb{R}^3$ acting on the drone in the body frame, we add residual force $f_{\mathrm{res}} : \mathbb{R}^4 \times \mathbb{R}^3 \to \mathbb{R}^3$ to the total thrust, with

$$f(v_\mathcal{B}, f_z) := \\ T(f_z) + f_{\mathrm{res}}(v_\mathcal{B}, \Omega^2) = T(f_z) + f_{\mathrm{res}}(v_\mathcal{B}, \frac{1}{c_l} f_z),$$

where the velocity $v_\mathcal{B} \in \mathbb{R}^3$ in the body frame $\mathcal{B}$ can be obtained via $v_\mathcal{B} = q \odot v$.

The residual force $f_{\mathrm{res}}(v_\mathcal{B}, \Omega^2)$ models aerodynamic effects and is obtained by fitting selected features of Ω^2 and velocities $v_\mathcal{B}$ to measured data. The literature proposes either higher order polynomials [27] or linear features of only the velocity $v_\mathcal{B}$ [14]. We use higher order polynomial for a high-fidelity residual model, cf. App. A.

In addition to translation forces, we compute rotational counter torques of the rotors via

$$\tau_{\mathrm{c},i}(f_{z,i}) = \kappa_i \begin{bmatrix} 0 & 0 & f_{z,i} \end{bmatrix}^\top.$$

In symmetric setups, the coefficients κ_i are equal in absolute value and have pairwise the same sign, i.e., $\kappa_1 = -\kappa_2 = \kappa_3 = -\kappa_4$. By adding the torque originating from the single rotor thrusts to the counter torques, the collective torque is

$$\tau(f_z) = \sum_{i=1}^4 \tau_{\mathrm{c},i}(f_{z_i}) + r_{\mathrm{p},i} \times \begin{bmatrix} 0 & 0 & f_{z,i} \end{bmatrix}^\top.$$

With the gravity vector $g^\top = \begin{bmatrix} 0 & 0 & 9.81\,\mathrm{m/s^2} \end{bmatrix}$ and the state the high-fidelity state

$$x = \begin{bmatrix} p^\top & q^\top & v^\top & \omega^\top & f_z^\top \end{bmatrix}^\top \in \mathbb{R}^{17},$$

the quadrotor dynamics are finally written as

$$\dot{x} = \begin{bmatrix} \dot{p} \\ \dot{q} \\ \dot{v} \\ \dot{\omega} \\ \dot{f}_z \end{bmatrix} = \begin{bmatrix} v \\ q \cdot \begin{bmatrix} 0 & \frac{\omega}{2} \end{bmatrix}^\top \\ \frac{1}{m}\Big(q \odot f\big(q^{-1} \odot v,\, f_z\big)\Big) - g \\ J^{-1}\big(\tau(f_z) - \omega \times J\omega\big) \\ u \end{bmatrix} =: f(x,u). \tag{1}$$

A simplified version of (1) uses linear residual forces, which was shown to be differentiable flat [14], cf. App. B.

B. Quadrotor System Constraints

Quadrotors are typically constrained by maximum and minimum rotor speeds, resulting in maximum and minimum single rotor thrusts $\overline{f_z}$ and $\underline{f_z}$. Most quadrotors are equipped with a low-level tracking controller that takes the body rates and collective thrust as inputs. Therefore, constraints are expressed as constrains $\overline{\omega}_{xy} \in \mathbb{R}$ on the body rate for roll and pith, and the total collective thrust $\overline{f}_{\mathrm{th}}$ and $\underline{f}_{\mathrm{th}}$, written as

$$\mathbb{S}_\omega := \big\{\omega \in \mathbb{R}^3 \mid \omega_x \le \overline{\omega}_{xy} \wedge \omega_y \le \overline{\omega}_{xy}\big\}, \tag{2a}$$

$$\mathbb{S}_\mathrm{f} := \big\{f_z \in \mathbb{R}^4 \mid \underline{f}_{\mathrm{th}} \le T(f_z) \le \overline{f}_{\mathrm{th}}\big\}. \tag{2b}$$

We constrain the quadrotor to an operating region with a maximum vertical speed of $\overline{v}_z \in \mathbb{R}$ and horizontal speed $\overline{v}_{xy} \in \mathbb{R}$, with

$$\mathbb{S}_v = \left\{v \in \mathbb{R}^3 \middle| \|v_z\| \le \overline{v}_z \wedge \|[v_x, v_y]\| \le \overline{v}_{xy}\right\}.$$

Assumption 1: We assume that the residual force $f_{\mathrm{res}}(v_\mathcal{B}, \Omega^2)$ is bounded by $\overline{f}_{\mathrm{res}} \in \mathbb{R}^3$ with $f_{\mathrm{res}}(v_\mathcal{B}, \Omega^2) \le \overline{f}_{\mathrm{res}}$ for all $v_\mathcal{B} \in \mathbb{S}_v$ and $c_l\Omega^2 \in \mathbb{S}_\Omega$.

C. Obstacle Avoidance Constraints

We consider convex obstacle shapes that can be described by p-norms, including ellipsoids and cubes. An obstacle is described by the tuple $\theta = (p_o, R_o, d_o, \alpha_0)$ with the obstacle center $p_o \in \mathbb{R}^3$, the orientation normal vectors $u_o, v_o, w_o \in \mathbb{R}^3$, with the rotation matrix $R_o = \begin{bmatrix} u_o & v_o & w_o \end{bmatrix}$, the obstacle

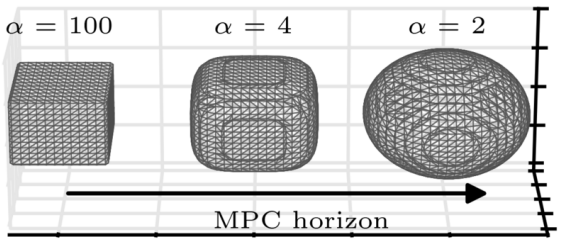

Fig. 2. Visualization of obstacles with varying smoothness parameters $\alpha = 100$, 8, and 2. The shape transitions from a cube to an ellipsoid. The arrow indicates the MPC horizon.

dimensions $d_o^\top = \begin{bmatrix} d_x & d_y & d_z \end{bmatrix}$, and the norm or shape parameter $\alpha_o \geq 2$. The transformation

$$\eta(p) = \begin{bmatrix} \eta_1(p) & \eta_2(p) & \eta_3(p) \end{bmatrix}^\top := \operatorname{diag}(d_o)^{-1}(R_o^\top p - p_0)$$

is used to obtain normalized obstacle coordinates $\eta(p)$ and describe the occupied space via the set

$$\mathcal{O} := \left\{ x \in \mathbb{R}^3 \middle| \left(\frac{1}{3} \sum_{i=1}^{3} |\eta_i(p)|^\alpha \right)^{\frac{1}{\alpha}} \leq 1 \right\}, \tag{3}$$

where a parameter $\alpha \geq 2$ defines the shape of the obstacle. The smoothness parameter $\alpha = 2$ corresponds to an ellipsoid and $\alpha \to \infty$ to a cuboid. Note that (3) is a scaled norm, where lower values of alpha are always over-approximations of larger alpha, i.e., $2 \leq \alpha_1 \leq \alpha_2 \Rightarrow \mathcal{O}(\alpha_2) \subseteq \mathcal{O}(\alpha_1)$, cf. [6].

D. Nominal MPC

With enough computational resources and given a sufficiently good initial guess, we would solve a nominal MPC problem, with the highest fidelity model available. We formulate an multiple-shooting nonlinear program (NLP) to solve the MPC problem with state variables $X = \begin{bmatrix} x_0 & \dots & x_M \end{bmatrix}$ and controls variables $U = \begin{bmatrix} u_0 & \dots & u_{M-1} \end{bmatrix}$, and use numerical integration to obtain the consecutive states x_{k+1} at discrete times kt_Δ via $x_{k+1} = F(x_k, u_k; t_\Delta)$, particularly an RK4 scheme with sampling time t_Δ. A terminal value function $\Phi(x)$ and a terminal set $h_M(x_M)$ approximate the infinite horizon and guarantee recursive feasibility. With state selection matrices $P_p \in \mathbb{R}^{3\times 17}$, $P_f \in \mathbb{R}^{4\times 17}$, $P_v \in \mathbb{R}^{3\times 17}$ and $P_\omega \in \mathbb{R}^{3\times 17}$ that select the states p, f, v and ω of the full state x, the nominal MPC problem is defined as

$$\min_{X,U} \quad \sum_{k=0}^{M-1} l(x_k, u_k) + \Phi(x_M) \tag{4a}$$

s.t.

$$x_0 = \hat{x}_0, \tag{4b}$$

$$x_{k+1} = F(x_k, u_k; t_\Delta), \qquad k \in \mathbb{N}_{[0,M-1]}, \tag{4c}$$

$$P_\omega x_k \in \mathbb{S}_\omega,\ P_f x_k \in \mathbb{S}_f,\ P_v x_k \in \mathbb{S}_v \quad k \in \mathbb{N}_{[0,M-1]}, \tag{4d}$$

$$P_p x_k \notin \mathbb{O}, \qquad k \in \mathbb{N}_{[0,M-1]}, \tag{4e}$$

$$0 \geq h_M(x_M). \tag{4f}$$

We define the nonlinear least-squares cost

$$l(x, u) = t_\Delta (y(x) - \tilde{y})^\top \operatorname{diag}(w)(y(x) - \tilde{y}), \tag{5}$$

with the tracking states $y^\top(x) = \begin{bmatrix} p & \phi(q, \tilde{q}) & v & \omega & f_z & u \end{bmatrix} \in \mathbb{R}^{20}$ that includes the quaternion error $\phi : \mathbb{R}^{4\times 4} \to \mathbb{R}^3$ transformed to Euler angles, the reference $\tilde{y} \in \mathbb{R}^{20}$ and the weight vector $w \in \mathbb{R}^{20}$.

IV. Low-Fidelity Model

For long-horizon trajectory planning, low-fidelity models are inevitable due to the necessary trade-off between accuracy and computation time per planner iteration. Due to the differential flatness of the quadrotor motion, a chain of integrators is often used to approximately describe the quadrotor motion [4], [13], [25]. In this work, we use a chain of three integrators to model the trajectory of a quadrotor, as proposed in [4] for rapid planning of feasible trajectories. Three integrators combined with specific constraints are necessary to generate feasible trajectories for quadrotors [4]. Due to readability, we overload the notation for positions p and velocities v as for the quadrotor model, and add acceleration states $a \in \mathbb{R}^3$ and jerk inputs $u^{\text{lf}} := j \in \mathbb{R}^3$. to obtain the low-fidelity model

$$\dot{x}^{\text{lf}} = \begin{bmatrix} \dot{p} & \dot{v} & \dot{a} \end{bmatrix}^\top = \begin{bmatrix} v & a & j \end{bmatrix}^\top := f^{\text{lf}}(x^{\text{lf}}, u^{\text{lf}}), \tag{6}$$

with the low-fidelity state $x^{\text{lf}} = \begin{bmatrix} p & v & a \end{bmatrix}^\top$.

