

Temporal Cascading of Planning and Control for Quadrotor MPC

Rudolf Reiter¹, Chao Qin², Leonard Bauersfeld¹ and Davide Scaramuzza¹

Abstract—Many aerial tasks involving quadrotors demand both instant reactivity and long-horizon planning for obstacle avoidance, energy efficiency, or trajectory tracking. High-fidelity models enable accurate control but are too slow for long horizons. Low-fidelity planners scale but cannot control the system directly, necessitating cascaded architectures. Prevailing hierarchical approaches plan with a simplified model and use a high-fidelity controller for tracking, yet this decomposition is inherently suboptimal. The controller is limited by the coarse plan, and conventional MPC alternatives shorten the horizon to stay real-time feasible. We present UNIQUE, an MPC architecture that replaces this hierarchical stacking with temporal cascading. The planning problem is formulated as the second tail horizon of a single multi-phase MPC, rather than solved separately. We align costs across horizons, derive feasibility constraints for the point-mass planning model, and introduce transition constraints that convert high-fidelity states into meaningful low-fidelity states. Parallel point-mass and mixed-integer solvers address nonconvexities while incorporating progressive 3D obstacle smoothing over the planning horizon. In simulation and real flights, under equal computational budgets, UNIQUE improves closed-loop tracking by up to 75% compared with standard MPC and hierarchical baselines. Ablations and Pareto analyses confirm performance gains across variations in horizon, constraint approximations, and smoothing schedules.

Index Terms—Aerial systems, model predictive control, motion planning, multi-fidelity optimization, quadrotors.

SUPPLEMENTARY MATERIAL

Video available at <https://youtu.be/Ceflqpld9N8>. Supplementary derivations and parameters are provided in an accompanying document.

I. INTRODUCTION

Fast and safe autonomous quadrotor flight depends on two competing requirements: *fast feedback* to disturbances and dynamics, 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, causing the long-standing separation between planning and control for quadrotors [1]–[3].

Hierarchical architectures compute a long-horizon plan using a simplified model, such as a point mass or a polynomial

¹Authors are with the Robotics and Perception Group, Department of Informatics, University of Zurich, Andreasstrasse 15, 8050 Zurich, Switzerland (<https://rpg.ifi.uzh.ch>). Rudolf Reiter is the corresponding author (e-mail: rreiter@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).

²Institute for Aerospace Studies, University of Toronto, 4925 Dufferin Street, Toronto, ON M3H 5T6, Canada.




Fig. 1. Conceptual sketch of UNIQUE. We employ a multi-phase MPC with two horizons: a high-fidelity control horizon and a long horizon with a lower-fidelity point-mass model, enabling long-horizon planning. Local minima of the gradient-based optimizer are avoided by (i) smoothing obstacle shapes along the horizon, and (ii) running parallel planners solely for the computationally 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.

model [1], [4], and track it with an MPC. This hierarchical decomposition is inherently suboptimal. The asynchronously computed low-fidelity plan cannot account for the true vehicle dynamics, leading to (i) performance limitations imposed by the coarse plan, (ii) redundant optimization in the initial horizon, and (iii) conflicting inner or 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 a terminal cost and set that are difficult to obtain for complex scenes involving moving obstacles or nonlinear constraints and are prone to local minima in nonsmooth environments. Remarkably, the planning problem remains NP-hard [5], yet heuristics and tailored algorithms often make it computationally tractable.

Our pivotal idea is to replace the conventional hierarchical planner-controller decomposition with a *temporal cascading* of planning and control inside a single MPC. Rather than computing a plan separately and tracking it, we embed the

planning problem as the second tail horizon of a multi-phase MPC, as sketched in Fig. 1. This is motivated by the fact that an MPC, in general, aims to optimize the same underlying task objective and aims to satisfy the same constraints as the planning problem. We observe that (i) the geometric planning around obstacles is predominantly a long-horizon problem with negligible influence on the short control horizon, and (ii) a point-mass planner can be solved very fast within each closed-loop iteration. Therefore, we formulate the planning problem as finding the optimal tail trajectory on the second horizon, given that we can effectively convert high-fidelity states into meaningful low-fidelity states and map constraints from the high- to the low-fidelity model, which we show to be possible for quadrotors. 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. 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 improves real time iteration (RTI) convergence [7].

The two-phase formulation alone converges well around local minima, but planning problems often involve severe nonconvexities. Therefore, we propose a parallel computation strategy in which we compare the tail trajectory of the main MPC with alternative solutions that may achieve a lower predicted cost. Randomly initialized point-mass solvers run in parallel and, as we show, finish faster than the two-phase MPC, enabling synchronous cost comparison after each RTI step. For harder combinatorial problems, we additionally deploy a potentially slower global planner, in our case a mixed-integer planner as used by many state-of-the-art planning algorithms [8], which first initializes a guaranteed fast parallel point-mass solver. Only the latter's converged cost is compared to the main MPC, ensuring synchronous evaluation. Once the cost of one parallel MPC is lower than the cost along the second horizon of the multi-phase MPC, the solver variables are initialized with the lower-cost solution of the corresponding parallel MPC. In summary, our contributions are as follows:

- Temporal planning paradigm. We propose to formulate the planning problem not hierarchically but as the tail horizon of an MPC, yielding a sequential controller-planner structure within a single optimization.
- Two-phase MPC architecture for quadrotors. 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. We show that high-fidelity states can be effectively converted into meaningful low-fidelity states for quadrotors.
- 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.
- Parallel tail-horizon re-planning. A parallel computation

strategy that compares the tail trajectory of the main MPC against alternative solutions from randomly initialized point-mass solvers.

- Pareto-superior control performance. Under equal computational budgets, the two-phase formulation achieves better closed-loop performance than extending the high-fidelity MPC horizon.
- Comprehensive evaluation. Simulations and real flights show substantial closed-loop gains over standard MPC and hierarchical baselines under equal compute budgets, including Pareto analyses across horizon variations and feasibility approximations, as well as demonstrations of long-horizon planning with iteration times below 5 ms.

The UNIQUE framework lowers the planning-control performance gap by treating planning as a tail-horizon problem of MPC rather than a separate hierarchical stage. 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 work in Sec. II, Sec. 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 the numerical approximations for real-time efficiency. Section VI describes the experimental setup, Section VII reports real-world experiments, and Section VIII presents simulated ablations. We conclude with limitations and future directions in Section IX.

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 [9]–[11] and sampling-based methods such as rapidly exploring random trees (RRTs) [12] 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.

Motion planning with surrogate models. Approximating the quadrotor as a low-fidelity integrator chain [1] and allowing small deviations from its true dynamics [13], [14] 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 [15] 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 [16], [17] without significantly compromising model fidelity.

Beyond time-optimal navigation, a broader line of work has studied planning and control for quadrotors via dynamically feasible point-mass trajectory generation, with a strong focus on the combinatorial aspects of the problem. Safe-flight-corridor methods generate polynomial trajectories inside convex free-space decompositions for real-time flight in cluttered environments [18], while FASTER uses a similar mixed-integer quadratic program (MIQP) formulation to this work to navigate specifically unknown environments, including safe backup trajectories [19]. The authors in [20] propose a very fast planner by eliminating geometric constraints through smooth maps. Classical exact cell decompositions [21] and their modern convex-set variants represent free space as unions or graphs of convex regions, including safe boxes and graphs of convex sets [8], [22], [23] and propose MIQP solvers to provide globally optimal trajectories. As long as obstacles and constraints can be represented as cuboids and quadratic costs are considered, the planners can be used as a backend for our framework. Remarkably, these works demonstrate the strength of planning with simplified or geometry-aware models, but they are still primarily planner-focused and are typically deployed upstream of a separate tracking controller, rather than as a single optimization that jointly allocates high-fidelity near-horizon control and low-fidelity long-horizon planning. One key limitation arises from their inherently smooth, lower-fidelity plans, which cannot account for more complex 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 [15].

Nevertheless, these simplified models form the foundation of the long-horizon component of our unified MPC framework, where we address the limitations mentioned above by introducing a multi-phase structure that precedes the point-mass phase, enforcing feasibility-preserving constraints, and ensuring consistency with high-fidelity quadrotor dynamics. Instead of proposing a competing planning algorithm to the above approaches, we propose a fundamentally different planning paradigm: not hierarchically, but along the second horizon of model predictive control (MPC).

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 [24], [25] or supervised learning for convex quadratic surrogates [26]. If a policy is available, long-term behavior can

be approximated by closed-loop costing [27], [28], including recent neural-policy rollouts [29]. However, these methods rely on pretrained policies or value functions, which may be environment-specific. Other research has studied recursive feasibility under multiple phases [30] and developed multi-phase MPC formulations within modern solvers such as *acados* [31]. Li et al. [32] introduced a dual-model MPC for legged and aerial robots, although it was only tested in offline simulations and did not provide a recursive-feasibility guarantee. The same authors extended the multi-model formulation for quadrupeds [33], addressing the switching dynamics of the legged robot’s gait pattern and providing a tailored iLQR numerical solution scheme for the optimal control problem (OCP). Instead, we adapt the state-of-the-art sequential quadratic programming (SQP) algorithm shown in [31], which provides the necessary efficiency for the state constraints inherent in the targeted motion planning subproblem.

