



Delft University of Technology

Embedded Hierarchical MPC for Autonomous Navigation

Benders, Dennis; Kohler, Johannes; Niesten, Thijs; Babuska, Robert; Alonso-Mora, Javier; Ferranti, Laura

DOI

[10.1109/TRO.2025.3567529](https://doi.org/10.1109/TRO.2025.3567529)

Publication date

2025

Document Version

Final published version

Published in

IEEE Transactions on Robotics

Citation (APA)

Benders, D., Kohler, J., Niesten, T., Babuska, R., Alonso-Mora, J., & Ferranti, L. (2025). Embedded Hierarchical MPC for Autonomous Navigation. *IEEE Transactions on Robotics*, 41, 3556-3574.
<https://doi.org/10.1109/TRO.2025.3567529>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Embedded Hierarchical MPC for Autonomous Navigation

Dennis Benders^{ID}, *Graduate Student Member, IEEE*, Johannes Köhler^{ID}, *Member, IEEE*, Thijs Niesten^{ID}, Robert Babuška^{ID}, Javier Alonso-Mora^{ID}, *Senior Member, IEEE*, and Laura Ferranti^{ID}, *Senior Member, IEEE*

Abstract—To efficiently deploy robotic systems in society, mobile robots must move autonomously and safely through complex environments. Nonlinear model predictive control (MPC) methods provide a natural way to find a dynamically feasible trajectory through the environment without colliding with nearby obstacles. However, the limited computation power available on typical embedded robotic systems, such as quadrotors, poses a challenge to running MPC in real time, including its most expensive tasks: constraints generation and optimization. To address this problem, we propose a novel hierarchical MPC scheme that consists of a planning and a tracking layer. The planner constructs a trajectory with a long prediction horizon at a slow rate, while the tracker ensures trajectory tracking at a relatively fast rate. We prove that the proposed framework avoids collisions and is recursively feasible. Furthermore, we demonstrate its effectiveness in simulations and lab experiments with a quadrotor that needs to reach a goal position in a complex static environment. The code is efficiently implemented on the quadrotor’s embedded computer to ensure real-time feasibility. Compared to a state-of-the-art single-layer MPC formulation, this allows us to increase the planning horizon by a factor of 5, which results in significantly better performance.

Index Terms—Embedded autonomous mobile robots, hierarchical model predictive control (MPC), obstacle avoidance, real-time motion planning and tracking.

I. INTRODUCTION

AUTONOMOUS mobile robots play an increasingly important role in our society [1]. The application domains are widespread, including self-driving cars [2] and environment

Received 29 November 2024; accepted 19 April 2025. Date of publication 6 May 2025; date of current version 9 June 2025. The work of Laura Ferranti was supported by the Dutch Science Foundation NWO-TTW Foundation within the Veni Project HARMONIA (nr. 18165). The work of Johannes Köhler was supported by the Swiss National Science Foundation under NCCR Automation under Grant 51NF40 180545. This work was supported by the National Police of The Netherlands. This article was recommended for publication by Associate Editor D. Panagou and Editor T. Bretl upon evaluation of the reviewers’ comments. (Corresponding author: Dennis Benders.)

Dennis Benders, Thijs Niesten, Javier Alonso-Mora, and Laura Ferranti are with the Department of Cognitive Robotics, Delft University of Technology, 2628, CD Delft, The Netherlands (e-mail: d.benders@tudelft.nl; t.e.j.niesten@tudelft.nl; j.alonsomora@tudelft.nl; l.ferranti@tudelft.nl).

Johannes Köhler is with the Institute for Dynamic Systems and Control, CH-8092 Zürich, Switzerland (e-mail: jkoehle@ethz.ch).

Robert Babuška is with the Department of Cognitive Robotics, Delft University of Technology, 2628, CD Delft, The Netherlands, and also with the Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, 16000 Praha, Czechia (e-mail: r.babuska@tudelft.nl).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TRO.2025.3567529>, provided by the authors.

Digital Object Identifier 10.1109/TRO.2025.3567529

exploration in search and rescue operations [3]. To successfully perform its task in such applications, a robot typically needs to navigate through a complex environment and reach a goal [4], [5]. This requires solving the motion planning and control problem: the robot needs to plan a smooth, collision-free, and dynamically feasible trajectory to avoid unnecessary braking and remain safe.