A. Relation to High-Fidelity Model

Due to the differentiable flatness of the quadrotor model, cf. App. B, all quadrotor states of a simple model with linear residuals [14] can be obtained from derivatives of the position and the orientation of the quadrotor model. The point-mass model, which is a chain of integrators, can be seen as a lower fidelity model from which we can derive the most important states for motion planning. We are particularly interested in the quadrotor position p, velocity v, body rate ω, and thrust T. These states are typically associated with the optimization objective and constraints related to physical feasibility and safety, as in (2a), (2b) and (3). A point-mass model formulates derivatives of the position, in our case, up to the jerk. These physical quantities are aligned with the quadrotor motion. From the quadrotor model (1), the acceleration can be directly obtained via

$$a_{\text{quad}}(q, v, f_z) = \frac{1}{m}\Big(q \odot f\big(q^{-1} \odot v,\ f_z\big)\Big) - g$$

and the jerk via the derivative of the acceleration

$$j_{\text{quad}}(q, v, \omega, f_z, \dot{f}_z) = \dot{a}_{\text{quad}}(q, v, f_z),$$

with details in App. C. Without the consideration of residual forces, i.e., assuming $f_{\text{res}} = 0$, the quadrotor acceleration simplifies to

$$\hat{a}_{\text{quad}}(q, f_z) = \frac{1}{m}\Big(q \odot T(f_z)\Big) - g,$$

with the total thrust $T(f_z) = \|T(f_z)\| = e_z^\top T(f_z)$, for $T \geq 0$. Taking the derivative of each single rotor thrust leads to $\dot{T}(f_z) = T(\dot{f}_z) = \|T(\dot{f}_z)\|$ and the jerk follows as

$$\hat{j}_{\text{quad}}(q, \omega, f_z, \dot{f}_z) = q \odot \begin{bmatrix} T(f_z)\omega_y \\ -T(f_z)\omega_x \\ T(\dot{f}_z) \end{bmatrix}.$$

B. Point-Mass System Constraints

The authors in [4] derived constraints

$$m\|a+g\| \leq \overline{f}_{\text{th}}$$

on a third-order point mass model that guarantees feasible trajectories for a quadrotor model with constraints defined in (2) but without residual forces. We extend these constraints to include feasibility for the residual model by utilizing the upper bound $\overline{f}_{\text{res}}$ from Ass. 1. While [4] required decoupled constraints per axis due to their computationally efficient planning algorithm, we utilize nonlinear programming and successive linearization, which allows for coupled, less restrictive constraints. Particularly, the constraint $\mathbb{S}_{\text{f}}$ on the collective thrust (2b) can be formulated with the individual acceleration components a_x, a_y and a_z per axis in the world frame for the low-fidelity planning model. The collective constraint for the low-fidelity model is therefore

$$\mathbb{S}_{\text{f}}^{\text{lf}} := \left\{a \in \mathbb{R}^3 \mid m\|a+g\| \leq \overline{f}_{\text{th}} - \overline{f}_{\text{res}}\right\}. \tag{7}$$

Body rate constraints $\mathbb{S}_{\omega}$ formulated in (2b) for the high-fidelity model can be incorporated in the low-fidelity point mass model via constraints on the time derivative of the unit vector of the gravity shifted acceleration, i.e.,

$$\left\| \frac{d}{dt}\left(\frac{a+g}{\|a+g\|}\right)\right\| \leq \overline{\omega}_{xy}. \tag{8}$$

Resolving the time derivative in (8) and utilizing the explicit notation of jerk with $j = \dot{a} + \dot{g} = \dot{a}$, the body rate feasibility constraint for the low-fidelity model can be written as

$$\mathbb{S}_{\omega}^{\text{lf}} := \left\{ \begin{bmatrix} a \\ j \end{bmatrix} \in \mathbb{R}^6 \middle| \frac{\|j\|}{\|a+g\|} \leq \overline{\omega}_{xy} \right\}. \tag{9}$$

The auxiliary constraints on the thrust (7) and the body rate (9) are highly nonlinear. However, they can be conservatively reformulated into more optimization-friendly subsets, as shown in the following, with details in [4].

First, we simplify the lower bound of the thrust constraint $\mathbb{S}_{\text{f}}^{\text{lf}}$ defined in (7). The lower bound of (7) is satisfied if the simpler inequality

$$a_z \geq \underline{a}_z \geq \frac{1}{m}\underline{f}_{\text{th}} - g,$$

holds, where $\underline{a}_z$ is an auxiliary, more restrictive, bound that is also used in the following to shape the simplified constraints. We refer to [4] for a proof.

Second, the upper bound of (7) can be simplified, where a more restrictive constraint guarantees the feasibility of (7). The simplified constraint involves two design parameters $\alpha_x, \alpha_y \in (0,1)$ that shape and trade off the more optimization-friendly constraints. By utilizing the design parameters, box-constraints for $\overline{a}^\top = \begin{bmatrix}\overline{a}_x & \overline{a}_y & \overline{a}_z\end{bmatrix}$ can be computed iteratively by

$$\begin{aligned}
\overline{a}_z &:= \alpha_z \left(\frac{1}{m}\overline{f}_{\text{th}} - g\right), \\
\overline{a}_x &:= \alpha_x \sqrt{\left(\frac{1}{m}\overline{f}_{\text{th}}\right)^2 - (\overline{a}_z + g)^2}, \\
\overline{a}_y &:= \sqrt{\left(\frac{1}{m}\overline{f}_{\text{th}}\right)^2 - \overline{a}_x^2 - (\overline{a}_z + g)^2}.
\end{aligned}$$

The simplified box constraints, expressed as a set constraining the acceleration states of the low-fidelity model, can finally be written as

$$\hat{\mathbb{S}}_{\text{f}}^{\text{lf}} := \left\{a \in \mathbb{R}^3 \mid \begin{bmatrix}-\overline{a}_x & -\overline{a}_y & \underline{a}_z\end{bmatrix} \leq a+g \leq \overline{a}\right\} \subset \mathbb{S}_{\text{f}}^{\text{lf}}, \tag{10}$$

which is subset of the thrust constraints (7).

Next, we simplify the constraints of $\mathbb{S}_{\omega}^{\text{lf}}$, defined in (9). A more optimization-friendly subset of (9) can be found by exploiting the lower bound on the acceleration

$$\|a+g\| \geq \underline{a}_z$$

to obtain simplified convex constraints on the jerk j

$$\hat{\mathbb{S}}_{\omega}^{\text{lf}} := \left\{j \in \mathbb{R}^3 \middle| \|j\| \leq (\underline{a}_z + g)\overline{\omega}_{xy}\right\} \subseteq \mathbb{S}_{\omega}^{\text{lf}}. \tag{11}$$

Arbitrary box constraints for each individual component of the jerk can be computed, forming a subset of (11) and, therefore, decouple the constraints. A valid choice would be $\overline{j}_{xyz} = \frac{\underline{a}_z + g}{\sqrt{3}}\overline{\omega}_{xy}$ to obtain the convex decoupled box constraints

$$\widehat{\mathbb{S}}_{\omega}^{\text{lf}} := \left\{j \in \mathbb{R}^3 \middle| |j_{\{x,y,z\}}| \leq \overline{j}_{xyz}\right\} \subseteq \hat{\mathbb{S}}_{\omega}^{\text{lf}}. \tag{12}$$

V. UNIQUE: THE UNIFIED MPC FORMULATION

In the following, we propose the UNIQUE, a multi-phase MPC framework that combines the high-fidelity MPC from Sect. III-D for short horizons and utilizes the low-fidelity point mass model from Sect. IV for long-horizon planning.-horizon planning. The merging of two different models along the MPC horizons requires (i) the alignment of the cost function, (ii) recursive feasible constraint formulations, (iii) transition function between the models. Very long prediction horizons result in MPC optimization problems with ubiquitous local minima, which may degrade the performance of the multi-phase MPC, particularly in cluttered environments with nonsmooth obstacle shapes. Two proposed strategies effectively improve the performance as part of the UNIQUE framework, i.e., progressive smoothing and parallel low-fidelity optimization from random initial seeds. These strategies and the requirements for merging the two models along the horizon are detailed in the following.

A. Cost Alignment

The high fidelity weights

$$w^\top = [w_p^\top, w_\phi^\top, w_v^\top, w_\omega^\top, w_f^\top, w_u^\top]$$

of the nonlinear-least square cost (5) are split into groups for physical states, with $w_p \in \mathbb{R}^3$ for positions, $w_v \in \mathbb{R}^3$ for velocities, $w_\phi \in \mathbb{R}^3$ for Euler angles, $w_\omega \in \mathbb{R}^3$ for body rates, $w_f \in \mathbb{R}^4$ for single rotor thrusts and $w_u \in \mathbb{R}^4$ for its derivatives. The cost for the low-fidelity model is formulated as a linear least squares cost for the full low-fidelity state via

$$l^{\text{lf}}(x^{\text{lf}}, u^{\text{lf}}) = t_\Delta^{\text{lf}}(y^{\text{lf}} - \bar{y}^{\text{lf}})^\top \text{diag}(w^{\text{lf}})(y^{\text{lf}} - \bar{y}^{\text{lf}}), \tag{13}$$

with $y^{\text{lf}} = [x^{\text{lf}}, u^{\text{lf}}]^\top$ and the weights

$$w^{\text{lf}} = \begin{bmatrix} w_p^\top & w_v^\top & w_a^\top & w_j^\top \end{bmatrix}^\top,$$

that are partially aligned with the high-fidelity weights. Particularly, the weights w_p and w_v are related to the same physical states and therefore chosen identically.

In the high fidelity cost (5), single rotor vertical thrusts are penalized by a sum of squares $w_f f_z^\top f_z$. In the low-level model, we aim to penalize the equivalent squared sum of forces

$$w_a a^\top a \stackrel{!}{=} w_f f_z^\top f_z = w_f m a^\top m a = w_f m^2 a^\top a,$$

which we obtain from the accelerations. To align the low-level and high-level weights, we obtain the acceleration weight w_a from the high-level weight w_f via $w_a := m^2 w_f$.

The jerk weight w_j is related to the body rate weight w_ω and the force derivative input weight w_f. Similar to the acceleration weight, we take the input weight relation $w_j \approxeq m^2 w_u$ as an approximation and increase it in closed-loop experiments to account for the body rate weight.

For the terminal value function $\Phi^{\mathrm{lf}} : \mathbb{R}^9 \to \mathbb{R}$, we define a least-square penalty on the final reference position $p_T^{\mathrm{lf}} \in \mathbb{R}^3$ and the weight $q_T \in \mathbb{R}^3$, i.e.,

$$\Phi^{\mathrm{lf}}(x^{\mathrm{lf}}) := (P_p x^{\mathrm{lf}} - p_T^{\mathrm{lf}})^\top \mathrm{diag}(q_T)(P_p x^{\mathrm{lf}} - p_T^{\mathrm{lf}})$$

and a terminal safe set $h^{\mathrm{lf}}(x^{\mathrm{lf}})$. With the simple assumption of considering hovering as a safe state, we can use the corresponding low-level states $P_{va} x^{\mathrm{lf}}$, with the projection matrix $P_{va} \in \mathbb{R}^{6\times 9}$ selecting velocities and accelerations, to establish a simple safe terminal constraint $h^{\mathrm{lf}}(x^{\mathrm{lf}}) = P_{va} x^{\mathrm{lf}} = 0$.

B. Constraint Alignment

As shown in Sect. IV-B, we can find constraint sets $\mathbb{S}_\omega^{\mathrm{lf}}$ and $\mathbb{S}_f^{\mathrm{lf}}$ for the low-fidelity model that guarantee a feasible trajectory of the high-fidelity model that is constrained via the feasible sets $\mathbb{S}_\omega$ and $\mathbb{S}_f$. We propose two subset variants of the low-fidelity set $\mathbb{S}_\omega^{\mathrm{lf}}$, which are tighter constrained but numerically easier to solve. The set $\hat{\mathbb{S}}_\omega^{\mathrm{lf}}$ assumes a lower bound $\underline{a}_z$, leading a to an convex quadratic constraint on the jerk j (11) and the set $\widehat{\mathbb{S}}_\omega^{\mathrm{lf}}$ decouples the jerk components to box constraints (12). The constraint alignment guarantees that the plan in the second horizon is a feasible plan for the higher-fidelity drone model.

C. Transition Function

The conditions on the low-fidelity model of the previous Sect. V-B provide necessary conditions on the position trajectory and its derivatives of the high-fidelity model for feasibility. In this section, we develop necessary conditions for the final state of the high-fidelity model x_M, such that the low-fidelity plan for future times $t \geq M t_\Delta$ can be tracked. Our approach explicitly couples all derivatives of the low-fidelity trajectory up to the jerk with the corresponding high-fidelity states. First, we require the state and velocities of both fidelity models to be aligned via

$$\pi_{pv}^\top \left(x, x^{\mathrm{lf}}\right) := \begin{bmatrix} P_p x & P_v x \end{bmatrix} - \begin{bmatrix} P_p^{\mathrm{lf}} x^{\mathrm{lf}} & P_v^{\mathrm{lf}} x^{\mathrm{lf}} \end{bmatrix} = 0. \quad (14)$$

Next, we require the acceleration obtained from the rotor thrusts to be aligned with the acceleration of the low-fidelity model, via

$$\pi_a \left(x, x^{\mathrm{lf}}\right) := \frac{1}{m} \left(P_q x\right) \odot T \left(P_f x\right) - g - P_a x^{\mathrm{lf}} = 0. \quad (15)$$

Finally, the body rates need to be aligned with the jerk and the acceleration. According to the derivation in [4], the body rates of the quadrotor model can be obtained from the acceleration a and the jerk $j = u^{\mathrm{lf}}$ via

$$\begin{bmatrix} \omega_x & -\omega_y & 0 \end{bmatrix}^\top = q \odot \left(\frac{j}{\|a\|} - \frac{a a^\top j}{\|a\|^3} \right) \quad (16)$$

which leads to the constraint

$$\begin{aligned} &\pi_\omega \left(x, x^{\mathrm{lf}}, u^{\mathrm{lf}}\right) := \\ &\quad \mathrm{diag}\left([1, -1, 0]^\top\right) P_\omega x - \\ &\quad P_q x \odot \left(\frac{u^{\mathrm{lf}}}{\|P_a x^{\mathrm{lf}}\|} - \frac{P_a x^{\mathrm{lf}} (P_a x^{\mathrm{lf}})^\top u^{\mathrm{lf}}}{\|P_a x^{\mathrm{lf}}\|^3} \right) = 0. \end{aligned} \quad (17)$$

We require the acceleration a in (16) and (17), respectively, to be nonzero at the transition. Fortunately, we already defined a minimum acceleration for the low-fidelity model in (10), which mitigates this numerical problem. Additionally, the jerk related constraint (17) can be interpreted as a constraint to the preceding jerk control before the initial transition state x_0^{lf} thus it is not required for feasibility of the quadrotor model. Finally, we summarize the relevant coupling conditions of (14) and (15) by

$$\pi \left(x, x^{\mathrm{lf}}\right) := \begin{bmatrix} \pi_{pv} \left(x, x^{\mathrm{lf}}\right) & \pi_a \left(x, x^{\mathrm{lf}}\right) \end{bmatrix}^\top = 0.$$

D. Progressive Smoothing

Within UNIQUE, we aim to plan long-horizon trajectories via numerical optimization, which is inherently prone to getting stuck in local minima, particularly when encountering non-smooth obstacles, such as cubes. To improve the convergence, we adapt the progressive smoothing method proposed by [6] for autonomous driving to the three-dimensional planning domain. In progressive smoothing, potentially non-smooth obstacle shapes $\mathcal{O}(\alpha)$ from (3), are progressively smoothened along the planning horizon of the MPC by scheduling the smoothing parameter α towards a smooth 2-norm. Particularly, a monotonously decreasing scheduling function $\lambda(t)$ with $\lambda : \mathbb{R}^+ \to \mathbb{R}^+$ depending on the MPC prediction time t is used, with $\lambda(0) = \alpha_0$ and $\lambda(t_f) = 2$, where t_f is the total prediction horizon and α_0 describes the actual obstacle shape. The smoothing along the horizon allows the application of the RTI scheme [28], and recursive feasibility for the obstacle constraints is guaranteed [6].