Nonsmooth Obstacles in MPC. Exact collision avoidance can be imposed through smooth dual reformulations without conservative geometric approximations [34], while corridor-based methods use discrete polyhedral corridors [18], [19], smooth differentiable parametric corridors based on ellipses [35], or differentiable collision models for convex primitives [36]. In contrast, the proposed progressive smoothing, derived from [6], is not intended to replace these formulations but to provide a lightweight strategy for norm-based obstacle constraints within a single receding-horizon NLP, improving numerical convergence in cluttered 3D environments while preserving the single-optimization-problem structure. The proposed framework can be adapted to fit specific application domains where other obstacle approximations, such as [35], may be preferred.

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 [37] as a 6-degree-of-freedom rigid body with mass m and a diagonal moment of inertia matrix $J = \text{diag}([J_x \ J_y \ J_z])$.

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 $f_i^\top = [0 \ 0 \ f_{z,i}]$ with $i \in \{1, 2, 3, 4\}$ and the collective z-components of the single rotor thrusts $f_z^\top = [f_{z,1} \ f_{z,2} \ f_{z,3} \ f_{z,4}] \in \mathbb{R}^4$. Thrusts are generated by the rotors with the rotor speed Ω_i and the geometric location $r_{p,i} \in \mathbb{R}^3$ of the i -th rotor in the body frame via $f_{z,i} = \Omega_i^2 c_l$ using the thrust coefficient $c_l \in \mathbb{R}$.

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 $u^{\text{hf}} := f_z \in \mathbb{R}^4$. When now computing the total force $f : \mathbb{R}^4 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ acting on the drone in the body

frame, we add residual force $f_{\text{res}} : \mathbb{R}^3 \times \mathbb{R}^4 \rightarrow \mathbb{R}^3$ to the total thrust, with

$$f(f_z, v_B) := T(f_z) + f_{\text{res}}(v_B, \Omega^2) = T(f_z) + f_{\text{res}}(v_B, \frac{1}{c_l} f_z),$$

where the velocity $v_B \in \mathbb{R}^3$ in the body frame $\mathcal{B}$ can be obtained via $v_B = q^{-1} \odot v$.

The residual force $f_{\text{res}}(v_B, \Omega^2)$ models aerodynamic effects and is obtained by fitting selected features of Ω^2 and velocities v_B to measured data. The literature proposes either higher-order polynomials [37] or linear features of only the velocity v_B [16]. We use a higher-order polynomial for a high-fidelity residual model.

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

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

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

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

With the gravity vector $g^\top = [0 \ 0 \ 9.81 \text{ m/s}^2]$ and the high-fidelity state

$$x = [p^\top \ q^\top \ v^\top \ \omega^\top \ f_z^\top]^\top \in \mathbb{R}^{17},$$

the quadrotor dynamics are finally written as

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

A simplified version of (1) uses linear residual forces, which were shown to be differentially flat [16], cf. the supplementary material.

B. Quadrotor System Constraints

Quadrotors are typically constrained by maximum and minimum rotor speeds, resulting in maximum and minimum single rotor thrusts $\bar{f}_z$ 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 constraints $\bar{\omega}_{xy} \in \mathbb{R}$ on the body rate for roll and pitch, and the total collective thrust $\bar{f}_{\text{th}}$ and $\underline{f}_{\text{th}}$, written as

$$\mathbb{S}_\omega := \{ \omega \in \mathbb{R}^3 \mid |\omega_x| \leq \bar{\omega}_{xy} \wedge |\omega_y| \leq \bar{\omega}_{xy} \}, \quad (2a)$$

$$\mathbb{S}_f := \{ f_z \in \mathbb{R}^4 \mid \underline{f}_{\text{th}} \leq T(f_z) \leq \bar{f}_{\text{th}} \}. \quad (2b)$$

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

$$\mathbb{S}_v = \{ v \in \mathbb{R}^3 \mid \|v_z\| \leq \bar{v}_z \wedge \|[v_x, v_y]\| \leq \bar{v}_{xy} \}.$$

Assumption 1: We assume that the residual force $f_{\text{res}}(v_B, \Omega^2)$ is bounded by $\bar{f}_{\text{res}} \in \mathbb{R}_{\geq 0}$ with $\|f_{\text{res}}(v_B, \Omega^2)\| \leq \bar{f}_{\text{res}}$ for all $v_B \in \mathbb{S}_v$ and $c_l \Omega^2 \in \mathbb{S}_f$.

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 $\theta = (p_o, R_o, d_o, \alpha_o)$ 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 $R_o = [u_o \ v_o \ w_o]$, the obstacle scale parameters $d_o^\top = [d_x \ d_y \ d_z]$, and the norm or shape parameter $\alpha_o \geq 2$. The transformation

$$\eta(p) = [\eta_1(p) \ \eta_2(p) \ \eta_3(p)]^\top := \text{diag}(d_o)^{-1} R_o^\top (p - p_o)$$

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

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

where a parameter $\alpha \geq 2$ defines the shape of the obstacle. The smoothness parameter $\alpha = 2$ corresponds to an ellipsoid and $\alpha \rightarrow \infty$ to a cuboid, see Fig. 2. Note that (3) is a scaled norm, where lower values of α are always over-approximations of larger α , i.e., $2 \leq \alpha_1 \leq \alpha_2 \Rightarrow \mathcal{O}(\alpha_2) \subseteq \mathcal{O}(\alpha_1)$, cf. [6]. Accordingly, d_o denotes the scale before the factor $3^{1/\alpha}$ induced by this normalized norm.

D. Convex Configuration Space Decomposition

The collision constraints are most naturally expressed after a decomposition of the collision-free workspace, generally referred to as *exact cell decomposition* [5], [21]. The general procedure is as follows. The boundary primitives of the inflated obstacles represent a partition of the workspace. For polyhedral obstacles, these primitives are supporting hyperplanes of obstacle facets. For curved obstacles, a conservative polyhedral approximation is used. Intersecting this arrangement with the workspace yields a set of cells whose interiors are collision-free. Each cell is convex whenever it is defined as the intersection of half-spaces from one arrangement region, and adjacent cells that share the same active inequalities can be merged to reduce the overall number of convex cells. The resulting free space $\mathcal{F}$ is represented as

$$\mathcal{F} = \bigcup_{j=1}^{n_c} C_j, \quad C_j = \{ p \in \mathbb{R}^3 \mid H_j p \leq h_j \},$$

where C_j is a convex polytope, n_c are the number of convex sets, and $H_j \in \mathbb{R}^{m_j \times 3}$ and $h_j \in \mathbb{R}^{m_j}$ describe m_j halfspace constraints.

Without loss of generality, we consider axis-aligned box obstacles. After obstacle inflation and clipping to the workspace, all obstacle faces and workspace faces define coordinate-wise cut locations to rectangular voxels. Each elementary box induced by these cuts is classified as free or occupied according to whether its midpoint lies in obstacle-free space; contiguous free boxes are first grouped into maximal connected segments




Fig. 2. Obstacle representations. Visualization of a cubic obstacle smoothed along the high-fidelity horizon (blue) for $\alpha = 100, 8$, and 6 . Along the low-fidelity horizon, cubic obstacles can be further smoothed, as shown here for $\alpha = 6$ to 4 . The mixed-integer planner requires a polyhedral over-approximation due to the convex decomposition of the free planning space.

along one coordinate direction and are then merged across shared faces in the remaining directions until no further convex merging is possible. Because all cuts are axis-aligned, each final cell is an axis-aligned box

$$C_j = \{p \in \mathbb{R}^3 \mid \ell_j \leq p \leq u_j\},$$

which is a particularly simple polyhedral representation of the exact cell decomposition induced by the inflated obstacles. Additionally, a cell adjacency graph is constructed: two cells C_i and C_j are adjacent if their axis-aligned bounding boxes touch or overlap on every coordinate axis.

a) *Mixed-integer encoding.*: Let $\beta_{k_b, j} \in \{0, 1\}$ denote whether the trajectory at binary step k_b is assigned to cell C_j . To reduce the number of binary variables, cell-assignment binaries are introduced only at a subset of nodes $\mathcal{K}_b = \{s, 2s, \dots\} \cup \{N\} \subseteq \{1, \dots, N\}$ with stride $s \geq 1$. Each node k inherits the assignment from the nearest preceding binary step $\kappa(k) = \max\{k_b \in \mathcal{K}_b : k_b \leq k\}$, so that the total number of binary variables is $|\mathcal{K}_b| n_c$ instead of $N n_c$. Exactly one cell is selected at each binary step,

$$\sum_{j=1}^{n_c} \beta_{k_b, j} = 1, \quad k_b \in \mathcal{K}_b, \quad (4)$$

and the position at every node is forced into the selected cell via indicator constraints,

$$\beta_{\kappa(k), j} = 1 \Rightarrow \ell_j \leq p_k \leq u_j, \quad j = 1, \dots, n_c, \quad k = 1, \dots, N. \quad (5)$$