A natural way to solve this problem is to leverage nonlinear model predictive control (MPC) [6]. MPC is an effective method that, based on a nonlinear system model and an online observed environment, can optimize for a smooth, dynamically feasible, and collision-free trajectory. Using MPC, one can incorporate model, actuator, system, and obstacle avoidance constraints in a straightforward way, which, if constructed properly, provides safety guarantees of the closed-loop system [7], [8], [9].

A common approach is to use a single-layer MPC (SMPC) scheme that executes trajectory planning and tracking as a single optimization problem [7], [8], [9], [10], [11], [12], [13]. However, the main drawback of this approach is the time to solve the problem, which increases with increasing prediction horizon (for long-term predictions) and decreasing sampling times (for fast feedback), meaning that the user has to make a tradeoff between long-term planning and accurate tracking. This tradeoff leads to suboptimal solutions, especially on mobile robots that carry embedded hardware with limited computing power.

To overcome this problem, recent works consider decoupling the planning and tracking tasks in the form of a hierarchical MPC (HMPC) scheme, where a planning MPC (PMPC) computes a long-term trajectory, which is tracked by a fast tracking MPC (TMPC) [14], [15], [16], [17], [18], [19]. In HMPC, the PMPC must generate a dynamically feasible trajectory that the TMPC is able to track. Although necessary for reliable operation, some of the mentioned works do not consider the codesign of PMPC and TMPC, i.e., designing the PMPC based on the tracking capabilities of the TMPC. Furthermore, to ensure runtime feasibility, these works either consider linearizing the nonlinear model or optimizing a plan from a predefined set of motion primitives, thereby limiting the maneuverability of the nonlinear system.

To address both problems, we propose a novel HMPC framework. In this framework, PMPC and TMPC use the same nonlinear robot model but are sampled at different rates. This allows planning a dynamically feasible trajectory far ahead and tracking it accurately (see Fig. 1). The TMPC and PMPC are codesigned so that we can formally prove collision avoidance

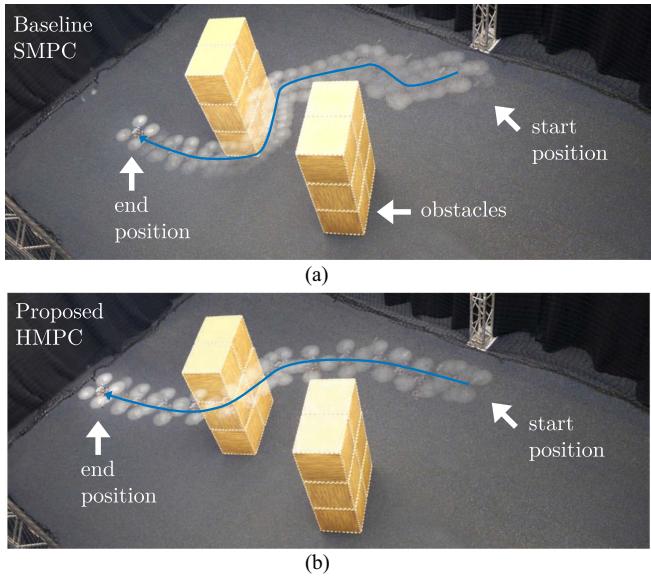


Fig. 1. Closed-loop trajectory of (a) baseline SMPC scheme and (b) proposed HMPC scheme. Whereas SMPC solves the planning and tracking tasks in a single-layer formulation, HMPC decouples the tasks using a PMPC and a TMPC that use the same nonlinear model but run at different frequencies. SMPC uses the same nonlinear model as PMPC and TMPC and is sampled at the same rate as TMPC to ensure reliable operation. Both schemes fly from start to goal position and avoid the obstacles. Compared to SMPC, HMPC has a significantly longer planning horizon, given the limited computational resources on the onboard computer, and does not have to balance planning and tracking tasks. Therefore, HMPC reaches the goal faster, is less sensitive to model mismatch, and maintains altitude better than SMPC.

and recursive feasibility for the overall HMPC framework in static environments. Furthermore, we show the runtime feasibility of the framework on embedded hardware.