E. Multi-Phase MPC formulation

The multi-phase formulation optimizes the short-horizon high-fidelity states (X, U) together with a long-horizon low-fidelity states $(X^{\mathrm{lf}}, U^{\mathrm{lf}})$ in the single NLP:

Main Multi-Phases MPC

$$\min_{\substack{X,U,\\ X^{\mathrm{lf}},U^{\mathrm{lf}}}} \quad \sum_{k=0}^{M-1} l(x_k,u_k) + \sum_{k=0}^{N-1} l^{\mathrm{lf}}(x_k^{\mathrm{lf}},u_k^{\mathrm{lf}}) + \Phi^{\mathrm{lf}}(x_N^{\mathrm{lf}}) \tag{18a}$$

s.t.

$$x_0 = \hat{x}_0, \tag{18b}$$
$$0 = \pi\left(x_M, x_0^{\mathrm{lf}}\right), \tag{18c}$$
$$x_{k+1} = F\left(x_k,u_k;t_\Delta\right), \qquad k\in\mathbb{N}_{[0,M-1]}, \tag{18d}$$
$$x_{k+1}^{\mathrm{lf}} = F^{\mathrm{lf}}\left(x_k^{\mathrm{lf}},u_k^{\mathrm{lf}};t_\Delta^{\mathrm{lf}}\right), \qquad k\in\mathbb{N}_{[0,N-1]}, \tag{18e}$$
$$P_f x_k \in \mathbb{S}_f, P_\omega x_k \in \mathbb{S}_\omega, P_v x_k \in \mathbb{S}_v \quad k\in\mathbb{N}_{[0,M-1]}, \tag{18f}$$
$$P_{a,j} x_k^{\mathrm{lf}} \in \mathbb{S}_f^{\mathrm{lf}},\ P_j x_k^{\mathrm{lf}} \in \mathbb{S}_\omega^{\mathrm{lf}}, \qquad k\in\mathbb{N}_{[0,N-1]}, \tag{18g}$$
$$P_p x_k \notin \mathbb{O}\left(\lambda(kt_\Delta)\right), \qquad k\in\mathbb{N}_{[0,M-1]}, \tag{18h}$$
$$P_p^{\mathrm{lf}} x_k^{\mathrm{lf}} \notin \mathbb{O}\left(\lambda(Mt_\Delta + kt_\Delta^{\mathrm{lf}})\right), \qquad k\in\mathbb{N}_{[0,N]}, \tag{18i}$$
$$0 = h^{\mathrm{lf}}\left(x_N^{\mathrm{lf}}\right). \tag{18j}$$

The cost function accumulates stage costs from both phases and a terminal cost on the long horizon. System dynamics are enforced by the respective discrete time integration functions F and F^{lf}, and the transition constraint $\pi(x_M, x_0^{\mathrm{lf}}) = 0$ guarantees consistency of position, velocity, and thrust-induced acceleration. Actuator and body-rate feasibility is maintained through the sets $\mathbb{S}_f, \mathbb{S}_\omega$ for the high-fidelity model and their inner-approximated counterparts $\mathbb{S}_f^{\mathrm{lf}}, \mathbb{S}_\omega^{\mathrm{lf}}$ for the low-fidelity model. Both horizons are kept collision-free by enforcing position constraints against the obstacle set $\mathbb{O}(\lambda(\cdot))$, whose recursive feasibility is guaranteed [6]. Finally, the terminal equality $h^{\mathrm{lf}}(x_N^{\mathrm{lf}}) = 0$ enforces a simple safe set at the goal. The influence of the terminal safe set would be much larger when applied directly after the high-fidelity horizon.

F. Parallel Low-Fidelity MPCs

Above, we introduced the multi-phase formulation with progressive smoothing to plan for long horizons. In cluttered environments, solving the nonlinear program (18) is still prone to getting stuck in local minima due to the long planning horizon and the gradient-based local optimization of MPC solvers [29]. The MPC decision variables related to long-horizon planning are part of the second phase subproblem. Formulating this subproblem as a separate NLP allows to optimize parallel instances starting from different initial guesses and escaping the local minima of the dominant second planning horizon, and sacrificing only minor computational resources. Starting from the same initial point-mass state $\hat{x}_0^{\mathrm{lf}}$ obtained from the multi-phase MPC (18), the parallel MPC subproblems are

Algorithm 1 UNIQUE: Unified Multi-Fidelity MPC

Inputs: HF state $\hat{x}_0$, reference $\tilde{y}(t)$, obstacles $\{\mathcal{O}_i\}$, smoothing schedule $\lambda(t)$, number of parallel MPC restarts S, number of RTI P before reinitializing
Set $\alpha_i = \lambda(it_\Delta)$ for $i \in I_0$
Repeat at each control step k:
  1. Obstacles update:
    Set obstacle param. θ_i for prediction steps $i \in I_k$
  2. Main multi-phase MPC:
    Solve QP of multi-phase MPC at $x_k = \hat{x}_k$
    Evaluate second horizon cost: $J_{\mathrm{lf},k}$
  3. Parallel Point-Mass MPCs:
    Get optimal state $\hat{x}_0^{\mathrm{lf}} \leftarrow x_0^{\mathrm{lf}}$ from multi-phase MPC
    Solve $j = 1:S$ point-mass MPCs at $(x_0^{\mathrm{lf}})_j := \hat{x}_0^{\mathrm{lf}}$
    Evaluate parallel point-mass MPC costs: $(J_{\mathrm{lf},k})_j$
    if $(J_{\mathrm{lf},k})_j < J_{\mathrm{lf},k}$ **then**
      Reinitialize multi-phase MPC
        $(x_i^{\mathrm{lf}})_0 \leftarrow (x_i^{\mathrm{lf}})_j \ \forall i \in \{0,\dots,M+1\}$
    if $k \mod P = 0$ **then**
      Randomly reinitialize parallel point-mass MPCs
  4. Apply and shift:
    Execute first main multi-phase control u_0
    Shift primal variables of all MPCs

Parallel Point-Mass MPC

$$\begin{aligned} J_{\mathrm{lf}} = \min_{X^{\mathrm{lf}},U^{\mathrm{lf}}} \quad & \sum_{k=0}^{N-1} l^{\mathrm{lf}}(x_k^{\mathrm{lf}},u_k^{\mathrm{lf}}) + \Phi^{\mathrm{lf}}(x_N^{\mathrm{lf}}) \\ \text{s.t.} \quad & \\ & x_0^{\mathrm{lf}} = \hat{x}_0^{\mathrm{lf}}, \\ & x_{k+1}^{\mathrm{lf}} = F^{\mathrm{lf}}\left(x_k^{\mathrm{lf}},u_k^{\mathrm{lf}};t_\Delta^{\mathrm{lf}}\right), \qquad k\in\mathbb{N}_{[0,N-1]}, \\ & P_{a,j} x_k^{\mathrm{lf}} \in \mathbb{S}_f^{\mathrm{lf}},\ P_j x_k^{\mathrm{lf}} \in \mathbb{S}_\omega^{\mathrm{lf}}, \qquad k\in\mathbb{N}_{[0,N-1]}, \\ & P_p^{\mathrm{lf}} x_k^{\mathrm{lf}} \notin \mathbb{O}\left(\lambda(Mt_\Delta + kt_\Delta^{\mathrm{lf}})\right), \ k\in\mathbb{N}_{[0,N]}, \\ & 0 = h^{\mathrm{lf}}\left(x_N^{\mathrm{lf}}\right). \end{aligned} \tag{19}$$

Due to the feasibility of the initial state from (18), all MPC subproblems provide recursive feasibility when combined with the high-fidelity horizon of (18). Once any of the MPC subproblems initialized with random primal variables $({}^0X^{\mathrm{lf}}, {}^0U^{\mathrm{lf}})^i$ exhibits a feasible lower cost $J_{\mathrm{lf}}\left(({}^0X^{\mathrm{lf}}, {}^0U^{\mathrm{lf}})^i, \hat{x}_0^{\mathrm{lf}}\right)$ than the second horizon of (18), the second horizon of (18) is reinitialized with the optimal lower cost decision variables of i-th parallel MPC subproblem. Therefore, the cost improvement is guaranteed with the initialization strategy. See [16] for details on parallel initialization and evaluation strategies.

G. Algorithm and Numerical Efficiency

The complete UNIQUE algorithm is stated in Alg. 1. We use the notation $I_k = k, k+1, \dots, N+M+1$ for the prediction indexes of the MPC. In order to solve (18) numerically efficiently, we use the solver ACADOS [29] with the interior-point quadratic program (QP) solver HPIPM [30], deploy the RTI scheme [31], and use slack variables penalized by exact L1 and L2 constraints [32] for obstacle avoidance constraints

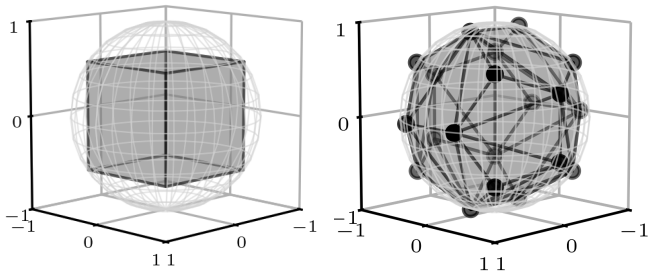

Fig. 3. Visualization of inner norm approximation. An inscribed cube covers less volume than the dodecahedron with 12 flat faces.

and approximate convex quadratic constraints of the sets $\hat{\mathbb{S}}_{\mathrm{f}}^{\mathrm{lf}}$ and $\hat{\mathbb{S}}_{\omega}^{\mathrm{lf}}$ via a polyhedron $\{x \in \mathbb{R}^3 | b_{\mathrm{p}} \leq A_{\mathrm{p}} x\}$, with $b_{\mathrm{p}} \in \mathbb{R}^{12}$ and $A \in \mathbb{R}^{3\times 12}$. Particularly, we use a dodecahedron with 12 surfaces compared to the six surfaces of an approximation via a cube as in $\widehat{\mathbb{S}}_{\omega}^{\mathrm{lf}}$ and $\hat{\mathbb{S}}_{\mathrm{f}}^{\mathrm{lf}}$. to lower the conservatism of the inner approximation by adding more constraints, as shown in Fig. 10. The approximation via a polyhedron allows the QP solver to safely consider the norm constraints, which is not possible when linearizing the constraints within one QP subroutine of sequential quadratic programming (SQP). Notably, the velocity constraint $\mathbb{S}_v$ spans the whole practical operating region of the quadrotor and can therefore be removed in (18).