To ensure spatial consistency across the horizon, cell transitions between consecutive binary steps are restricted to neighboring cells in the adjacency graph,

$$\sum_{i \in \mathcal{N}(j) \cup \{j\}} \beta_{k'_b, i} \geq \beta_{k_b, j}, \quad j = 1, \dots, n_c, \quad (6)$$

for each pair of successive binary steps $k_b, k'_b \in \mathcal{K}_b$, where $\mathcal{N}(j)$ denotes the set of cells adjacent to C_j .

b) *Online pruning.*: Before each solve, conservative per-axis reachable position bounds are propagated from the current state x_0 over the planning horizon. At each binary step k_b , cells whose bounding boxes do not overlap the reachable region are fixed to $\beta_{k_b, j} = 0$, which substantially reduces the effective number of active binary variables without altering feasibility. For node p_k and binary variables $\beta_{\kappa(k)}^\top = [\beta_{\kappa(k), 1}, \dots, \beta_{\kappa(k), n_c}]$ the constraints (4)–(6) and the reachability fixing are summarized by $(P_p^{\text{lf}} x_k^{\text{lf}}, \beta_{\kappa(k)}) \in \mathcal{F}$.

E. Nominal MPC

With sufficient computational resources and a sufficiently good initial guess, we would solve a nominal MPC problem using the highest-fidelity model available. We formulate a multiple-shooting nonlinear program (NLP) to solve the MPC problem with state variables $X = [x_0 \dots x_M]$ and control variables $U = [u_0 \dots u_{M-1}]$, 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 can establish recursive feasibility under the standard terminal-set invariance assumptions. 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_z , v and ω of the full state x , the nominal MPC problem is defined as

$$\min_{X, U} \sum_{k=0}^{M-1} l(x_k, u_k) + \Phi(x_M) \quad (7a)$$

s.t.

$$x_0 = \hat{x}_0, \quad (7b)$$

$$x_{k+1} = F(x_k, u_k; t_\Delta), \quad k \in \mathbb{N}_{[0, M-1]}, \quad (7c)$$

$$P_\omega x_k \in \mathbb{S}_\omega, \quad P_f x_k \in \mathbb{S}_f, \quad P_v x_k \in \mathbb{S}_v \quad k \in \mathbb{N}_{[0, M-1]}, \quad (7d)$$

$$P_p x_k \notin \mathcal{O}, \quad k \in \mathbb{N}_{[0, M-1]}, \quad (7e)$$

$$0 \geq h_M(x_M). \quad (7f)$$

We define the nonlinear least-squares cost

$$l(x, u) = t_\Delta (y(x, u) - \tilde{y})^\top \text{diag}(w) (y(x, u) - \tilde{y}), \quad (8)$$

with the tracking states $y^\top(x, u) = [p \ \phi(q, \tilde{q}) \ v \ \omega \ f_z \ u] \in \mathbb{R}^{20}$ that includes the quaternion error $\phi: \mathbb{R}^{4 \times 4} \rightarrow \mathbb{R}^3$ 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], [15], [32]. 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]. For 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

$$\dot{x}^{\text{lf}} = [\dot{p} \quad \dot{v} \quad \dot{a}]^\top = [v \quad a \quad j]^\top := f^{\text{lf}}(x^{\text{lf}}, u^{\text{lf}}), \quad (9)$$

with the low-fidelity state $x^{\text{lf}} = [p \quad v \quad a]^\top$.

A. Relation to High-Fidelity Model

Due to the differential flatness of the quadrotor model [16], all quadrotor states of a simple model with linear residuals [16] 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

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

and the jerk via the derivative of the acceleration

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

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

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

with the scalar collective thrust $T_z(f_z) := e_z^\top T(f_z) = \sum_{i=1}^4 f_{z,i} \geq 0$. Taking its derivative gives $\dot{T}_z(f_z) = \sum_{i=1}^4 \dot{f}_{z,i}$, and the jerk follows as

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

B. Point-Mass System Constraints

The authors in [4] derived constraints

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

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 $\bar{f}_{\text{res}}$ 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 $\mathbb{S}_f$ 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

$$\mathbb{S}_f^{\text{lf}} := \{a \in \mathbb{R}^3 \mid \underline{f}_{\text{th}} + \bar{f}_{\text{res}} \leq m\|a + g\| \leq \bar{f}_{\text{th}} - \bar{f}_{\text{res}}\}. \quad (10)$$

Body rate constraints $\mathbb{S}_\omega$ formulated in (2a) 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.,

$$\left\| \frac{d}{dt} \left(\frac{a + g}{\|a + g\|} \right) \right\| \leq \bar{\omega}_{xy}. \quad (11)$$

Resolving the time derivative in (11) and utilizing the explicit notation of jerk with $j = \dot{a} + \dot{g} = \dot{a}$, a conservative sufficient body rate feasibility constraint for the low-fidelity model can be written as

$$\mathbb{S}_\omega^{\text{lf}} := \left\{ \begin{bmatrix} a \\ j \end{bmatrix} \in \mathbb{R}^6 \mid \frac{\|j\|}{\|a + g\|} \leq \bar{\omega}_{xy} \right\}. \quad (12)$$

The auxiliary constraints on the thrust (10) and the body rate (12) 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 $\mathbb{S}_f^{\text{lf}}$ defined in (10). The lower bound of (10) is satisfied if the simpler inequality

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

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 (10) can be simplified, where a more restrictive constraint guarantees the feasibility of (10). The simplified constraint involves two design parameters $\alpha_x, \alpha_z \in (0, 1)$ that shape and trade off the more optimization-friendly constraints. Define the available thrust margin as $\bar{f}_{\text{av}} := \bar{f}_{\text{th}} - \bar{f}_{\text{res}}$. By utilizing the design parameters, box-constraints for $\bar{a}^\top = [\bar{a}_x \quad \bar{a}_y \quad \bar{a}_z]$ can be computed iteratively by

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

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

$$\hat{\mathbb{S}}_f^{\text{lf}} := \{a \in \mathbb{R}^3 \mid [-\bar{a}_x \quad -\bar{a}_y \quad \underline{a}_z] \leq a \leq \bar{a}\} \subset \mathbb{S}_f^{\text{lf}}, \quad (13)$$

which is a subset of the thrust constraints (10).

Next, we simplify the constraints of $\mathbb{S}_\omega^{\text{lf}}$, defined in (12). A more optimization-friendly subset of (12) can be found by exploiting the lower bound on the acceleration

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

to obtain simplified convex constraints on the jerk j

$$\hat{\mathbb{S}}_\omega^{\text{lf}} := \{j \in \mathbb{R}^3 \mid \|j\| \leq (\underline{a}_z + g)\bar{\omega}_{xy}\} \subseteq \mathbb{S}_\omega^{\text{lf}}. \quad (14)$$

Arbitrary box constraints for each individual component of the jerk can be computed, forming a subset of (14) and, therefore, decoupling the constraints. A valid choice would be $\bar{j}_{xyz} = \frac{\underline{a}_z + g}{\sqrt{3}}\bar{\omega}_{xy}$ to obtain the convex decoupled box constraints

$$\hat{\mathbb{S}}_\omega^{\text{lf}} := \{j \in \mathbb{R}^3 \mid |j_{\{x,y,z\}}| \leq \bar{j}_{xyz}\} \subseteq \hat{\mathbb{S}}_\omega^{\text{lf}}. \quad (15)$$

V. UNIQUE: THE TEMPORAL CASCADING MPC FORMULATION

In the following, we propose UNIQUE, a multi-phase MPC framework that embeds the planning problem as the tail horizon of the controller. Rather than solving the planning problem hierarchically on a separate model and tracking the result, we append the low-fidelity point-mass model from Sect. IV as a second phase after the high-fidelity MPC from Sect. III-E, forming a single optimization that jointly plans and controls. The merging of two different models along the MPC horizons requires (i) the alignment of the cost function, (ii) recursive feasible constraint formulations, (iii) a 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 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 over the horizon are detailed below.

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 (8) 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

$$l^{\text{lf}}(x^{\text{lf}}, u^{\text{lf}}) = t_\Delta^{\text{lf}}(y^{\text{lf}} - \bar{y}^{\text{lf}})^\top \text{diag}(w^{\text{lf}})(y^{\text{lf}} - \bar{y}^{\text{lf}}), \quad (16)$$

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. In particular, the weights w_p and w_v are associated with the same physical states and are therefore chosen to be identical.