A. Related Work

The problem of planning and tracking a trajectory for a mobile robot has been addressed in the literature using different methods. Popular planning methods include reactive [20], [21], sampling-based [22], optimization-based [23], or a combination of sampling and optimization-based [24] methods. The planner output can take various forms of plans, including paths [25], splines [26], and time-parameterized kinodynamic trajectories [23]. A general overview of motion planning methods can be found in [2], [27], and [28].

Depending on the type of plan, various tracking methods might be preferred. An overview of trajectory tracking methods on road and aerial vehicles is given in [29] and [30], respectively. Popular methods to design tracking controllers are proportional-integral-derivative [31], linear-quadratic regulator [32], sum-of-squares (SOS) [33], and MPC [16].

Most of the methods above focus on either the planning or the tracking task and do not consider the tracker capabilities in the planner. Consequently, they cannot guarantee the satisfaction of input and state constraints for general nonlinear robotic systems. Therefore, we consider works describing the planner and tracker codesign in the remainder of this section. We first discuss *general*

planner-tracker schemes that do not use MPC in both the planner and tracker. These schemes can be separated into approaches that perform planning *offline* and *online*. Afterward, we focus on existing *HMPC* approaches.

1) *Planning Offline*: Offline planning is useful in static scenarios such as car racing, [34] quadrotor racing [35], and industrial ground robot applications [36].

In the case of racing, there is usually no codesign of planner and tracker, and the assumption is that the planner and tracker are designed suitably and that the tracker is sampled at a sufficiently high frequency that it can accurately track the plan. Recent MPC schemes with tracking guarantees can be found in [37] and [38]. However, ensuring recursive feasibility requires restrictive terminal equality constraints.

Li et al. [36] showed the performance of a hierarchical scheme in which a trajectory is planned and the velocity is smoothed offline, and the plan is tracked with MPC online on an industrial autonomous ground vehicle.

It should be noted that a significant number of works in the offline planning field show embedded hardware results since it is easier to obtain runtime feasibility for these algorithms compared to online planning methods. However, they can only be applied to solve a predefined task in a known environment.

2) *Planning Online*: Online planning is relevant in scenarios where the perceived environment information of the robot changes during runtime. The main challenges of online planning methods are ensuring runtime feasibility and constructing a plan such that the closed-loop system avoids collisions despite deviating from the plan. Three relevant approaches to quantify this deviation are Lyapunov-based, hamilton-Jacobi (HJ) reachability-based, and contraction-based methods.

Common Lyapunov-based methods design funnels around offline-designed motion primitives [39], [40] from which the composable ones are selected during runtime [41]. The disadvantage is that the robot is limited to the set of motion primitives, which can lead to conservative plans. This problem is addressed in recent works using parameterized funnels [42] or a combination of sampling and invariant set analysis [43]. Other relevant Lyapunov-based methods include explicit reference governors [44] and control barrier functions [45], [46], [47].

One of the possible applications of HJ reachability-based methods is to derive a tracking error bound (TEB) between a low-fidelity planning and high-fidelity tracking model in order to reduce planning computation time [14], [48]. However, the TEB is computationally expensive and can be overly conservative. To address these problems, Althoff et al. [49] reduced computation time by linearizing around operating points and constructs corresponding zonotopic reachable sets, and Fridovich-Keil et al. [50] reduced conservatism by switching between a nonconservative “slow” plan and a conservative “fast” plan.

Contraction-based methods are based on contraction theory, which concludes about the convergence of two trajectories, the planned and closed-loop trajectories, based on the evolution of two infinitesimally close trajectories [51]. The approach in [52] leverages contraction theory and SOS programming to generate a control contraction metric controller with a corresponding

funnel around the trajectory but suffers from expensive computation time, both offline and online, and a conservative plan. The method proposed in [53] combines the approaches in [52] and [48] to reduce offline computations compared to [48] and online planning time compared to [52].

3) *Hierarchical MPC*: Similar to general planner-tracker methods, HMPC schemes also suffer from long planning computation times and propose similar solutions: using a simpler model or optimizing from a predefined set of motion primitives.

Following a similar idea to [48], Falcone et al. [14] used a low-fidelity nonlinear PMPC model and a high-fidelity linear time-varying TMPC from [54]. Other works leverage a linear PMPC model to reduce computation time. For example, Ibrahim et al. [18] showed how to efficiently use a linear mixed-integer PMPC and nonlinear TMPC to cover an area with a quadrotor in simulation.