VI. Experimental Setup

The proposed algorithm is evaluated in a real-world and a simulation environment with a realistic model that considers aerodynamic effects, as detailed in Appendix A. For both real-world experiments and simulations, we utilize the high-end *Offboard* drone. The specifications are summarized in Tab. I. The drone is controlled from a ground station via the low-latency Crossfire protocol (CRSF). The overall latency of the pipeline, from state estimate to control command, is typically around 20 ms, including all processing and transmission delays with CRSF. The drone is equipped with the `Betaflight` [34] lowest-level onboard controller, which tracks input body rates and total thrust. Therefore, the total thrust obtained from the states, i.e., the thrust derivatives, is used as the reference for the lowest-level controller. The flight experiments are conducted in a large industrial hall equipped with a motion-capture system providing state-estimation with millimeter accuracy at 400 Hz across a volume spanning $25 \times 12 \times 4\,\mathrm{m}$. In the final evaluations, we employ two distinct experimental settings: the *constant velocity environment* and the *tracking environment*, which are described below. In both experiments, the obstacles can be smooth ellipses or non-smooth cuboids.

Constant Velocity Environment. In this simulation environment, we evaluate the control performance, including eight randomly moving dynamic obstacles, when aiming to fly at a constant speed of 15 $\frac{\mathrm{m}}{\mathrm{s}}$ at a specific height. The environment evaluates how early adaptation towards long-horizon prediction and short-horizon adaptation due to moving obstacles affects the closed-loop performance of maintaining

TABLE I
Specifications of the "Offboard" Quadrotor

Parameter		Value	Unit
Geometry	$[r_{p,i}]_x$	$\pm$ 7.5	cm
	$[r_{p,i}]_y$	$\pm$ 10.0	cm
Mass	m	0.6	kg
Inertia	J_x	2.4	gm^2
	J_y	1.8	gm^2
	J_z	3.8	gm^2
Motor Const.	c_l	1.6	$\mu\mathrm{Ns}^2$
	κ	11.0	s
Quad. Diag.	l	38.2	cm
Max Thrust	$\overline{f}_{\mathrm{th}}$	34.0	N
Max Rate	$\overline{\omega}_{xy}$	10.0	s^{-1}
	$\overline{\omega}_z$	6.0	s^{-1}

a reference speed. More details on the environment are given in App. D.

Reference Tracking Environment. In contrast to the previous environment, where the path is of minor importance, the reference-tracking environment requires the drone to track periodic trajectories that are occluded by varying numbers of static objects, with obstacle counts ranging from 3 to 34. The path of the trajectory resembles shapes such as a triangle, a rounded rectangle, a figure eight, a butterfly, or sinusoidal waves. The tracking environment cost emphasizes the position tracking of the reference, and due to its periodic flight structure, it can be readily evaluated in real-world experiments. More details on the environment are given in App. E.

Both environments are evaluated via the quadratic closed-loop cost using the diagonal matrix weight $Q^{\mathrm{eval}} \in \mathbb{R}^{n_x \times n_x}$ and weights $R^{\mathrm{eval}} \in \mathbb{R}^{n_u \times n_u}$ to promote smooth flight. For the total simulation time of $t_{\mathrm{sim}} = 24\,\mathrm{s}$ with a sampling time of $t_{\Delta,\mathrm{sim}} = 40\,\mathrm{ms}$ we evaluate the closed-loop cost for N_{sim} steps with the simulated states $\hat{X} = [\hat{x}_0, \ldots, \hat{x}_{N_{\mathrm{sim}}-1}]$ and controls $\hat{U} = [\hat{u}_0, \ldots, \hat{u}_{N_{\mathrm{sim}}-1}]$ via

$$J^{\mathrm{eval}}(\hat{X}, \hat{U}) := t_{\Delta,\mathrm{sim}} \sum_{k=0}^{N_{\mathrm{sim}}-1} \|\hat{x}_k - \tilde{x}_k\|_{2,Q^{\mathrm{eval}}} + \|\hat{u}_k - \tilde{u}_k\|_{2,R^{\mathrm{eval}}}. \tag{20}$$

For the benchmark comparison, we normalize the accumulated cost via $\bar{J}^{\mathrm{eval}}(\hat{X}, \hat{U}) = J^{\mathrm{eval}}(\hat{X}, \hat{U}) / J^{\mathrm{eval}}(\hat{X}^{\mathrm{UNIQUE}}, \hat{U}^{\mathrm{UNIQUE}})$ by the cost from UNIQUE, and similarly we normalize the computation time t_{comp} taken by the MPC solver.

VII. Evaluation

We evaluate and ablate our approach by comparing it against two benchmarks: the *hierarchical approach* [1], [4] and the high-fidelity *standard MPC* formulation, which uses a longer horizon without utilizing the point-mass model. The hierarchical formulation first plans a trajectory using only a point-mass model, formulating it as an MPC problem, and then employs a high-fidelity MPC tracking controller. The planner updates the trajectory every tenth lower-level tracking iteration. Remarkably, all parameters of the point mass

planning MPC, the tracking MPC, and the extended horizon MPC are equal to the corresponding phases of UNIQUE for a fair comparison, besides the position tracking weight, which is adapted for the hierarchical lower-level tracking MPC.

The first comparison in Sect. VII-A involves real-world experiments in the reference tracking environment. We analyze performance on four different tracks with varying numbers of obstacles. Additionally, we evaluate progressive smoothing for cubic obstacles and parallel MPCs for the proposed initialization strategy, avoiding local minima. Secondly, in Sect. VII-B, we ablate different hyperparameters for the constant velocity environment to show the robustness of UNIQUE. Moreover, we evaluate how adding prediction steps to the second phase of the MPC (18) affects its evolution on the Pareto front. Finally, in Sect. VII-C, we evaluate the different constraint approximations of the second horizon and provide some qualitative insights into progressive smoothing and the parallel initialization strategy.

A. Real-World Benchmark Comparison

The real-world flights are performed across three representative trajectory shapes, with varying numbers of obstacles and different track periods, resulting in both agile and slower flight speeds. The evaluations involve (i) the flat $15 \times 8\,\mathrm{m}$ Agile Sinusoidal track with seven obstacles, a period of $10\,\mathrm{s}$ and accelerations of $20\,\frac{\mathrm{m}}{\mathrm{s}^2}$, (ii) the flat $15 \times 8\,\mathrm{m}$ Agile Butterfly track with seven obstacles and accelerations of $35\,\frac{\mathrm{m}}{\mathrm{s}^2}$, (iii) the flat $15 \times 8\,\mathrm{m}$ Cluttered Figure-Eight-Track-1 with 30 obstacles and a period of $40\,\mathrm{m}$ and (iv) a second similar Cluttered Figure-Eight-2 track with 34 obstacles. The evaluations shown in Tab. II and Fig. 4 demonstrate that UNIQUE consistently achieves a lower tracking error while requiring comparable or even reduced online computation time. The hierarchical approach increases the tracking error by up to 274.4% and standard MPC by up to 301.2%, corresponding to a reduction of the evaluation cost by 75% using UNIQUE. The significant increase in the tracking error of the standard MPC formulation in the cluttered environments emerges from the abundance of local minima that necessitate long-horizon planning. The hierarchical approach particularly struggles in the most agile Butterfly scenario, as the point-mass model can only approximate agile flight due to its conservative bounds. The trajectory visualizations in Fig. 5 show that UNIQUE tracks tighter paths with fewer deviations, especially in high-curvature segments where the hierarchical approach struggles. These results confirm that cascading models within a single optimization not only improves accuracy but also avoids the overhead of long-horizon high-fidelity MPC, leading to superior closed-loop performance in practice.

Progressive Smoothing. The influence of progressive smoothing and its significance for cubic obstacles is evaluated on the Sinusoidal track. We compare variants of UNIQUE with fixed L2-norm, infinity norm, or progressive smoothing. The trajectories are shown in Fig. 6 and reveal a superior performance of UNIQUE with progressive smoothing (mean/max tracking error of $0.6/1.4\,\mathrm{m}$), compared to the conservative L2-norm approximation (mean/max tracking error of $0.8/2.0\,\mathrm{m}$) and the numerically unstable but tight infinity norm. Remarkably, the computation time is not increased by progressive smoothing, since the norms and related parameters along the horizon are fixed. More details on progressive smoothing with different schedules and the plotted MPC predictions is shown in Appendix F.

Parallel Point-Mass MPC. We evaluate the influence of a

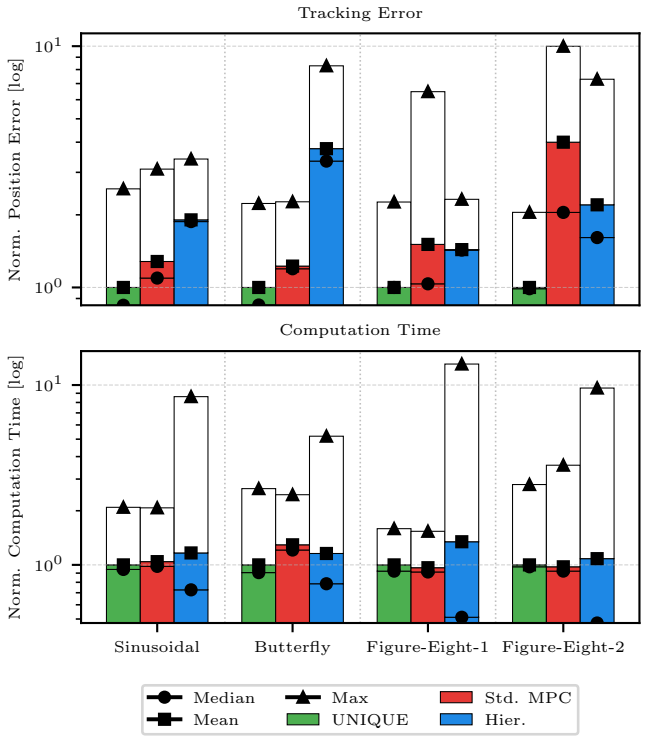

Fig. 4. Comparison between UNIQUE, the standard MPC formulation, and the phases MPC formulation regarding the online computation time and the tracking error (position). The tracking error is significantly reduced by UNIQUE, with online computation time either unchanged or slightly reduced.

TABLE II
TRACKING COST IN REAL-WORLD EXPERIMENTS.

Method	c.time (ms)	tracking error (m)
	mean/med/max	mean/med/max
	Sinusoidal	
UNIQUE	2.12 / 2.00 / 4.43	**0.32** / **0.27** / **0.82**
Standard	2.21 / 2.08 / 4.40	0.41 / 0.35 / 0.99 (+28.1%)
Hierarchical	2.47 / 1.54 / 18.27	0.61 / 0.60 / 1.09 (+90.6%)
	Butterfly	
UNIQUE	1.91 / 1.73 / 5.07	**0.39** / **0.33** / **0.87**
Standard	2.47 / 2.31 / 4.69	0.48 / 0.47 / 0.88 (+23.1%)
Hierarchical	2.21 / 1.50 / 9.92	1.46 / 1.30 / 3.23 (+274.4%)
	Figure-Eight-1	
UNIQUE	7.02 / 6.48 / 11.16	**0.79** / **0.79** / **1.79**
Standard	6.77 / 6.41 / 10.79	1.20 / 0.82 / 5.14 (+51.9%)
Hierarchical	9.44 / 3.59 / 91.84	1.14 / 1.13 / 1.84 (+44.3%)
	Figure-Eight-2	
UNIQUE	7.83 / 6.98 / 21.90	**0.83** / **0.82** / **1.71**
Standard	7.64 / 7.23 / 28.05	3.33 / 1.71 / 8.32 (+301.2%)
Hierarchical	8.48 / 3.72 / 75.34	1.83 / 1.34 / 6.08 (+120.5%)