In the high-fidelity cost (8), single-rotor vertical thrusts are penalized by a weighted sum of squares. For equal rotor-thrust component weights $w_{f,0}$ and equal sharing of a collective-thrust variation, $\sum_{i=1}^4 (\Delta f_{z,i})^2 = m^2 \|\Delta a\|^2 / 4$. We therefore use the approximate acceleration-weight alignment $w_a := m^2 w_{f,0} / 4$; attitude, gravity, and residual forces prevent an exact global equivalence.

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

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

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

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

B. Constraint Alignment

As shown in Sect. IV-B, we can find constraint sets $\mathbb{S}_\omega^{\text{lf}}$ and $\mathbb{S}_f^{\text{lf}}$ 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 $\mathbb{S}_\omega^{\text{lf}}$ that are more tightly constrained yet numerically easier to solve. The set $\hat{\mathbb{S}}_\omega^{\text{lf}}$ assumes a lower bound $\underline{a}_z$, leading to a convex quadratic constraint on the jerk j (14) and the set $\hat{\mathbb{S}}_\omega^{\text{lf}}$ decouples the jerk components to box constraints (15). 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 $t \geq Mt_\Delta$ can be tracked. Our approach explicitly couples all derivatives of the low-fidelity trajectory, up to and including jerk, to the corresponding high-fidelity states. First, we require the state and velocities of both fidelity models to be aligned via

$$\pi_{pv}(x, x^{\text{lf}}) := \begin{bmatrix} P_p x \\ P_v x \end{bmatrix} - \begin{bmatrix} P_p^{\text{lf}} x^{\text{lf}} \\ P_v^{\text{lf}} x^{\text{lf}} \end{bmatrix} = 0. \quad (17)$$

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

$$\pi_a(x, x^{\text{lf}}) := a_{\text{quad}}(P_q x, P_v x, P_f x) - P_a x^{\text{lf}} = 0. \quad (18)$$

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 gravity-shifted acceleration $s(a) := a + g$ and the jerk $j = u^{\text{lf}}$ via

$$\begin{bmatrix} \omega_x & -\omega_y & 0 \end{bmatrix}^\top = q \odot \left(\frac{j}{\|s(a)\|} - \frac{s(a)s(a)^\top j}{\|s(a)\|^3} \right). \quad (19)$$

With $s_{\text{lf}} := P_a x^{\text{lf}} + g$, this leads to the constraint

$$\pi_\omega \left(x, x^{\text{lf}}, u^{\text{lf}} \right) := \text{diag} \left([1, -1, 0]^\top \right) P_\omega x - P_q x \odot \left(\frac{u^{\text{lf}}}{\|s_{\text{lf}}\|} - \frac{s_{\text{lf}} s_{\text{lf}}^\top u^{\text{lf}}}{\|s_{\text{lf}}\|^3} \right) = 0. \quad (20)$$

We require $\|a+g\| > 0$ in (19) and (20). This follows from the positive lower collective-thrust bound in (10). Additionally, the jerk-related constraint (20) 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 the feasibility of the quadrotor model. Finally, we summarize the relevant coupling conditions of (17) and (18) by

$$\pi \left(x, x^{\text{lf}} \right) := \left[\pi_{pv} \left(x, x^{\text{lf}} \right) \quad \pi_a \left(x, x^{\text{lf}} \right) \right]^\top = 0.$$

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 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 smoothed along the MPC planning horizon by scheduling the smoothing parameter α towards a smooth 2-norm. Particularly, a monotonically decreasing scheduling function $\lambda(t)$ with $\lambda : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ 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 improves numerical behavior of the RTI scheme [7]. Recursive feasibility additionally requires the usual shift, terminal-set, and hard-constraint assumptions.

E. Multi-Phase MPC formulation

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

Main Multi-Phase MPC

$$\min_{\substack{X, U, \\ X^{\text{lf}}, U^{\text{lf}}}} \sum_{k=0}^{M-1} l(x_k, u_k) + \sum_{k=0}^{N-1} l^{\text{lf}}(x_k^{\text{lf}}, u_k^{\text{lf}}) + \Phi^{\text{lf}}(x_N^{\text{lf}}) \quad (21a)$$

s.t.

$$x_0 = \hat{x}_0, \quad (21b)$$

$$0 = \pi \left(x_M, x_0^{\text{lf}} \right), \quad (21c)$$

$$x_{k+1} = F \left(x_k, u_k; t_\Delta \right), \quad k \in \mathbb{N}_{[0, M-1]}, \quad (21d)$$

$$x_{k+1}^{\text{lf}} = F^{\text{lf}} \left(x_k^{\text{lf}}, u_k^{\text{lf}}; t_\Delta^{\text{lf}} \right), \quad k \in \mathbb{N}_{[0, N-1]}, \quad (21e)$$

$$P_f x_k \in \mathbb{S}_f, P_\omega x_k \in \mathbb{S}_\omega, P_v x_k \in \mathbb{S}_v \quad k \in \mathbb{N}_{[0, M-1]}, \quad (21f)$$

$$P_a x_k^{\text{lf}} \in \mathbb{S}_f^{\text{lf}}, [P_a x_k^{\text{lf}}, u_k^{\text{lf}}] \in \mathbb{S}_\omega^{\text{lf}}, \quad k \in \mathbb{N}_{[0, N-1]}, \quad (21g)$$

$$P_p x_k \notin \mathcal{O} \left(\lambda(k t_\Delta) \right), \quad k \in \mathbb{N}_{[0, M-1]}, \quad (21h)$$

$$P_p^{\text{lf}} x_k^{\text{lf}} \notin \mathcal{O} \left(\lambda(M t_\Delta + k t_\Delta^{\text{lf}}) \right), \quad k \in \mathbb{N}_{[0, N]}, \quad (21i)$$

$$0 = h^{\text{lf}} \left(x_N^{\text{lf}} \right). \quad (21j)$$

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 $\pi(x_M, x_0^{\text{lf}}) = 0$ 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 $\mathbb{S}_f^{\text{lf}}, \mathbb{S}_\omega^{\text{lf}}$ for the low-fidelity model. Both horizons are kept collision-free by enforcing position constraints against the obstacle set $\mathcal{O}(\lambda(\cdot))$, whose recursive feasibility is guaranteed [6]. Finally, the terminal equality $h^{\text{lf}}(x_N^{\text{lf}}) = 0$ 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 (21) is still prone to getting stuck in local minima due to the long planning horizon and the gradient-based local optimization of MPC solvers [38]. Dealing with such nonconvexities is the pivotal point of many planning algorithms discussed in Section II. Since the planning-relevant decision variables are tied dominantly to the second (tail) phase of our formulation, we can exploit this structure. A key design goal is to keep the main controller fast: the high-fidelity first phase must not be delayed by the combinatorial complexity of planning. Therefore, we offload the search for better tail trajectories to parallel asynchronous solvers that run concurrently with the main MPC. Formulating this subproblem as a separate NLP allows us to optimize parallel instances from different initial guesses and escape local minima of the dominant second planning horizon while sacrificing only minor computational resources. As an initialization strategy, we roll out $M_{p, \text{ini}} = 64$ trajectories from $\hat{x}_0^{\text{lf}}$ by simulating the low-fidelity model with an RK4 integrator and applying random jerks sampled from a uniform distribution $j \sim \mathcal{U}(j, \bar{j})$. We then select N_p candidates, prioritizing feasible trajectories and otherwise using low-cost infeasible ones.

To guarantee continuity, we start from the same initial point-mass state $\hat{x}_0^{\text{lf}}$ obtained from the multi-phase MPC (21); the parallel MPC subproblems are

Parallel Point-Mass MPC

$$J_{\text{lf}}(x_0^{\text{lf}}) = \min_{X^{\text{lf}}, U^{\text{lf}}} \sum_{k=0}^{N-1} l^{\text{lf}}(x_k^{\text{lf}}, u_k^{\text{lf}}) + \Phi^{\text{lf}}(x_N^{\text{lf}}) \quad (22)$$

s.t.

$$x_0^{\text{lf}} = \hat{x}_0^{\text{lf}},$$

$$x_{k+1}^{\text{lf}} = F^{\text{lf}} \left(x_k^{\text{lf}}, u_k^{\text{lf}}; t_\Delta^{\text{lf}} \right), \quad k \in \mathbb{N}_{[0, N-1]},$$

$$P_a x_k^{\text{lf}} \in \mathbb{S}_f^{\text{lf}}, [P_a x_k^{\text{lf}}, u_k^{\text{lf}}] \in \mathbb{S}_\omega^{\text{lf}}, \quad k \in \mathbb{N}_{[0, N-1]},$$

$$P_p^{\text{lf}} x_k^{\text{lf}} \notin \mathcal{O} \left(\lambda(M t_\Delta + k t_\Delta^{\text{lf}}) \right), \quad k \in \mathbb{N}_{[0, N]},$$

$$0 = h^{\text{lf}} \left(x_N^{\text{lf}} \right).$$

Crucially for planning problems with severe nonconvexities, we add a particular parallel point-mass MPC variant that explicitly considers the nonconvexities in the free planning space by a decomposition into convex cells linked via binary indicator variables. The resulting MIQP formulation requires a dedicated solver.