Optimizing based on a set of motion primitives can be done via mixed-integer programming. This approach was taken in [15], demonstrating results in simulation and on a real passenger car driving on a slippery surface.

Other approaches combine a linear model and a set of motion primitives. One of these works is [16], in which the authors implement a runtime feasible HMPC scheme on small race cars by linearizing the nonlinear dynamics around the previously optimal trajectory and selecting trajectories from a predefined set of motion primitives in the PMPC. Another work is [17], in which simulation results are shown of a PMPC optimizing from a set of motion primitives and aggressive modes, and a TMPC being a tube-based formulation.

To avoid limited maneuverability, Kögel et al. [19] proposed a mixed-integer PMPC formulation that explicitly considers the distance to obstacles based on the aggressiveness mode and a TMPC being a cyclic horizon MPC. Quadrotor simulation results demonstrate that the method can safely fly from start to goal in a known static environment. However, the PMPC computation time is already at its limit for real-time feasibility while only using a linear model, and it was not implemented on embedded hardware.

In conclusion, designing an HMPC that computes a plan using the nonlinear robot dynamics online and accurately tracks it in real time on a real robot with limited computation power is challenging. The methods above solve this problem by reducing the model complexity or optimizing a plan using a set of predefined motion primitives. In contrast, our method uses the same nonlinear model in both MPC layers and guarantees tracking of the planned trajectory, similar to the idea of decoupling PMPC and TMPC while maintaining recursive feasibility presented in [55]. Leveraging recently developed MPC tracking [56] that guarantees computationally efficient tracking performance and recursive feasibility, we propose an HMPC framework for a mobile robot navigating in a static environment. The framework ensures collision avoidance and recursive feasibility, thereby taking an important step toward closing the gap between real-world mobile robot deployment and theoretical MPC feasibility and tracking guarantees.

B. Contributions

The main contribution of this article is a novel HMPC framework to enable embedded motion planning and tracking on a mobile robot. In particular, we propose the following:

- 1) a combined design methodology for PMPC and TMPC leveraging the same nonlinear model, in which the PMPC constructs a dynamically feasible trajectory that the TMPC is guaranteed to track by tightening the constraints in accordance with the offline terminal ingredients design of the TMPC. The proposed PMPC can solve the long-horizon planning problem completely independent of the TMPC over a longer, user-chosen time interval;
- 2) theoretical guarantees that both PMPC and TMPC are recursively feasible and the HMPC framework avoids static obstacles in complex environments at all times;
- 3) an efficient way to implement the convex free-space decomposition scheme presented in [57] for an MPC-based planner.

While the framework can be used on any mobile robotic platform, we implemented it on a quadrotor's NVIDIA Jetson Xavier NX embedded computer. Simulations and experiments validate the theoretical guarantees and practicality of the proposed HMPC framework.

The code base, written in C++ as a robot operating system package, implements different MPC formulations (i.e., SMPC, TMPC, and PMPC) and solves these problems leveraging the ForcesPro NLP solver [58], [59]. The only change in code for the different simulations and experiments is the interface, which allows for easier debugging and enhances sim-to-real transfer. The package, together with simulation and analysis code, is contained in a Docker environment, in line with the recommendation in [60], and is available at <https://github.com/dbenders1/hmpc>.

C. Outline

The rest of this article is organized as follows. Section II describes the considered robot and environment. Furthermore, it formulates the general trajectory planning and tracking problem and explains how the hierarchical framework follows from this scheme. Section III presents the TMPC design, including terminal ingredients (terminal cost and set) design with corresponding tracking properties. Section IV provides the PMPC design to construct a reference trajectory and the corresponding obstacle avoidance constraints for the TMPC to ensure recursive feasibility and obstacle avoidance of the closed-loop system. The overall HMPC framework is described in Section V, including concise pseudoalgorithms, the theoretical analysis of the framework and its components with proofs given in Appendices A and B, and a summary of the theoretical properties that the framework provides. Section VI provides the implementation details and experimental results. Finally, Section VII concludes this article.