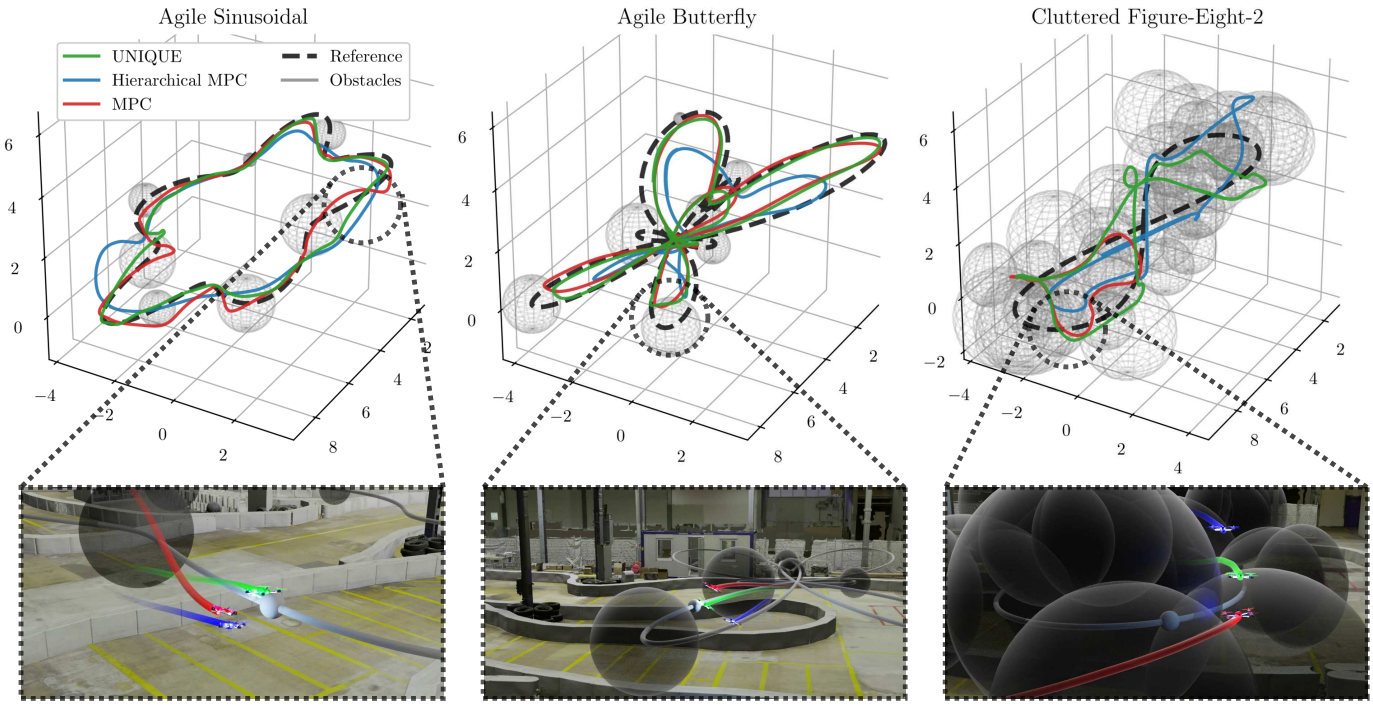

Fig. 5. Visualization and rendered images of trajectories from different controllers in real-world experiments in a flying arena. We compare the proposed UNIQUE framework with the hierarchical and MPC settings. The top plots show agile flights with seven smooth obstacles, therefore, less susceptible to local minima. Yet, UNIQUE still outperforms the two baselines due to the increased long-horizon prediction. Standard MPC struggles once an obstacle appears on the horizon, attempting to evade it with more costly maneuvers. The hierarchical MPC planner and controller setting is limited by the coarse path of the higher-level planner (point mass model), which is apparent in the agile flight of the Butterfly track. The lower plots illustrate the cluttered environment, which exhibits numerous local minima. The standard MPC is not even capable of following the reference coarsely, as it gets stuck in the Figure-Eight-2 track on the left. The hierarchical framework escapes some local minima due to the long-horizon planning, but still exhibits a large tracking error. Only the UNIQUE framework is able to track the cluttered Figure-Eight tracks due to its reinitialization strategy, which utilizes parallel and long-horizon planning. The axes are labeled in SI units.

parallel point-mass MPC in the Cluttered-Figure-Eight-2 track, where we activate or deactivate the parallel point-mass MPC. We perform 7 RTI steps on the parallel point-mass MPC before we randomly reinitialize. An active re-initialization and the related costs are plotted in Fig. 7. Notably, the parallel point-mass MPC computation time is evaluated assuming parallel cores, i.e., the total UNIQUE computation time is $\max(t_{\text{comp},0}, \dots, t_{\text{comp,P}})$, where $t_{\text{comp},0}$ is the computation time of the multiphase MPC and $t_{\text{comp},1}, \dots, t_{\text{comp,P}}$ are the computation times of the parallel point-mass MPCs. The mean and median computation time do not increase in our experiments compared to the multi-phase MPC without parallelization. However, the maximum computation time is increased from $10.9\,\text{ms}$ to $21.9\,\text{ms}$, which is still below the benchmark approaches compared, cf. Fig. 5 and Tab. II. Considering the closed-loop cost, the difference when using parallel point-mass MPCs is remarkable in this obstacle-rich environment, as the mean cost decreases by 77.1% from $3.5/2.0/8.3$ (mean/median/max) when using UNIQUE without parallel roll-outs to $0.8/0.8/1.7$ when using the parallel roll-outs. This is due to the multi-phase MPC getting trapped in the local minima, similarly to the MPC in Fig. 5. In the Agile Sinusoidal or Agile Butterfly tracks, the influence of the parallel point-mass MPCs was evaluated to be negligible due to the absence of dominant local minima.

B. Evaluation of Hyperparameters

The proposed UNIQUE formulation comprises numerous hyperparameters, which we compare in simulated experiments within a constant velocity environment by evaluating their closed-loop cost. The most critical hyperparameters are the particular horizon lengths, number of shooting nodes of the short and long horizon, and the choice of the low-fidelity force constraint approximation $\mathbb{S}_f^{\text{lf}}$ or $\hat{\mathbb{S}}_f^{\text{lf}}$ and the low-fidelity body rate constraint approximation $\mathbb{S}_\omega^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$ or $\widehat{\mathbb{S}}_\omega^{\text{lf}}$. Additionally, the discretization time can be spaced by a factor c_{space} for future horizons, for example by $t_{\Delta,i+1} = c_{\text{space}} t_{\Delta,i}$. The following section aims to answer crucial hyperparameter evaluation questions.

How do hyperparameters affect the performance? Fig. 8 shows the comparison on the Pareto front of closed loop performance and online computation time for the RTI scheme and the fully converged solvers, with parameters displayed in Tab. III. The comparison includes standard MPC with an increasing spacing of the high-fidelity horizon, which constitutes an alternative method for planning coarser in the distant future. However, due to the nonlinearity of the high-fidelity model, the spacing leads to convergence problems. The three different horizons of the standard MPC hint a barrier in the Pareto plot in Fig. 8 and also, the hierarchical approach clusters in a specific region. Both benchmarks are clearly outperformed by UNIQUE.

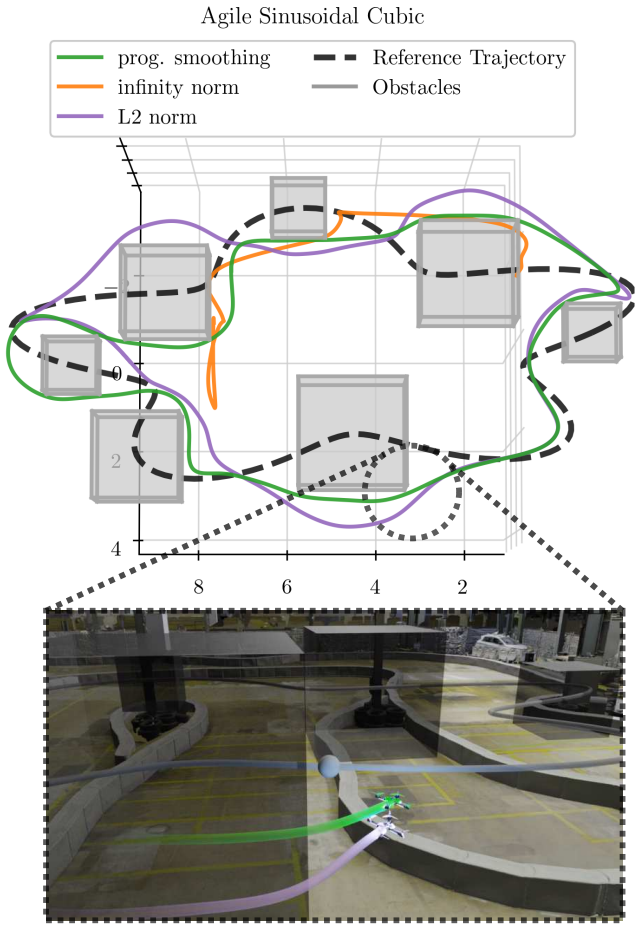

Fig. 6. Real-world comparison of UNIQUE with different formulations of cubic virtual obstacles. We compare the proposed progressive smoothing against a standard L2 norm and tight infinity norm. The L2 norm is conservative, while the infinity norm leads to numerically unstable behavior. Progressive smoothing allows for a numerically stable, tight cubic obstacle representation. The lower rendering visualizes the L2 norm and progressive smoothing variant on a particular snap shot.

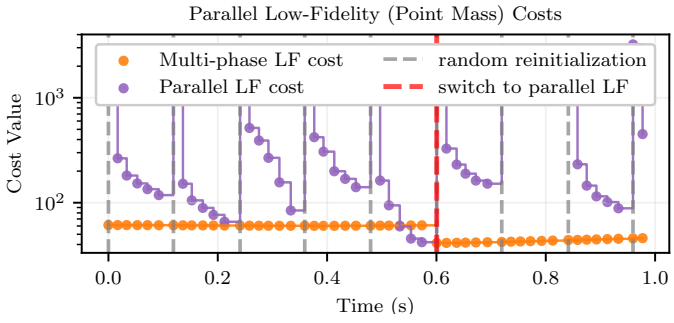

Fig. 7. Comparison of the parallel point-mass MPC initialization approach. One parallel point-mass MPC is structurally identical to the second phase of the main multi-phase MPC and randomly initialized every seventh step. Once the parallel MPC cost is lowest, the second phase of the main MPC is initialized with the corresponding lower-cost decision variables (at 0.6 s). Due to the random initialization and RTI, the parallel MPC cost decreases significantly over the seven iterations, which are clearly visible as "stairs".

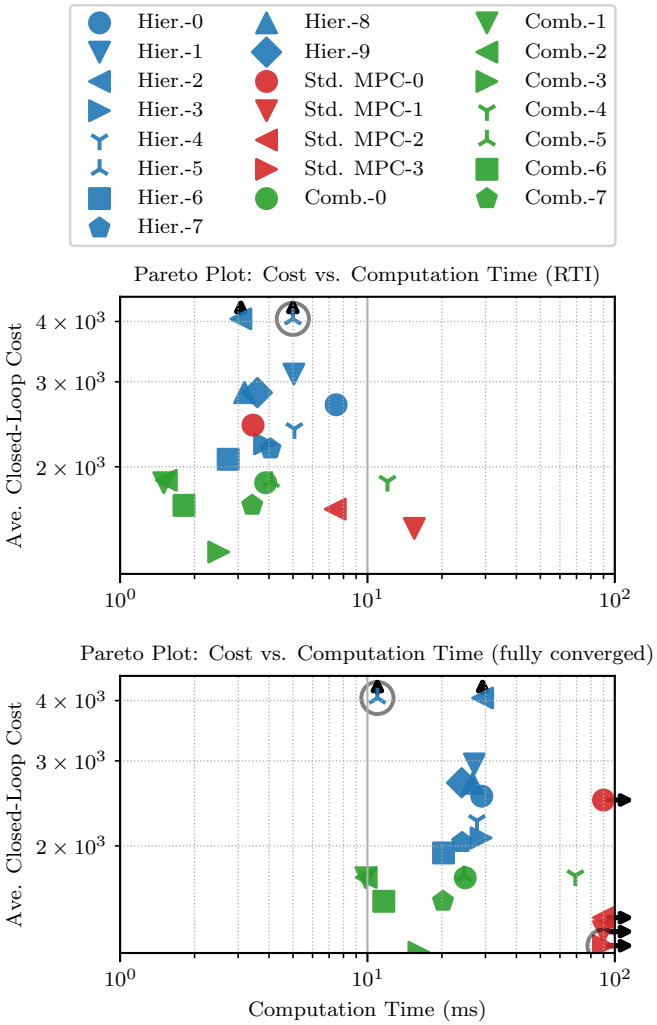

Fig. 8. Pareto comparison of the closed-loop cost and average computation time per iteration in the constant velocity environment. We compare different hyperparameter settings for the hierarchical, the standard MPC, and the unified formulation. The top plot shows a comparison for the real-time iteration scheme, while the bottom plot shows a comparison for the fully converged solver, which is nearly an order of magnitude slower. The numerically challenging long-horizon standard MPC formulations require comparably much higher computation times to fully converge than UNIQUE, which is also Pareto optimal in both plots.

How do hyperparameters affect convergence? Although the RTI scheme is the only applicable approach in real-world evaluations, the fully converged evaluation shown in Fig. 8 and Tab. III contrasts the different approaches in terms of their numerical complexity. Remarkably, the converged UNIQUE formulation requires ten times more computation time. In contrast, the standard MPC formulation increases computation time by a factor of one hundred due to the generally more complex optimization problem that emerges from the increased horizon. This highlights two favorable properties: multi-phase MPC (i) converges fully up to 100 times faster than std. MPC (cf. Tab. III), and (ii) the required solver time difference to the RTI is smaller by a factor of 10.

How does the second phase influence performance? In Fig. 9, we investigate how adding steps of the second phase

TABLE III
PARAMETER AND EVALUATION OF UNIQUE, THE HIERARCHICAL SETTING, AND THE STANDARD MPC CONTROLLER IN THE RANDOMIZED CONSTANT VELOCITY SIMULATION EXPERIMENT.