Parallel Mixed-Integer MPC

$$\begin{aligned}
 J_{\text{mi}}(x_0^{\text{lf}}) = & \min_{X^{\text{lf}}, U^{\text{lf}}, B} \sum_{k=0}^{N-1} l^{\text{lf}}(x_k^{\text{lf}}, u_k^{\text{lf}}) + \Phi^{\text{lf}}(x_N^{\text{lf}}) \\
 \text{s.t.} \quad & x_0^{\text{lf}} = \hat{x}_0^{\text{lf}}, \\
 & x_{k+1}^{\text{lf}} = F^{\text{lf}}(x_k^{\text{lf}}, u_k^{\text{lf}}; t_{\Delta}^{\text{lf}}), \quad 0 \leq k < N, \\
 & P_a x_k^{\text{lf}} \in \mathbb{S}_f^{\text{lf}}, \quad 0 \leq k < N, \\
 & [P_a x_k^{\text{lf}}; u_k^{\text{lf}}] \in \mathbb{S}_{\omega}^{\text{lf}}, \quad 0 \leq k < N, \\
 & (P_p^{\text{lf}} x_k^{\text{lf}}, \beta_{\kappa(k)}) \in \mathcal{F}, \quad 0 \leq k \leq N, \\
 & 0 = h^{\text{lf}}(x_N^{\text{lf}}).
 \end{aligned} \tag{23}$$

Due to the feasibility of the initial state from (21), all MPC subproblems provide recursive feasibility when combined with the high-fidelity horizon of (21). A proof under the standard shift and terminal-set assumptions is provided in the supplementary material. Once any of the MPC subproblems initialized with random primal variables $(^0X^{\text{lf}}, ^0U^{\text{lf}})^i$ exhibits a feasible lower cost $J_{\text{lf}}((^0X^{\text{lf}}, ^0U^{\text{lf}})^i, \hat{x}_0^{\text{lf}})$ than the second horizon of (21), the second horizon of (21) is reinitialized with the optimal lower cost decision variables of i -th parallel MPC subproblem. Therefore, the accepted initialization has a lower evaluated tail cost before the next RTI step. Since the first horizon is unchanged during reinitialization, this reduces the evaluated objective at that instant; subsequent RTI steps and closed-loop shifts need not preserve a monotonic decrease. See [24] 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 (21) numerically efficiently, we use the solver ACADOS [38] with the interior-point quadratic program (QP) solver HPIPM [39] without condensing, and deploy the RTI scheme [7]. We do not use condensing, since for the considered problem size we did not observe an improvement. We use slack variables with exact L_1 and quadratic L_2 penalties [40] for obstacle avoidance constraints and approximate convex quadratic constraints of the sets $\hat{\mathbb{S}}_f^{\text{lf}}$ and $\hat{\mathbb{S}}_{\omega}^{\text{lf}}$ via a polyhedron $\{x \in \mathbb{R}^3 \mid b_p \leq A_p x\}$, with $b_p \in \mathbb{R}^{12}$ and $A_p \in \mathbb{R}^{12 \times 3}$. In particular, we use a dodecahedron with 12 surfaces instead of the six surfaces of a cube approximation as in $\hat{\mathbb{S}}_{\omega}^{\text{lf}}$ and $\hat{\mathbb{S}}_f^{\text{lf}}$ to reduce the conservatism of the inner approximation by adding more constraints, as shown in Fig. 3. The approximation via a polyhedron allows the QP solver to safely handle the norm constraints, which is not possible when linearizing the constraints within a single QP subroutine of SQP. Notably, the velocity constraint $\mathbb{S}_v$ spans the whole

Algorithm 1 UNIQUE: Unified Multi-Fidelity MPC

- 1: **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
- 2: Set $\alpha_i = \lambda(t_i)$ at each prediction time t_i , $i \in I_0$
- 3: Decompose obstacles $\{\mathcal{O}_i\}$ to free space $\mathcal{F}$
- 4: **Repeat at each control step k :**
- 5: **1. Obstacles update:**
- 6: Set obstacle param. θ_i for prediction steps $i \in I_k$
- 7: **2. Main multi-phase MPC:**
- 8: Solve QP of multi-phase MPC at $x_k = \hat{x}_k$
- 9: Evaluate second horizon cost: $(J_{\text{lf},k})_0$
- 10: Get initial l.f. state $\hat{x}_0^{\text{lf}} \leftarrow x_0^{\text{lf}}$
- 11: **3. Parallel Mixed-Integer Planner:**
- 12: Preprocess binary variables
- 13: Solve MIQP of point-mass planner at $(x_0^{\text{lf}})_{\text{mi}} := \hat{x}_0^{\text{lf}}$
- 14: Evaluate second horizon cost: $(J_{\text{lf},k})_{\text{mi}}$
- 15: **4. Parallel Point-Mass MPCs:**
- 16: Solve $j = 1 : S$ point-mass MPCs at $(x_0^{\text{lf}})_j := \hat{x}_0^{\text{lf}}$
- 17: Evaluate parallel point-mass MPC costs: $(J_{\text{lf},k})_j$
- 18: **5. Evaluate & Reinitialize:**
- 19: Get lowest cost $\lambda^* = \arg \min_{\lambda \in 0, \dots, S, \text{mi}} (J_{\text{lf},k})_{\lambda}$
- 20: Reinitialize multi-phase MPC
- 21: $(x_i^{\text{lf}})_0 \leftarrow (x_i^{\text{lf}})_{\lambda^*} \quad \forall i \in \{0, \dots, N\}$
- 22: $(u_i^{\text{lf}})_0 \leftarrow (u_i^{\text{lf}})_{\lambda^*} \quad \forall i \in \{0, \dots, N-1\}$
- 23: **if $k \bmod P = 0$ then**
- 24: Randomly reini. parallel p.m. MPCs $j = 1 : S-1$
- 25: Reini. S -th parallel p.m. MPC with MIQP solution
- 26: **6. Apply and shift:**
- 27: Execute first main multi-phase control u_0
- 28: Shift primal variables of all MPCs




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

practical operating region of the quadrotor and can therefore be removed in (21). Remarkably, we normalize quaternions before using them to initialize the MPC problem, but do not explicitly constrain them to the unit sphere within the MPC, as we observed convergence problems and no performance benefit for converged solutions.

VI. EXPERIMENTAL SETUP

The proposed algorithm is evaluated in both real-world and simulation environments using a realistic model that accounts for aerodynamic effects, as detailed in [41] and the supplementary material. For both real-world experiments and simulations,

TABLE I
SPECIFICATIONS OF THE “OFFBOARD” QUADROTOR

Parameter		Value	Unit
Geometry	$[r_{p,i}]_x$	± 7.5	cm
	$[r_{p,i}]_y$	± 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	μNs^2
	κ	11.0	mm
Quad. Diag.	l	38.2	cm
Max Thrust	$\bar{f}_{th}$	34.0	N
Max Rate	$\bar{\omega}_{xy}$	10.0	s ⁻¹
	$\bar{\omega}_z$	6.0	s ⁻¹



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 [42] lowest-level onboard controller, which tracks input body rates and total thrust. Therefore, the total thrust obtained from the rotor-thrust states 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 $25 \times 12 \times 4$ m. In the final evaluations, we employ three distinct experimental settings: the *constant velocity environment*, the *tracking environment*, and the *motion-planning maze environment*, which are described below. In these experiments, the obstacles can be smooth ellipses or nonsmooth 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 $15 \frac{\text{m}}{\text{s}}$ at a specific height. The environment evaluates how early adaptation towards long-horizon prediction and short-horizon adaptation due to moving obstacles affect the closed-loop performance of maintaining a reference speed. More details on the environment are given in the supplementary material.

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 trajectory resembles shapes such as a triangle, a rounded rectangle, a figure-eight, a butterfly, or a sinusoidal wave. The tracking environment cost emphasizes the position tracking of the reference and, due to its periodic flight structure, can be readily evaluated in real-world experiments. More details on the environment are given in the supplementary material.

Motion Planning Maze Environment. The motion planning environments require the consideration of major nonconvexities that cannot be handled by smoothing alone. The goal is to traverse a 6×6 m tunnel at a target speed of $1 \frac{\text{m}}{\text{s}}$ with

stage costs defined by $x_k^\top Q x_k$ and $u_k^\top R u_k$. The tunnel has a length of 17 m for real-world experiments and 40 m for simulated experiments. We place up to 14 random axis-aligned walls, see Fig. 14, thereby forming a maze-like configuration space.