D. Notation

$\mathbb{N}_{[a,b)}$ represents the natural numbers between a and b , including a and excluding b . Vectors are expressed in bold, matrices in capitals. $\|\cdot\|$ represents the 2-norm, and the quadratic

vector norm for a positive definite matrix Q is written as $\|\mathbf{x}\|_Q^2 = \mathbf{x}^\top Q \mathbf{x}$, $Q \succ 0$. All subscripts relate to time or indices of vectors/matrices, the rest of the information is given as superscript, i.e., \mathbf{x}_t^P represents the state at time t for PMPC P. The superscripts are omitted if the meaning is clear from the context. For the MPC problem at time t , the predicted state-input pair is denoted by $(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t})$, and its optimal solution by $(\mathbf{x}_{\tau|t}^*, \mathbf{u}_{\tau|t}^*)$ with predicted time $\tau \in [0, T]$ and prediction horizon T . Finally, $\mathbf{x}_{\cdot|t}$ represents all states over the prediction horizon.

II. PROBLEM FORMULATION

A. Robot Description

The mobile robot system is described by the following Lipschitz continuous nonlinear dynamics:

$$\dot{\mathbf{x}}_t = f(\mathbf{x}_t, \mathbf{u}_t) \quad (1)$$

with states $\mathbf{x}_t \in \mathbb{R}^{n^x}$, inputs $\mathbf{u}_t \in \mathbb{R}^{n^u}$, time $t \in \mathbb{R}$, and continuous-time dynamics $f(\mathbf{x}_t, \mathbf{u}_t)$. The system constraints are given by a compact polytopic set \mathcal{Z}

$$\mathcal{Z} := \mathcal{X} \times \mathcal{U}$$

$$= \{(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{n^x+n^u} \mid g_j^s(\mathbf{x}, \mathbf{u}) \leq 0, j \in \mathbb{N}_{[1, n^s]}\} \quad (2)$$

where $g_j^s(\mathbf{x}, \mathbf{u}) = L_j^s \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} - \mathbf{l}_j^s, L_j^s \in \mathbb{R}^{1 \times (n^x+n^u)}, \mathbf{l}_j^s \in \mathbb{R}, j \in \mathbb{N}_{[1, n^s]}$.

B. Obstacle Avoidance Constraints Formulation

Without loss of generality, the position states of a mobile robot are given by $\mathbf{p} = C\mathbf{x} \in \mathbb{R}^{n^p}$, with dimensions depending on the robot configuration space, where C is a matrix selecting the corresponding states from \mathbf{x} . Position \mathbf{p} is used to formulate obstacle avoidance constraints, ensuring that the robot region \mathcal{R} does not intersect with obstacle regions \mathcal{O} . The obstacle avoidance constraints at time t are given by the following polytopic free space set:

$$\mathcal{F}_{\tau|t} := \left\{ \mathbf{p}_{\tau|t} \mid g_{j,\tau|t}^o(\mathbf{p}_{\tau|t}) \leq 0, j \in \mathbb{N}_{[1, n^o]}\right\} \quad (3)$$

where $g_{j,\tau|t}^o(\mathbf{p}_{\tau|t}) = L_{j,\tau|t}^o \mathbf{p}_{\tau|t} - \mathbf{l}_{j,\tau|t}^o, L_{j,\tau|t}^o \in \mathbb{R}^{1 \times n^p}, \mathbf{l}_{j,\tau|t}^o \in \mathbb{R}, j \in \mathbb{N}_{[1, n^o]}$. Without loss of generality, we set $\|L_{j,\tau|t}^o\| = 1, j \in \mathbb{N}_{[1, n^o]}$ by scaling the constraints.

Section IV-B provides details on constructing the obstacle avoidance constraints.

C. Trajectory Planning and Tracking

The problem we aim to solve is to generate a dynamically feasible and collision-free trajectory for a mobile robot navigating in a static environment while considering the limited computation resources of the robot's embedded computer. To solve this problem, we rely on MPC. The SMPC scheme to solve this problem is given by

$$\min_{\mathbf{x}_{\cdot|t}, \mathbf{u}_{\cdot|t}} \mathcal{J}(\mathbf{x}_{\cdot|t}, \mathbf{u}_{\cdot|t}, \mathbf{p}^g) \quad (4a)$$