Experiment Unit	Hor. / t_Δ (s)	Nodes	Spacing	Hor. / t_Δ^{lf} (s)	Nodes	Spacing	Interval	Constr.	Cost Mean Val.	Comp. Time Mean Val (ms)
	High-Fid. Part			Low-Fidelity Part					RTI / Converged	
UNIQUE-0	0.52 / 0.04	13		17.60 / 0.80	22	0.10		$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	1853.74 / 1717.60	3.88 / 24.85
UNIQUE-1	0.52 / 0.04	13		17.60 / 0.80	22	0.10		$\hat{\mathbb{S}}_f^{\text{lf}}$, $\widehat{\mathbb{S}}_\omega^{\text{lf}}$	1853.74 / 1717.60	**1.50 / 9.84**
UNIQUE-2	0.52 / 0.04	13		16.00 / 0.80	20	0.20		$\hat{\mathbb{S}}_f^{\text{lf}}$, $\widehat{\mathbb{S}}_\omega^{\text{lf}}$	1875.22 / 1720.68	1.54 / 9.88
UNIQUE-3	0.80 / 0.04	20		7.20 / 0.80	9	0.00		$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	**1331.95 / 1203.88**	2.50 / 16.15
UNIQUE-4	0.80 / 0.04	20		15.20 / 0.20	76	0.00		$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	1865.09 / 1734.07	12.05 / 69.26
UNIQUE-5	0.80 / 0.04	20		15.20 / 0.20	76	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}$, $\widehat{\mathbb{S}}_\omega^{\text{lf}}$	1865.12 / 1734.09	4.08 / 24.59
UNIQUE-6	0.80 / 0.04	20		3.20 / 0.20	16	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}$, $\widehat{\mathbb{S}}_\omega^{\text{lf}}$	1664.14 / 1535.86	1.81 / 11.64
UNIQUE-7	0.80 / 0.04	20		3.20 / 0.20	16	0.00		$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	1665.23 / 1536.51	3.42 / 20.21
	Controller			Planner						
Hier.-0	2.00 / 0.04	50		24.00 / 0.20	120		10	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2688.66 / 2535.47	7.47 / 28.98
Hier.-1	2.00 / 0.04	50		24.00 / 0.40	60		10	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	3112.53 / 2957.01	5.04 / 27.06
Hier.-2	2.00 / 0.04	50		8.00 / 0.20	40		10	$\hat{\mathbb{S}}_f^{\text{lf}}$, $\widehat{\mathbb{S}}_\omega^{\text{lf}}$	7170.31 / 7019.67	3.08 / 29.21
Hier.-3	2.00 / 0.04	50		8.00 / 0.20	40		10	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2231.18 / 2077.67	3.84 / 28.81
Hier.-4	2.00 / 0.04	50		8.00 / 0.20	40		4	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2397.97 / 2255.47	5.07 / 27.72
Hier.-5	0.40 / 0.04	10		8.00 / 0.20	40		4	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	22189.62 / 30800.65	4.99 / 10.97
Hier.-6	1.20 / 0.04	30		8.00 / 0.20	40		10	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2072.90 / 1934.08	2.73 / 20.31
Hier.-7	1.20 / 0.04	30		8.00 / 0.20	40		4	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2178.72 / 2033.14	4.06 / 24.13
Hier.-8	2.00 / 0.04	50		4.00 / 0.20	20		10	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2848.82 / 2695.63	3.18 / 26.60
Hier.-9	2.00 / 0.04	50		4.00 / 0.20	20		4	$\mathbb{S}_f^{\text{lf}}$, $\hat{\mathbb{S}}_\omega^{\text{lf}}$	2848.93 / 2703.57	3.59 / 24.04
	Controller									
Std. MPC-0	2.00 / 0.04	50	0.00						2439.76 / 2489.68	3.45 / 474.95
Std. MPC-1	8.00 / 0.04	200	0.00						1490.52 / 1328.99	15.47 / 2102.14
Std. MPC-2	4.00 / 0.04	100	0.00						1634.34 / 1420.70	7.41 / 1013.03
Std. MPC-3	7.92 / 0.04	55	0.10						nan / 1242.62	5.23 / 152.10

influences the Pareto performance in comparison to how adding standard MPC steps influences the performance. We take standard MPC with $n_{\text{hf}} = \{10, 15, 20, 30, 50\}$ shooting nodes and $t_\Delta^{\text{hf}} = 40\,\text{ms}$ (red line in Fig. 9) and append to each standard high-fidelity MPC a the low-fidelity part of $n_{\text{lf}} = \{3, 5, 7\}$ with $t_\Delta^{\text{if}} = 400\,\text{ms}$ (black dotted line). The reference tracking involves three obstacles and shapes of a Triangle, a rounded Rectangle, a Figure Eight, and a Butterfly. Additionally, we add a small sinusoidal z-component. The evaluation on the Pareto front reveals that it is more efficient to add the low-fidelity horizon of UNIQUE than to increase the MPC horizon.

C. Evaluation of Constraint Approximation

In this Section, we evaluate the influence of the different constraint formulations on the low-fidelity model phase. Particularly, we assess the rise time, i.e., the time it takes to reach 90% of the set velocity of $14\,\frac{\text{m}}{\text{s}}$ from hovering using the different low-fidelity force constraint approximation $\mathbb{S}_f^{\text{lf}}$ (polyhedral) or $\hat{\mathbb{S}}_f^{\text{lf}}$ (box) and the low-fidelity body rate constraint approximation $\mathbb{S}_\omega^{\text{lf}}$ (nonlinear), $\hat{\mathbb{S}}_\omega^{\text{lf}}$ (polyhedral) or $\widehat{\mathbb{S}}_\omega^{\text{lf}}$ (box)[1]. Fig. 10 shows the corresponding velocity state v_y, the related body rate ω_x, and the evaluation of the three different constraint functions that approximate the body rates. In addition, all MPC predictions of both phases are plotted, as well as one particular prediction at step at $k = 5$ equal to $0.2\,\text{s}$. As anticipated, the more conservative constraints lead to an increased rise time. The conservatism can be verified in the lower three rows showing their constraint function evaluations. Yet, despite the visible larger conservatism in the second horizon, the re-optimization along the first high-fidelity horizon can improve the closed-loop performance.

Trivially, if the first horizon is increased, the second horizon and the related constraints become less relevant. We evaluated this dependency in Fig. 11, where we increased the first high-fidelity horizon and evaluated the rise time and computation time for RTIs and the fully converged solution. If the first horizon is increased over $20\,\text{ms}$, the constraint formulation on the second horizon becomes irrelevant for the particular maneuver. Both the converged and the RTI evaluations show an increasing difference in the constraint formulation for shorter horizons of the first phase. Using box constraints on body rates with $\hat{\mathbb{S}}_\omega^{\text{lf}}$ and thrust $\widehat{\mathbb{S}}_\omega^{\text{lf}}$ shows particular in RTI the lowest computation time. The constraint formulation for the thrust shows no difference in the evaluation.

VIII. CONCLUSION AND DISCUSSION

The proposed UNIQUE framework unifies quadrotor motion planning and control by embedding models of different fidelity within a single MPC formulation. A long-horizon, low-fidelity point-mass model provides foresight, while a short-horizon, high-fidelity model ensures precise control, achieving

[1]We use a color coding for body rate constraint types to simplify the reading

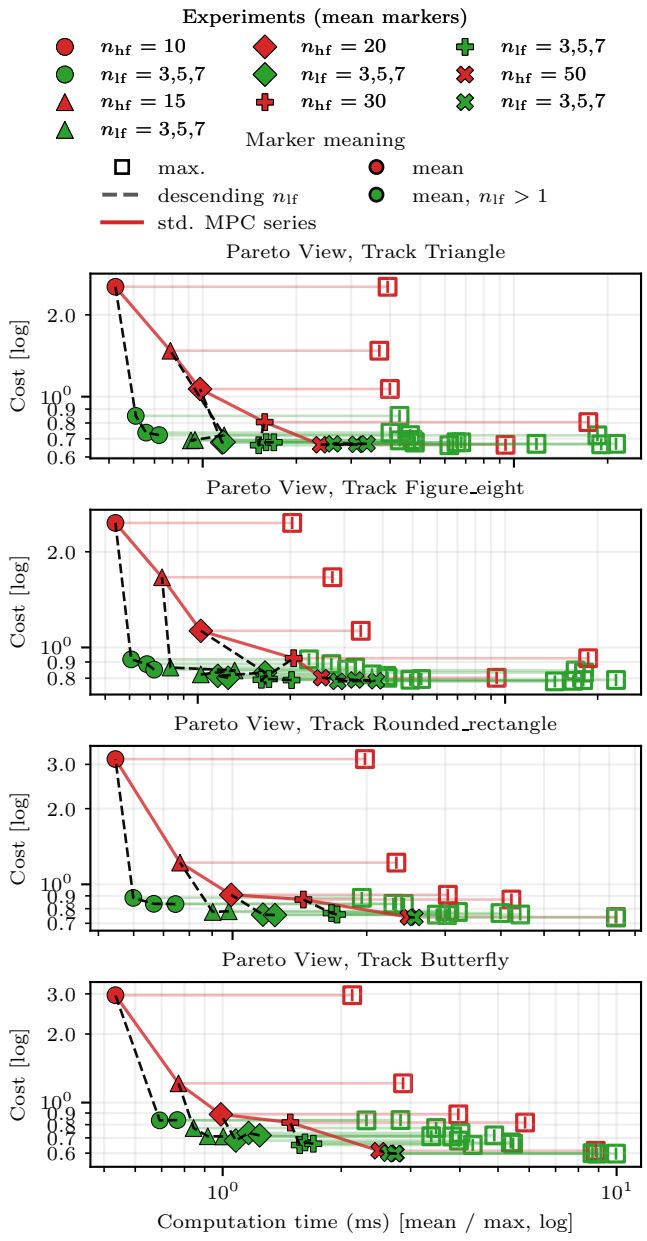

Fig. 9. Pareto comparison between increasing the horizon of the standard MPC formulation (red line), as opposed to increasing the horizon with the planner formulation (green markers and dashed gray line) on different test tracks. Remarkably, adding the second low-fidelity horizon outperforms MPC with increased horizons.

real-time performance without hierarchical decomposition. By coupling both models through feasibility-preserving transition constraints, UNIQUE maintains recursive feasibility across phases. This integration avoids the redundancy and limitations of the point-mass model typical of planner–tracker pipelines, and mitigates the limited foresight of single-model MPCs.

The key mechanisms of *progressive obstacle smoothing* and *parallel low-fidelity MPCs* enhance convergence in cluttered environments, preventing convergence to highly suboptimal local minima without compromising computation speed. Experiments show a reduction of up to 75% in closed-loop cost compared to hierarchical and standard MPC baselines. Both mechanisms still can not provide guarantees of finding the global optimal solution, which is common for NMPC. However, the experiments reveal the practical low-cost closed-loop performance. Notably, evaluating the multiple shooting cost of an RTI trajectory is nontrivial due to infeasible gaps in the dynamics. We simply use the cost of the potentially infeasible trajectory as it resembles the closed-loop cost sufficiently in our experiments. More advanced cost evaluation strategies are presented in [16].

Another limitation is that tuning the cost of the point-mass model remains task-dependent. As the transformation between point-mass and quadrotor states is not bijective, a generic task-independent tuning of the costs is impossible. Extending the framework to automatic cost tuning for particular tasks is considered in future work.

UNIQUE is the first recursive feasible real-time MPC that unifies planning and control through multi-fidelity cascading. It provides long-horizon foresight and short-horizon accuracy in a single computationally efficient NLP.

REFERENCES

[1] A. Romero, R. Penicka, and D. Scaramuzza, "Time-optimal online replanning for agile quadrotor flight," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7730–7737, July 2022.

[2] A. Romero, S. Sun, P. Foehn, and D. Scaramuzza, "Model predictive contouring control for time-optimal quadrotor flight," *IEEE Transactions on Robotics*, vol. PP, pp. 1–17, 12 2022.

[3] M. Krinner, A. Romero, L. Bauersfeld, M. Zeilinger, A. Carron, and D. Scaramuzza, "MPCC++: Model predictive contouring control for time-optimal flight with safety constraints," in *Robotics: Science and System*, D. Kulic, G. Venture, K. Bekris, and E. Coronado, Eds. s.l.: Robotics Science and Systems Foundation, 2024, Conference Paper, p. 109, robotics: Science and Systems Conference (RSS 2024); Conference Location: Delft, Netherlands; Conference Date: July 15-19, 2024.

[4] M. Hehn and R. D'Andrea, "Real-time trajectory generation for quadrocopters," *IEEE Transactions on Robotics*, vol. 31, no. 4, pp. 877–892, 2015.

[5] 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.

[6] R. Reiter, K. Baumgärtner, R. Quirynen, and M. Diehl, "Progressive smoothing for motion planning in real-time nmpc," in *2024 European Control Conference (ECC)*, 2024, pp. 1816–1823.