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

$$J^{\text{eval}}(\hat{X}, \hat{U}) := t_{\Delta, \text{sim}} \sum_{k=0}^{N_{\text{sim}}-1} (\|\hat{x}_k - \tilde{x}_k\|_{2, Q^{\text{eval}}} + \|\hat{u}_k - \tilde{u}_k\|_{2, R^{\text{eval}}}). \quad (24)$$

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

VII. REAL-WORLD EXPERIMENTS

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 set to the corresponding phases of UNIQUE for a fair comparison, except for the position-tracking weight, which is adapted for the hierarchical lower-level tracking MPC.

The 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 within the proposed initialization strategy to avoid local minima. In the first comparison, we consider a use case with many obstacles and few local minima, so random initial guesses can escape them reliably.

In Sect. VII-B, we compare the three variants in the maze environment with randomly placed horizontal and vertical walls.

A. Real-World Tracking 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 15×8 m Agile Sinusoidal track with seven obstacles, a period of 10 s and accelerations of $20 \frac{\text{m}}{\text{s}^2}$, (ii) the flat 15×8 m Agile Butterfly




Fig. 4. Comparison between UNIQUE, the standard MPC formulation, and the hierarchical 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.

track with seven obstacles and accelerations of $35 \frac{\text{m}}{\text{s}^2}$, (iii) the flat $15 \times 8 \text{ m}$ Cluttered Figure-Eight-Track-1 with 30 obstacles and a period of 40s 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 tracking error of the standard MPC formulation in cluttered environments stems from the abundance of local minima, which 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 improve accuracy but also avoid 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 per-

TABLE II
TRACKING ERROR 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%)

formance of UNIQUE with progressive smoothing (mean/max tracking error of 0.6/1.4 m), compared to the conservative L2-norm approximation (mean/max tracking error of 0.8/2.0 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 are shown in the supplementary material.

Parallel Point-Mass MPC. We evaluate the influence of a 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 $\max(t_{\text{comp},0}, \dots, t_{\text{comp},S})$, where $t_{\text{comp},0}$ is the computation time of the multiphase MPC and $t_{\text{comp},1}, \dots, t_{\text{comp},S}$ are the computation times of the parallel point-mass MPCs. The mean and median computation times do not increase in our experiments compared to the multi-phase MPC without parallelization. However, the maximum computation time is increased from 10.9 ms to 21.9 ms, which is still below the benchmark approaches compared, cf. Fig. 5 and Tab. II. Considering the tracking error, the difference when using parallel point-mass MPCs is remarkable in this obstacle-rich environment, as the mean tracking error decreases by 77.1% from 3.5/2.0/8.3 (mean/median/max, in meters) 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.




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.

B. Real-World Planning Comparison

We evaluate the three navigation algorithms (std., hier., UNIQUE) in a maze environment built from virtual obstacles in a flying arena equipped with a Vicon motion-capture system. As before, the hierarchical framework uses the same hyperparameters as the second horizon of UNIQUE to provide a reference for a tracking MPC, whose hyperparameters match the first horizon of UNIQUE except for a stronger position-tracking cost. We start at the position $p_0 = [-9, 0, 1]$ m and aim to traverse the tunnel in the positive x direction until $x_T \geq 12$ m. Various obstacles block the tunnel, requiring planners with an 8-second horizon to find low-cost detours, as shown in Fig. 8. The standard MPC formulation is excluded from the comparison because it does not finish. Figure 8 shows rapid cost increases for the hierarchical planner when a parallel point-mass planner finds a lower-cost trajectory that deviates from the previous plan and the tracking controller attempts to follow it. This replanning is much smoother when it is integrated into the second horizon of UNIQUE. Scene 2 reveals a failure case for UNIQUE: the parallel planners found a lower-cost trajectory around the blocking obstacle, but the slack variables were violated, preventing switching. As shown in Fig. 9, this increased the accumulated cost over the baseline. In the other scenarios, however, the sequential planning framework reduced the cost.

VIII. SIMULATED EXPERIMENTS

In the following simulation experiments, we ablate different hyperparameters in the constant-velocity and maze environments to assess the robustness and scaling properties of UNIQUE. First, we focus on control-oriented evaluations, emphasizing the closed-loop performance gains, feasibility, and computation time achieved via the multi-phase formulation. In Sect. VIII-B, we evaluate the different constraint approximations of the second horizon and provide some qualitative insights into progressive smoothing and the parallel initialization strategy. Secondly, we focus on the performance gain achieved via the parallel planning framework using the point-mass and MIQP planners. Specifically, in Sect. VIII-C, we evaluate how the computational burden scales with the number of obstacles in the mixed-integer planner, how different combinations of parallel solvers influence performance, and how UNIQUE compares against the hierarchical formulation when both use equal planners.

A. Control-Focused 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




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 a 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 in a particular snapshot.




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 is 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. Real-world evaluation showing trajectories and cumulative closed-loop costs on four randomly generated scenes for the hierarchical controller and UNIQUE.




Fig. 9. Real-world evaluation on four randomly generated scenes. UNIQUE achieves a 70.8%, 40.9%, and 28.7% lower accumulated closed-loop cost than the hierarchical approach in three scenarios, and a 10.9% higher cost in one scenario, corresponding to a 34.2% reduction after aggregating the costs across all four scenarios.

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 $\hat{\mathbb{S}}_\omega^{\text{lf}}$. 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 closed-loop performance gains? Fig. 10 shows a comparison of the Pareto front of closed-loop performance and online computation time for the RTI scheme and the fully converged solvers, with parameters shown in Tab. III. The comparison includes a standard MPC with increasing spacing of the high-fidelity horizon, which constitutes an alternative method for planning more coarsely 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. 10 and also, the hierarchical approach clusters in a specific region. Both benchmarks are clearly outperformed by UNIQUE.

How do hyperparameters affect convergence? Although the




Fig. 10. Pareto comparison of closed-loop cost and average per-iteration computation time 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.

RTI scheme is the only applicable approach in real-world evaluations, the fully converged evaluation shown in Fig. 10 and Tab. III contrasts the different approaches in terms of their numerical complexity. Across the reported configurations, full convergence increases computation time by factors of approximately 6–7 for UNIQUE and 29–137 for standard MPC relative to RTI. Thus, UNIQUE requires substantially less additional time for full convergence than the long-horizon standard MPC configurations.

How does the second phase influence control performance? In Fig. 11, we investigate how adding steps of the second phase influences the Pareto performance in comparison to how adding standard MPC steps influences the performance. We take standard MPC with $n_{hf} = \{10, 15, 20, 30, 50\}$ shooting nodes and $t_{\Delta}^{hf} = 40$ ms (red line in Fig. 11) and




Fig. 11. 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.

append to each standard high-fidelity MPC the low-fidelity part of $n_{lf} = \{3, 5, 7\}$ with $t_{\Delta}^{lf} = 400$ 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. This confirms that, even from a pure control perspective, the two-phase structure yields superior closed-loop performance under the same computational budget.

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		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1853.74 / 1717.60	3.88 / 24.85	
▼ UNIQUE-1	0.52 / 0.04	13		17.60 / 0.80	22	0.10		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1853.74 / 1717.60	1.50 / 9.84	
◀ UNIQUE-2	0.52 / 0.04	13		16.00 / 0.80	20	0.20		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1875.22 / 1720.68	1.54 / 9.88	
▶ UNIQUE-3	0.80 / 0.04	20		7.20 / 0.80	9	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1331.95 / 1203.88	2.50 / 16.15	
▼ UNIQUE-4	0.80 / 0.04	20		15.20 / 0.20	76	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1865.09 / 1734.07	12.05 / 69.26	
▲ UNIQUE-5	0.80 / 0.04	20		15.20 / 0.20	76	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1865.12 / 1734.09	4.08 / 24.59	
■ UNIQUE-6	0.80 / 0.04	20		3.20 / 0.20	16	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1664.14 / 1535.86	1.81 / 11.64	
◆ UNIQUE-7	0.80 / 0.04	20		3.20 / 0.20	16	0.00		$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	1665.23 / 1536.51	3.42 / 20.21	
Controller				Planner							
● Hier.-0	2.00 / 0.04	50		24.00 / 0.20	120		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2688.66 / 2535.47	7.47 / 28.98	
▼ Hier.-1	2.00 / 0.04	50		24.00 / 0.40	60		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	3112.53 / 2957.01	5.04 / 27.06	
◀ Hier.-2	2.00 / 0.04	50		8.00 / 0.20	40		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	7170.31 / 7019.67	3.08 / 29.21	
▶ Hier.-3	2.00 / 0.04	50		8.00 / 0.20	40		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2231.18 / 2077.67	3.84 / 28.81	
▼ Hier.-4	2.00 / 0.04	50		8.00 / 0.20	40		4	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2397.97 / 2255.47	5.07 / 27.72	
▲ Hier.-5	0.40 / 0.04	10		8.00 / 0.20	40		4	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	22189.62 / 30800.65	4.99 / 10.97	
■ Hier.-6	1.20 / 0.04	30		8.00 / 0.20	40		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2072.90 / 1934.08	2.73 / 20.31	
◆ Hier.-7	1.20 / 0.04	30		8.00 / 0.20	40		4	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2178.72 / 2033.14	4.06 / 24.13	
▲ Hier.-8	2.00 / 0.04	50		4.00 / 0.20	20		10	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	2848.82 / 2695.63	3.18 / 26.60	
◆ Hier.-9	2.00 / 0.04	50		4.00 / 0.20	20		4	$\hat{\mathbb{S}}_f^{\text{lf}}, \hat{\mathbb{S}}_\omega^{\text{lf}}$	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							n.c. / 1242.62	5.23 / 152.10




Fig. 12. Visualization of different safe set approximations for body rates (three columns: **box constraints** $\hat{S}_\omega^{lf}$, **polyhedral** $\hat{S}_\omega^{lf}$, **nonlinear** S_ω^{lf}) for a step response to a horizontal set speed ($14 \frac{m}{s}$, 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, with the lower white region being 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.

than randomly initialized point-mass solvers. In the simpler, sparser environment, the parallel point-mass solvers can find an equally good solution. Figure 15 shows a quantitative comparison across 40 randomized mazes and different configurations. Remarkably, with 10 parallel randomly initialized solvers, a solution of similar quality to that of the MIQP solver was found.

How does UNIQUE compare against the hierarchical configuration with the same hyperparameters? In Fig. 15, we evaluate, for some configurations, the hierarchical setting, i.e., a planner with equal parameters and parallel evaluations together with a tracking MPC controller. The figure shows that, in terms of closed-loop cost, UNIQUE clearly outperforms the hierarchical baseline. Remarkably, if only progress is considered, the hierarchical configuration also achieves comparable performance, yet in a suboptimal way.

Can real-time control be achieved? In Fig. 16 the computation times of the parallel point-mass solvers, the main multi-phase MPC, and the MIQP are compared statistically. The time-critical solvers finish in time, while the MIQP has to run asynchronously because it exceeds the time budget.

Can the parallel evaluations finish on time? Since our algorithm assumes the parallel point-mass solver completes after

each iteration of the main multi-phase MPC, we evaluate the share of iterations in which the parallel solver finishes before the main multi-phase MPC in the maze environment over 200 iterations. We report a success rate of 99.5%. Moreover, we evaluate the same metric for a reinitialized full multi-phase MPC with randomized initial guesses and find that only 12% of the runs finish in time, highlighting the advantage of the considerably cheaper point-mass evaluations.

How does MIQP computation time scale with the number of obstacles? Most critically, the MIQP planner uses binary indicator constraints for obstacle-free convex cells, the number of which grows with the number of obstacles. In principle, this scales an NP-hard problem and could lead to a large increase in computation time. However, we show that, even with many obstacles, our pruning, preprocessing, and Gurobi’s internal optimization strategies keep computation times below one second, even for 50 obstacles, as detailed in the supplementary material.

IX. CONCLUSION AND DISCUSSION

The proposed UNIQUE framework introduces a temporal cascading paradigm that formulates the planning problem along the tail horizon of an MPC, replacing the conventional




Fig. 13. 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 horizontal axis reports the number of high-fidelity shooting nodes. 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. More accurate safe set formulations require longer computation times but improve performance, particularly for the shorter high-fidelity model phase.

hierarchical planner-controller decomposition with a sequential controller-planner structure within a single optimization. A long-horizon, low-fidelity point-mass model provides prospective planning, while a short-horizon, high-fidelity model ensures precise control, achieving real-time performance without hierarchical decomposition. This is enabled by the observation that the planning subproblem, primarily geometric obstacle avoidance, is a long-horizon concern that can be solved very efficiently with a point-mass model at each closed-loop iteration. 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 long-horizon planning 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 while preserving computational speed. Experiments show a reduction of up to 75% in closed-loop cost compared to hierarchical and standard MPC baselines. Both mechanisms still cannot guarantee finding the global optimal

solution, as is common in 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 [24].

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

Notably, motion planning is NP-hard in general [5]. However, motion planning with linear models, formulated as MIQPs or similar programs, performs sufficiently well in practice, as shown by various algorithms [8], [19], [22], [23], highlighting the focus on planning particularly along the second horizon of the proposed algorithm. For motion-planning problems of increasing complexity, such as larger mazes, the proposed approach cannot be applied without modifications. The main limiting factors are that the MPC horizon cannot be made arbitrarily long under a fixed computational budget and that numerical stability decreases as the horizon grows. Moreover, highly nonconvex planning costs such as time-optimal flight remain challenging to include in a numerically stable nonlinear program, although recent advances also aim to include more nonconvex costs [43]. A possible direction for future extensions would be to condense the long-horizon plan into a value function, potentially combined with the two stages.

The temporal cascading paradigm builds on two observations that hold well for quadrotors but may not transfer directly to all robotic systems: (i) that geometric planning around obstacles is predominantly a long-horizon concern with negligible influence on the short control horizon, and (ii) that meaningful state mappings from the high-fidelity to the low-fidelity model exist. For platforms where these conditions are less clear, e.g., legged robots with hybrid contact dynamics or manipulators with joint-space obstacles, the paradigm may still apply in principle, but the design of the transition function and the feasibility-preserving constraints require platform-specific analysis. Moreover, the effectiveness of temporal cascading relies on the low-fidelity planning model converging quickly once it reaches a good local minimum during closed-loop iteration. This property, which we empirically confirm for point-mass models solved via RTI, makes the second horizon computationally cheaper than the first. If the local planning model were too expensive to solve, as would be the case for a high-fidelity model on the tail horizon, the computational advantage over simply extending the MPC horizon would vanish. Importantly, when any parallel solver, including the MIQP, cannot finish within the main cycle, the algorithm degrades non-critically: the main multi-phase MPC retains its own feasible tail trajectory, and only completed, lower-cost solutions trigger reinitialization. This ensures that the real-time control loop is never blocked by the combinatorial planning search. A central design principle is to keep the




Fig. 14. Top view of simulation comparisons of UNIQUE with different hyperparameters. The left column shows a denser maze and the right column a sparser maze. The first row shows that UNIQUE without parallel planning gets easily stuck in local minima. The second row shows successful planning in the simpler right-hand scenario with parallel randomly initialized point-mass planners. The plots show how the free space is explored by alternating RTIs and reinitialization with initial guesses. The third row shows successful planning with the MIQP parallel planner, and the last row shows successful planning for all scenarios with additional point-mass planners.




Fig. 15. Closed-loop cost comparison on randomized instances of the maze environment. Including the mixed-integer planner yields the lowest closed-loop cost, provided sufficient computational resources. The hierarchical planner/controller architecture has a significantly larger closed-loop cost, despite also being able to navigate the mazes toward the goal.




Fig. 16. Timing comparison of the different solvers in the maze environment. Remarkably, the parallel, randomly initialized point-mass solvers can be solved much faster than the main multi-phase MPC and can therefore be synchronized after each iteration. However, the MIQP solver must be executed asynchronously due to its longer computation time.

high-fidelity controller fast and delegate the computationally harder planning search to parallel asynchronous tasks. By the principle of optimality, any improvement found on the tail trajectory reduces the overall cost without affecting the first-phase controller iteration time. The Pareto analyses confirm that this decomposition is beneficial not only for planning but also for control: under the same computational budget, the two-phase structure consistently outperforms a longer, high-fidelity MPC horizon.

UNIQUE is a real-time MPC that treats planning as a tail-horizon problem rather than a hierarchical stage, including long-horizon planning as an extension to short-horizon control in a single computationally efficient NLP through temporal cascading of planning and control.