[7] T. Cieslewski, E. Kaufmann, and D. Scaramuzza, "Rapid exploration with multi-rotors: A frontier selection method for high speed flight," in *IROS*, 2017.

[8] M. Naazare, D. Ramos, J. Wildt, and D. Schulz, "Application of graph-based path planning for UAVs to avoid restricted areas," in *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2019, pp. 139–144.

[9] W. Liu, Y. Ren, and F. Zhang, "Integrated planning and control for quadrotor navigation in presence of suddenly appearing objects and disturbances," *IEEE Robotics and Automation Letters*, 2024.

[10] M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "Path planning for motion dependent state estimation on micro aerial vehicles," in *2013 IEEE International Conference on Robotics and Automation*, 2013.

[11] F. Meyer, K. Glock, and D. Sayah, "Top-uav: Open-source time-optimal trajectory planner for point-masses under acceleration and velocity constraints," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 2838–2845.

[12] K. Teissing, M. Novosad, R. Penicka, and M. Saska, "Real-time planning of minimum-time trajectories for agile UAV flight," *IEEE Robotics and Automation Letters*, 2024.

[13] E. Tal and S. Karaman, "Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 1203–1218, 2021.

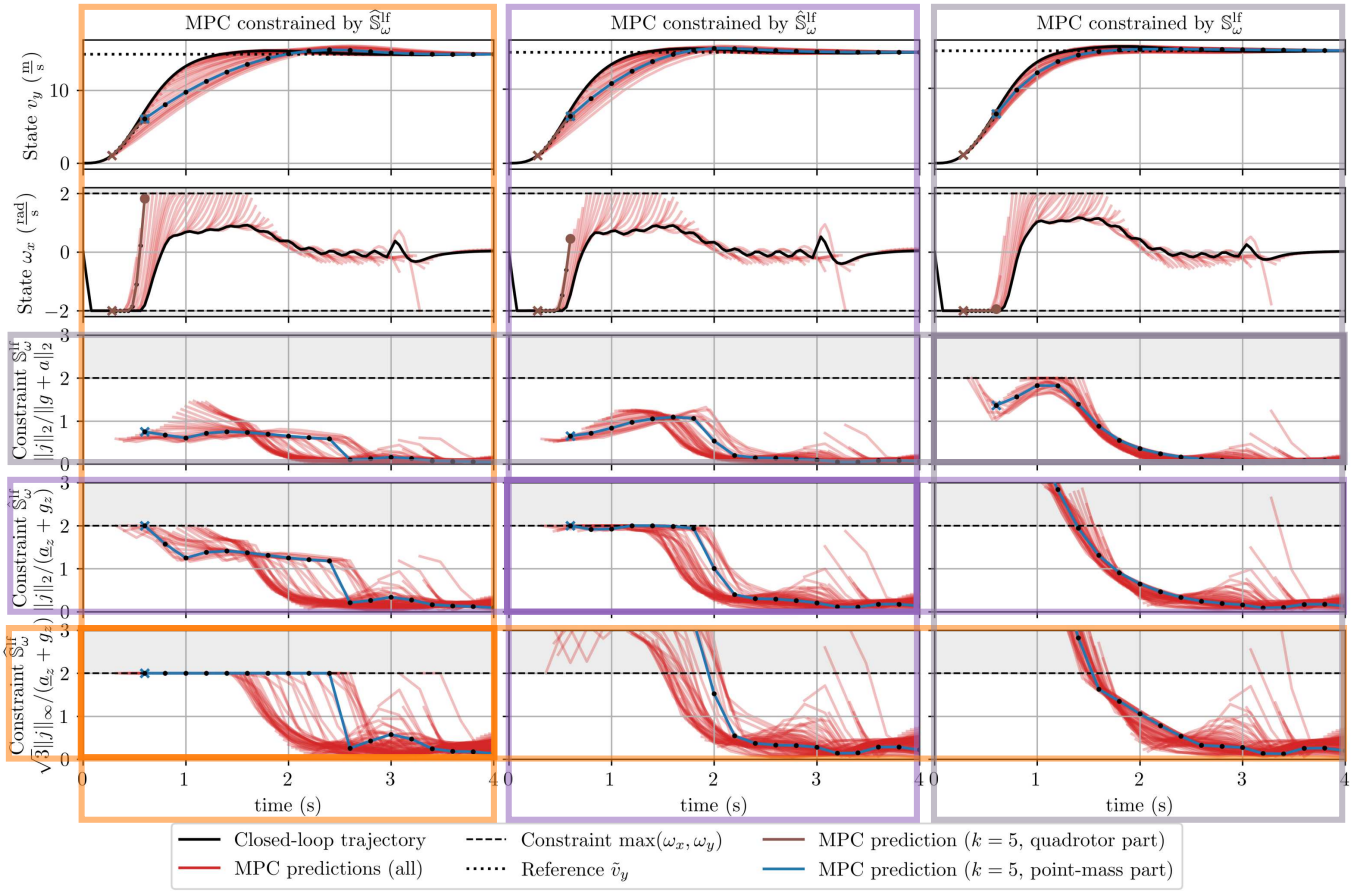

Fig. 10. Visualization of different safe set approximations for body rates (three columns: box constraints $\widehat{\mathbb{S}}_\omega^{\text{lf}}$, polyhedral $\hat{\mathbb{S}}_\omega^{\text{lf}}$, nonlinear $\mathbb{S}_\omega^{\text{lf}}$) for a step response to a horizontal set speed ($14\frac{\text{m}}{\text{s}}$, first row). The second row shows the body rates and the respective constraints for the high-fidelity model. The lower three rows show the different constraint function evaluations on the lower-fidelity part, where the lower white part is feasible. The third row shows the most accurate nonlinear constraints. It reveals that these are satisfied for all formulations, however, more conservatively, for the box constraints and the polyhedral constraints. The lower off-diagonal plots show the respective active constraint evaluation.

[14] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 620–626, 2018.

[15] M. Wang, Q. Wang, Z. Wang, Y. Gao, J. Wang, C. Cui, Y. Li, Z. Ding, K. Wang, C. Xu *et al.*, "Unlocking aerobatic potential of quadcopters: Autonomous freestyle flight generation and execution," *Science Robotics*, vol. 10, no. 101, p. eadp9905, 2025.

[16] R. Reiter, A. Ghezzi, K. Baumgärtner, J. Hoffmann, R. D. McAllister, and M. Diehl, "Ac4mpc: Actor-critic reinforcement learning for guiding model predictive control," *IEEE Transactions on Control Systems Technology*, pp. 1–16, 2025.

[17] K. Seel, A. B. Kordabad, S. Gros, and J. T. Gravdahl, "Convex neural network-based cost modifications for learning model predictive control," *IEEE Open Journal of Control Systems*, 2022.

[18] S. Abdufattokhov, M. Zanon, and A. Bemporad, "Learning Lyapunov terminal costs from data for complexity reduction in nonlinear model predictive control," *International Journal of Robust and Nonlinear Control*, 2024.

[19] G. De Nicolao, L. Magni, and R. Scattolini, "Stabilizing receding-horizon control of nonlinear time varying systems," *IEEE Trans. Automatic Control*, 1998.

[20] M. Diehl, L. Magni, and G. D. Nicolao, "Efficient NMPC of unstable periodic systems using approximate infinite horizon closed loop costing," *Annual Reviews in Control*, 2004.

[21] A. Ghezzi, R. Reiter, K. Baumgärtner, A. Bemporad, and M. Diehl, "A numerically efficient method to enhance model predictive control performance with a reinforcement learning policy," 2025.

[22] V. Behrunani, P. Heer, R. Smith, and J. Lygeros, "Recursive feasibility guarantees in multi-horizon mpc," in *IEEE 63rd Conference on Decision and Control (CDC)*. Piscataway, NJ: IEEE, 2024, Conference Paper, pp. 297 – 302.

[23] J. Frey, K. Baumgaertner, G. Frison, and M. Diehl, "Multi-phase optimal control problems for efficient nonlinear model predictive control with acados," *Optimal Control Applications and Methods*, vol. 46, no. 2, pp. 827–845, 2025.

[24] D. Leineweber, A. Schäfer, H. Bock, and J. Schlöder, "An efficient multiple shooting based reduced sqp strategy for large-scale dynamic process optimization: Part ii: Software aspects and applications," *Computers and Chemical Engineering*, vol. 27, pp. 167–174, 02 2003.

[25] H. Li, R. J. Frei, and P. M. Wensing, "Model hierarchy predictive control of robotic systems," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3373–3380, 2021.

[26] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Mueller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, pp. 982–987, 08 2023.

[27] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.

[28] 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.

[29] R. Verschueren, G. Frison, D. Kouzoupis, J. Frey, N. v. Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados—a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, no. 1, pp. 147–183, Mar. 2022.

[30] G. Frison and M. Diehl, "HPIPM: a high-performance quadratic programming framework for model predictive control," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6563–6569, Jan. 2020.

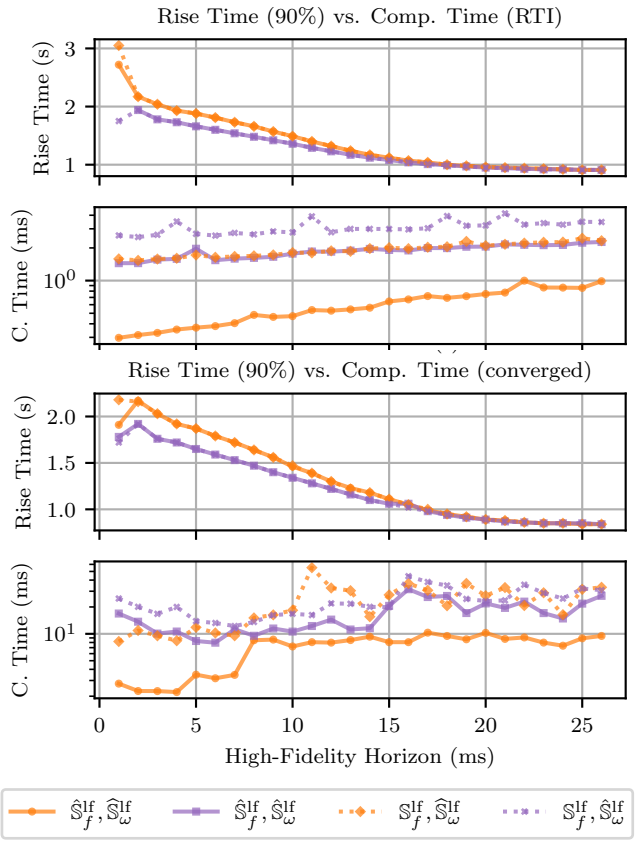

Fig. 11. Comparison of the rise time (90% of the set value) and computation time for a horizontal velocity step regarding different high-fidelity horizons using a large low-fidelity horizon. The experiments reveal the diminishing influence of the different safe set formulations on the closed-loop performance when the first high-fidelity model horizon is increased. The more accurate safe set formulations require higher computation times but increase the performance, particularly for the shorter high-fidelity model phase.

[31] M. Diehl, H. Bock, and J. Schlder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM Journal on Control and Optimization*, vol. 43, pp. 1714–1736, 07 2006.

[32] J. Nocedal and S. Wright, *Numerical optimization*, 2nd ed., ser. Springer series in operations research and financial engineering. New York, NY: Springer, 2006.

[33] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, and D. Scaramuzza, "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *Science Robotics*, vol. 7, no. 67, p. eabl6259, 2022.

[34] B. D. Team, "Betaflight — open source flight controller firmware," 2025, firmware for multi-rotor and fixed wing craft. GPL-3.0 licensed. [Online]. Available: https://github.com/betaflight/betaflight

APPENDIX

A. Residual Models for Aerodynamic Effects

Different types of residual forces related to their fidelity are modeled in the literature for quadrotors. We consider a polynomial higher-order model with increased accuracy but sophisticated nonlinearities for motion planning [26]. Aerodynamic effects are considered by a residual force $f_{\text{res}}(v_{\mathcal{B}}, \Omega^2) := \begin{bmatrix} f_x^{\text{res}} & f_y^{\text{res}} & f_z^{\text{res}} \end{bmatrix}^\top$. The aerodynamic effects are assumed to primarily depend on the velocity $v_{\mathcal{B}}$ in the body frame and the average squared motor speed $\overline{\Omega^2} = 1/4\,\Omega^\top \Omega$. The forces f_x^{res} and f_y^{res} represent induced drag, and f_z^{res} is induced lift. The quantities v_x, v_y, and v_z denote the three axial velocity components in the body frame, and v_{xy} denotes the velocity in the (x, y) plane of the quadrotor. Based on insights from the underlying physical processes, linear and quadratic combinations of the individual terms are selected. The forces are modeled by

$$
\begin{aligned}
f_x^{\text{res}}(v_{\mathcal{B}}, \Omega) &:= m \begin{bmatrix} 1 & v_x & v_x|v_x| & v_x\,\overline{\Omega^2} \end{bmatrix} c_x, \\
f_y^{\text{res}}(v_{\mathcal{B}}, \Omega) &:= m \begin{bmatrix} 1 & v_y & v_y|v_y| & v_y\,\overline{\Omega^2} \end{bmatrix} c_y, \\
f_z^{\text{res}}(v_{\mathcal{B}}, \Omega) &:= m c_z^\top \\
&\quad \begin{bmatrix} 1 & v_z & v_z^3 & v_{xy} & v_{xy}^2 & v_{xy}\,\overline{\Omega^2} & v_{xy}\,v_z\,\overline{\Omega^2} & \overline{\Omega^2} \end{bmatrix}^\top,
\end{aligned}
$$

where the weights $c_x \in \mathbb{R}^3$, $c_y \in \mathbb{R}^3$ and $c_z \in \mathbb{R}^8$ are obtained from fitting the model to real-world data.

TABLE IV
RESIDUAL PARAMETERS

Parameter	Value / Description
High-Fidelity Quadrotor Model	
c_x	[1.18e-02, -1.39e-01, -1.59e-03, -8.31e-08]
c_y	[-3.21e-02, -9.79e-02, -6.85e-03, -1.01e-07]
c_z	[-5.00e-01, 1.16e-01, -4.25e-01, -1.04e-06, 1.76e-07, -1.89e-08, -1.03e-03]

B. Differentiable flatness for quadrotors

Differentiable flatness refers to systems where there exists an output $y = h(x, u, \dot{u}, \ddot{u}, \ldots, u^{(r)})$, such that all states x and controls u can be described by a a function of derivatives of y, i.e.,

$$x = \gamma_1(y, \dot{y}, \ldots, y^{(k)}), \tag{21}$$

$$u = \gamma_1(y, \dot{y}, \ldots, y^{(k)}), \tag{22}$$

for some functions γ_1, γ_2 and constants $r, k \in \mathbb{N}$. As shown in [14], quadrotors with linear disturbance models, cf. App. A, are differentiable flat for an output $[p^\top, \psi]^\top$, where $\psi \in \mathbb{R}$ is the heading angle. The property of differentiable flatness implies that trajectory planning can be done directly in the space of flat outputs. This means that trajectories can be designed for the flat outputs, and the full system trajectory recovered.

C. Quadrotor jerk

In the following, we derive the jerk of the quadrotor similar to [14]. Let the time-dependent expression be defined as

$$ma = q \odot T + q \odot \left(C(q^{-1} \odot v)\right) - mg$$

We aim to compute the time derivative:

$$m\dot{a} = mj = \frac{d}{dt}\left(q \odot T\right) + \frac{d}{dt}\left(q \odot \left(C(q^{-1} \odot v)\right)\right)$$

The derivative of the first term is

$$\frac{d}{dt}(q \odot T) = q \odot \left(\dot{f}_{\text{th}} + \omega \times T\right)$$

Next, the derivative of the second term is computed. Let $u = q^{-1} \odot v$. Then

$$\frac{d}{dt}(q \odot (Cu)) = q \odot (C\dot{u} + \omega \times (Cu))$$

and with

$$\dot{u} = \frac{d}{dt}(q^{-1} \odot v) = q^{-1} \odot (\dot{v} - \omega \times v)$$

we obtain

$$\frac{d}{dt}\left(q \odot \left(C(q^{-1} \odot v)\right)\right) =$$
$$q \odot \left(C\left(q^{-1} \odot (\dot{v} - \omega \times v)\right) + \omega \times \left(C(q^{-1} \odot v)\right)\right)$$

The final expression for the quadrotor jerk can consequently be written as

$$\begin{aligned} j = j_{\text{quad}}(q, v, \omega, T, \dot{f}_{\text{th}}) = \\ & \frac{1}{m} q \odot \left(\dot{f}_{\text{th}} + \omega \times T\right) \\ & + \frac{1}{m} q \odot \Big(C\left(q^{-1} \odot (\dot{v}(q, v, T) - \omega \times v)\right) \\ & + \omega \times \left(C(q^{-1} \odot v)\right)\Big) \end{aligned}$$

with $\dot{v}$ as defined in (1).

D. Parameters for Constant Velocity Experiments

Table V shows the parameters used in the Constant Velocity experiments. We used equal parameters for UNIQUE, the standard MPC and the hierarchical setting, if not stated differently in Tab. II for the particular hyperparameter comparison. These experiments do not use the parallel point-mass MPCs.

E. Parameters for Tracking experiments

Table VI shows the parameters used in the Tracking experiments. We used equal parameters for UNIQUE, the standard MPC and the hierarchical setting. These experiments use one parallel point-mass MPC in the Cluttered-Figure-Eight environment. In the cubic obstacle environment, the equal obstacle positions and dimensions are used but with the maximum volume inscribed cube of the ellipsoid. Parameters not stated are equal to Tab. V

F. Obstacle avoidance with progressive smoothing

To give more qualitative insights into progressive smoothing, we plot the predicted MPC trajectories in the constant velocity environment with static obstacles in Fig. 12. We compare the predictions with 2-norm, infinity norm, and two progressive smoothing schedules, that linearly vary the norm from 10 to 2 in either 5 or 10 s. The evaluation reveals a slightly different behavior when using progressive smoothing. However, the closed-loop trajectories are similar and both traverse much tightly along the cubic obstacle shapes.

TABLE V
PARAMETERS FOR CONSTANT VELOCITY EXPERIMENTS (SI UNITS)

Parameter	**Value / Description**
High-Fidelity Quadrotor Model	
w	[0, 0.01, 1.0, 30, 30, 30, 10, 0, 0, 10, 10, 10, 3, 3, 3, 3e^{-5}, 3e^{-5}, 3e^{-5}, 3e^{-5}]
$\tilde{y}$	[0, 0, 0, 0, 0, 0, 15, 0, 0, 0, 0, 0, 1.47, 1.47, 1.47, 1.47, 0, 0, 0, 0]
Low-Fidelity Point-Mass Model	
w^{lf}	[0.0, 0.01, 1, 10, 0.0001, 0.0001, 1.08, 1.08, 1.08, 0.1, 0.1, 0.1]
$\tilde{y}^{\text{lf}}$	[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
Point Mass Constraint Parameters	
$\underline{a}_z$	-5.0 (requirement: $\underline{a}_z > -g + \underline{f}_{\text{th}}$)
$\overline{f}_{\text{res}}$	2.0
α_x	0.5
α_z	0.5
Tracking MPC Weights (different from above)	
$w[0:2]$	[10, 10, 10]
Environment Settings	
Ellipsoid dims.	[34, 34, 34], [25, 25, 25], [55, 55, 55], [52, 52, 52]
Ellipsoid pos.	[155, 20, 0], [50, -20, 0], [255, 20, 0], [450, -20, 0]
Max speeds	2
Movement type	Brownian motion
Simulation Settings	
Drone start	[0, 0, 0]
Simulation steps	10000
Time step	0.04
Evaluation	
$Q_{\text{eval}}, R_{\text{eval}}$	$\text{diag}(Q_{\text{eval}}, R_{\text{eval}}) = diag(w)$ (same as MPC costs)

TABLE VI
PARAMETERS FOR TRAJECTORY TRACKING EXPERIMENTS (SI UNITS)

Parameter	**Value / Description**
High-Fidelity Quadrotor Model	
w	[500, 500, 500, 10, 10, 10, 0, 0, 0, 10, 10, 10, 3, 3, 3, 3, 3e^{-5}, 3e^{-5}, 3e^{-5}, 3e^{-5}]
$\tilde{y}$	[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1.47, 1.47, 1.47, 1.47, 0, 0, 0, 0]
Low-Fidelity Point-Mass Model	
w^{lf}	[500, 500, 500, 0, 0, 0, 0.05, 0.05, 0.05, 0.1, 0.1, 0.1]
$\tilde{y}^{\text{lf}}$	[0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
Environment Agile Sinusoidal	
Ellipsoid radii	.5, .2, .5, .8, .5, .8, .9
Ellipsoid pos.	Equal spacing
UNIQUE Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	23, 10 / 0.02, 0.2
Hier. Appr. Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	23, 12 / 0.02, 0.2
Std. MPC Hor. N/t_Δ	30 / 0.02
Environment Agile Butterfly	
Ellipsoid radii	.5, .2, .5, .8, .5, .8, .9
Ellipsoid pos.	Equal spacing
UNIQUE Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	23, 5 / 0.02, 0.1
Hier. Appr. Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	23, 12 / 0.02, 0.2
Std. MPC Hor. N/t_Δ	35 / 0.02
Environment Cluttered Figure Eight-1	
Ellipsoid radii	1.5, 0.9, 1.5, 0.8, 0.9, 1.5, 0.9, 1.5, 0.9, 0.9, 1.5, 1.5, 1.5, 0.9, 1.5, 0.9, 1.5, 0.9, 1.5, 1.5, 1.5, 0.9, 0.9, 0.9, 1.5, 1.5, 1.5, 0.9, 0.9, 0.9, 1.5
Ellipsoid pos.	Equal spacing, random displacement by 1 m in each axis
UNIQUE Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	30, 30 / 0.02, 0.15
Hier. Appr. Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	30, 40 / 0.02, 0.15
Std. MPC Hor. N/t_Δ	50 / 0.02
Environment Cluttered Figure Eight-2	
Ellipsoid radii	1.5, 0.9, 1.5, 0.8, 0.9, 1.5, 0.9, 1.5, 0.9, 0.9, 1.5, 1.5, 1.5, 0.9, 1.5, 0.9, 1.5, 0.9, 1.5, 1.5, 1.5, 0.9, 0.9, 0.9, 1.5, 1.5, 1.5, 0.9, 0.9, 0.9, 1.5, 0.9, 0.9, 0.9
Ellipsoid pos.	Equal spacing, random displacement by 1 m in each axis
UNIQUE Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	30, 30 / 0.02, 0.3
Hier. Appr. Hor. $N, M/t_\Delta, t_\Delta^{\text{lf}}$	30, 30 / 0.02, 0.3
Std. MPC Hor. N/t_Δ	50 / 0.02
Simulation Settings	
Simulation steps	various (completing one lap)
Time step	0.02
Evaluation	
$Q_{\text{eval}}, R_{\text{eval}}$	diag($[1, 1, 1, \mathbf{0}]$) (positions only)

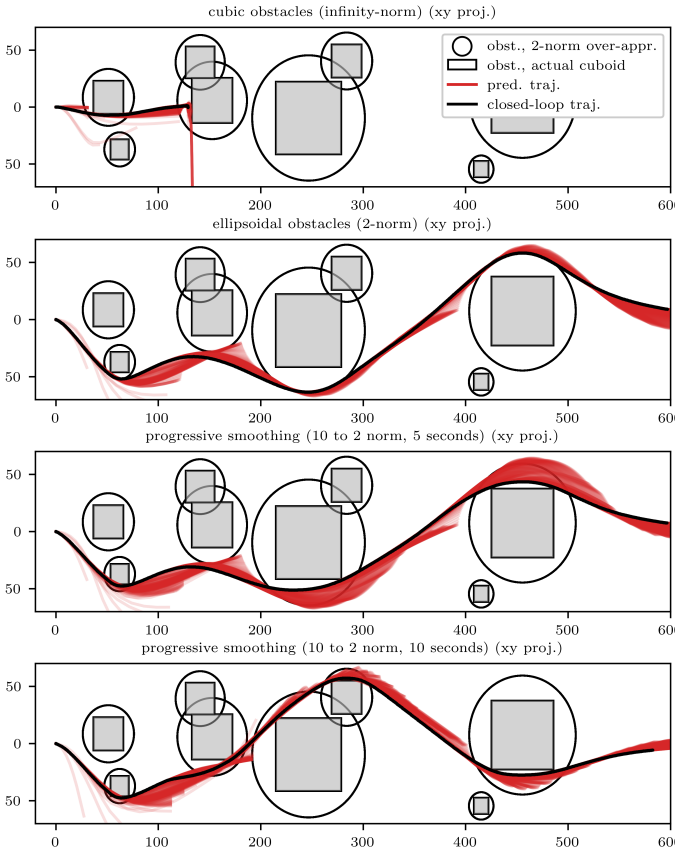

Fig. 12. Comparison of the influence of progressive smoothing in simulated experiments with cubic obstacles. In the first plot, the cubic obstacles are formulated via infinity norms, which leads to nonsmooth shapes and the drone getting stuck in local minima. The second plot shows conservative overapproximations with spheres, where the drone successfully navigates around obstacles, yet at a large distance. The final two plots show different parameterizations of progressive smoothing, which allows tight navigation without getting stuck in local minima.