REFERENCES

- [1] A. Romero, R. Penicka, and D. Scaramuzza, "Time-optimal online replanning for agile quadrotor flight," *IEEE Robot. Autom. Lett.*, 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 Trans. Robot.*, vol. 38, no. 6, pp. 3340–3356, Dec. 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 *Proc. Robot., Sci. Syst. (RSS)*, D. Kulis, G. Venture, K. Bekris, and E. Coronado, Eds. Robotics Science and Systems Foundation, 2024, Conference Paper, p. 109.
- [4] M. Hehn and R. D'Andrea, "Real-time trajectory generation for quadrotors," *IEEE Trans. Robot.*, vol. 31, no. 4, pp. 877–892, 2015.
- [5] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [6] R. Reiter, K. Baumgärtner, R. Quirynen, and M. Diehl, "Progressive smoothing for motion planning in real-time NMPC," in *Proc. Eur. Control Conf. (ECC)*, 2024, pp. 1816–1823.
- [7] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM J. Control Optim.*, vol. 43, no. 5, pp. 1714–1736, 2005.
- [8] T. Marcucci, P. Nobel, R. Tedrake, and S. Boyd, "Fast path planning through large collections of safe boxes," *IEEE Trans. Robot.*, vol. 40, pp. 3795–3811, 2024.
- [9] T. Cieslewski, E. Kaufmann, and D. Scaramuzza, "Rapid exploration with multi-rotors: A frontier selection method for high speed flight," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2017, pp. 2135–2142.
- [10] M. Naazare, D. Ramos, J. Wildt, and D. Schulz, "Application of graph-based path planning for UAVs to avoid restricted areas," in *Proc. IEEE Int. Symp. Saf. Secur. Rescue Robot. (SSRR)*, 2019, pp. 139–144.
- [11] W. Liu, Y. Ren, and F. Zhang, "Integrated planning and control for quadrotor navigation in presence of suddenly appearing objects and disturbances," *IEEE Robot. Autom. Lett.*, vol. 9, no. 1, pp. 899–906, 2024.
- [12] M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "Path planning for motion dependent state estimation on micro aerial vehicles," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2013, pp. 3926–3932.
- [13] F. Meyer, K. Glock, and D. Sayah, "TOP-UAV: Open-source time-optimal trajectory planner for point-masses under acceleration and velocity constraints," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2023, pp. 2838–2845.
- [14] K. Teissing, M. Novosad, R. Penicka, and M. Saska, "Real-time planning of minimum-time trajectories for agile UAV flight," *IEEE Robot. Autom. Lett.*, vol. 9, no. 11, pp. 10 351–10 358, 2024.
- [15] E. Tal and S. Karaman, "Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 3, pp. 1203–1218, 2021.
- [16] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE Robot. Autom. Lett.*, vol. 3, no. 2, pp. 620–626, 2018.
- [17] 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," *Sci. Robot.*, vol. 10, no. 101, p. eadp9905, 2025.
- [18] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. R. Kumar, "Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments," *IEEE Robot. Autom. Lett.*, vol. 2, no. 3, pp. 1688–1695, 2017.
- [19] J. Tordesillas, B. T. Lopez, M. Everett, and J. P. How, "Faster: Fast and safe trajectory planner for navigation in unknown environments," *IEEE Trans. Robot.*, vol. 38, no. 2, pp. 922–938, 2022.
- [20] Z. Wang, X. Zhou, C. Xu, and F. Gao, "Geometrically Constrained Trajectory Optimization for Multicopters," *IEEE Trans. Robot.*, vol. 38, no. 5, pp. 3259–3278, Oct. 2022.
- [21] J.-C. Latombe, "Exact cell decomposition," in *Robot Motion Planning*. Boston, MA: Springer US, 1991, pp. 200–247.
- [22] T. Marcucci, J. Umenberger, P. Parrilo, and R. Tedrake, "Shortest paths in graphs of convex sets," *SIAM J. Optim.*, vol. 34, no. 1, pp. 507–532, 2024.
- [23] T. Marcucci, M. Petersen, D. von Wrangel, and R. Tedrake, "Motion planning around obstacles with convex optimization," *Sci. Robot.*, vol. 8, no. 84, p. eadf7843, 2023.
- [24] 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 Trans. Control Syst. Technol.*, vol. 34, no. 1, pp. 395–410, 2026.

- [25] K. Seel, A. B. Kordabad, S. Gros, and J. T. Graydahl, "Convex neural network-based cost modifications for learning model predictive control," *IEEE Open J. Control Syst.*, vol. 1, pp. 366–379, 2022.
- [26] S. Abdufattokhov, M. Zanon, and A. Bemporad, "Learning Lyapunov terminal costs from data for complexity reduction in nonlinear model predictive control," *Int. J. Robust Nonlinear Control*, vol. 34, no. 13, pp. 8676–8691, 2024.
- [27] G. De Nicolao, L. Magni, and R. Scattolini, "Stabilizing receding-horizon control of nonlinear time varying systems," *IEEE Trans. Autom. Control*, vol. 43, no. 7, pp. 1030–1036, 1998.
- [28] M. Diehl, L. Magni, and G. D. Nicolao, "Efficient NMPC of unstable periodic systems using approximate infinite horizon closed loop costing," *Annu. Rev. Control*, vol. 28, no. 1, pp. 37–45, 2004.
- [29] 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, arXiv:2504.02710.
- [30] V. Behrunani, P. Heer, R. Smith, and J. Lygeros, "Recursive feasibility guarantees in multi-horizon mpc," in *Proc. IEEE Conf. Decis. Control (CDC)*. Piscataway, NJ: IEEE, 2024, Conference Paper, pp. 297 – 302.
- [31] J. Frey, K. Baumgaertner, G. Frison, and M. Diehl, "Multi-phase optimal control problems for efficient nonlinear model predictive control with acados," *Optim. Control Appl. Methods*, vol. 46, no. 2, pp. 827–845, 2025.
- [32] H. Li, R. J. Frei, and P. M. Wensing, "Model hierarchy predictive control of robotic systems," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 3373–3380, 2021.
- [33] H. Li and P. M. Wensing, "Cafe-MPC: A cascaded-fidelity model predictive control framework with tuning-free whole-body control," *IEEE Trans. Robot.*, vol. 41, pp. 837–856, 2025.
- [34] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 3, pp. 972–983, 2021.
- [35] J. Arrizabalaga, Z. Manchester, and M. Ryll, "Differentiable collision-free parametric corridors," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2024, pp. 1839–1846.
- [36] K. Tracy, T. A. Howell, and Z. Manchester, "Differentiable collision detection for a set of convex primitives," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 3663–3670.
- [37] 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, Aug. 2023.
- [38] 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," *Math. Program. Comput.*, vol. 14, no. 1, pp. 147–183, Mar. 2022.
- [39] 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.
- [40] J. Nocedal and S. Wright, *Numerical optimization*, 2nd ed., ser. Springer series in operations research and financial engineering. New York, NY: Springer, 2006.
- [41] 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," *Sci. Robot.*, vol. 7, no. 67, p. eabl6259, 2022.
- [42] Betaflight Development Team, "Betaflight: Open-Source Flight-Controller Firmware," <https://github.com/betaflight/betaflight>, 2026, Software repository, accessed Jul. 14, 2026.
- [43] K. Baumgärtner and M. Diehl, "The extended gauss-newton method for nonconvex loss functions and its application to time-optimal model predictive control," in *Proc. Amer. Control Conf. (ACC)*, 2022, pp. 4973–4978.



Rudolf Reiter is a postdoctoral researcher in the Robotics and Perception Group under Prof. Davide Scaramuzza at the University of Zurich, Zurich, Switzerland. He received his Ph.D. degree from the University of Freiburg, Freiburg, Germany, under Prof. Moritz Diehl, and specializes in control systems, optimization, machine learning, and robotics. His research focuses on learning- and optimization-based motion planning and control for autonomous systems.



Chao Qin received the B.Eng. degree in electrical engineering from Xidian University, Xi'an, China, in 2016, the M.S. degree in aerospace engineering from Shanghai Jiao Tong University, Shanghai, China, in 2019, and the Ph.D. degree in aerospace engineering from the University of Toronto, Toronto, ON, Canada, in 2025.



Leonard Bauersfeld received the M.Sc. degree in robotics, systems and control from ETH Zurich, Zurich, Switzerland, in 2020. He obtained his Ph.D. degree in the Robotics and Perception Group at the University of Zurich, Zurich, Switzerland, led by Prof. Davide Scaramuzza. His research interests are autonomous vision-based quadrotor flight and quadrotor simulations. He works on combining first-principles methods with modern data-driven models to advance agile quadrotor flight.



Davide Scaramuzza is a Professor of Robotics and Perception at the University of Zurich. He did his Ph.D. at ETH Zurich, a postdoc at the University of Pennsylvania, and was a visiting professor at Stanford University and NASA Jet Propulsion Laboratory. His research focuses on autonomous, agile navigation of mobile robots using standard and event-based cameras. He made fundamental contributions to visual-inertial state estimation, autonomous vision-based agile navigation of micro flying robots, and low-latency perception with event cameras, which were transferred to many products, from drones to automobiles, cameras, AR/VR headsets, and mobile devices. He pioneered autonomous, vision-based navigation of drones, which inspired the algorithm of the NASA Mars helicopter. In 2022, his team demonstrated that an AI-powered drone could outperform the world champions of drone racing. He received several awards, including an IEEE Technical Field Award, the IEEE Fellowship, the IEEE Robotics and Automation Society Early Career Award, a European Research Council Consolidator Grant, a Google Research Award, and many paper awards. In 2015, he co-founded Zurich-Eye, today Meta Zurich, which developed the head-tracking software of the Meta Quest. In 2020, he co-founded SUIND, which builds autonomous drones for precision agriculture. Many aspects of his research have been featured in the media, such as The New York Times, The Guardian, The Economist, and Forbes. He co-authored the book "Introduction to Autonomous Mobile Robots," published by MIT Press, which has sold over 10 thousand copies worldwide and is among the most used textbooks for teaching mobile robotics. He has been consulting the United Nations on disaster response, the Fukushima Action Plan, disarmament, and AI for good.