$$\text{s. t. } \mathbf{x}_{0|t} = \mathbf{x}_t \quad (4b)$$

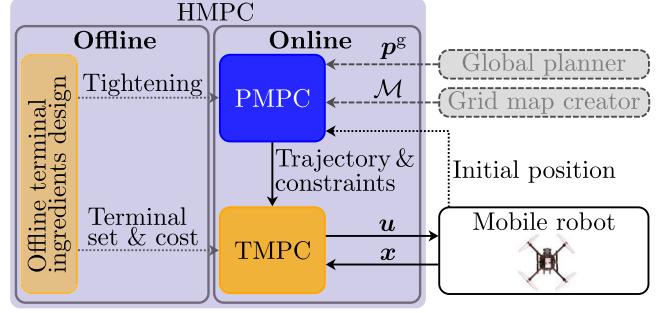


Fig. 2. HMPC framework overview. In the offline phase, the terminal ingredients of the TMPC, including the terminal set, cost, and tightening constants, are computed. During online operation, the PMPC plans a trajectory based on goal position \mathbf{p}^g and grid map \mathcal{M} , and the TMPC tracks this trajectory using state feedback \mathbf{x} .

$$\dot{\mathbf{x}}_{\tau|t} = f(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}), \quad \tau \in [0, T] \quad (4c)$$

$$(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}) \in \mathcal{Z}, \quad \tau \in [0, T] \quad (4d)$$

$$\mathbf{p}_{\tau|t} \in \mathcal{F}_{\tau|t}, \quad \tau \in [0, T] \quad (4e)$$

$$\mathbf{x}_{T|t} \in \mathcal{X}^f \quad (4f)$$

which optimizes a trajectory by minimizing cost function $\mathcal{J}(\mathbf{x}_{\cdot|t}, \mathbf{u}_{\cdot|t}, \mathbf{p}^g)$ (4a), while satisfying system constraints (4d) and obstacle avoidance constraints (4e), in order to steer the robot with dynamics (4c) from an initial state (4b) to a terminal set (4f) in the direction of goal position \mathbf{p}^g .

As discussed in the introduction, such an SMPC is challenging to implement on embedded hardware. Hence, our goal is to develop PMPC and TMPC formulations that, when combined, reach goal \mathbf{p}^g while satisfying constraints (4b)–(4f). PMPC and TMPC are combined so that (a) the TMPC is guaranteed to track the trajectory provided by the PMPC and (b) the overall scheme avoids collisions and is recursively feasible. In this setting, the PMPC minimizes a user-chosen planning cost, while the TMPC minimizes a tracking cost related to the reference generated by the PMPC. Both MPC formulations use the actual nonlinear dynamics (1) for the prediction.

To achieve this goal, the main TMPC design challenge is computing terminal ingredients, including terminal cost in (4a) and terminal set (4f), such that the closed-loop system tracks the reference computed by the PMPC.

The main PMPC design challenges are (a) constructing initial state constraint (4b), such that the communicated trajectory to the TMPC is continuous and dynamically feasible, (b) tightening the system and obstacle avoidance constraints, such that the robot does not collide with obstacles given the offline computed terminal set of the TMPC, and (c) choosing a proper terminal set ensuring recursive feasibility of the PMPC.

To summarize the codesign of PMPC and TMPC, Fig. 2 provides an overview of the proposed HMPC framework. The framework is detailed in Section V.

III. TMPC DESIGN

This section presents the TMPC formulation in Section III-A. Section III-B describes the offline design of suitable terminal

ingredients to prove recursive feasibility and convergence to the reference trajectory. Finally, Section III-C gives some properties that the reference trajectory and obstacle avoidance constraints generated by the PMPC should satisfy and formalizes the corresponding TMPC tracking properties in Proposition 1.

A. TMPC Formulation

The TMPC formulation is based on [56] with added obstacle avoidance capability

$$\min_{\mathbf{x}_{\cdot|t}, \mathbf{u}_{\cdot|t}} \mathcal{J}^{f,t}(\mathbf{x}_{T^t|t} - \mathbf{x}_{T^t|t}^r) + \int_{\tau=0}^{T^t} \mathcal{J}^{s,t}(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}, \mathbf{r}_{\tau|t}) d\tau \quad (5a)$$

$$\text{s. t. } \mathbf{x}_{0|t} = \mathbf{x}_t \quad (5b)$$

$$\dot{\mathbf{x}}_{\tau|t} = f(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}), \quad \tau \in [0, T^t] \quad (5c)$$

$$(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}) \in \mathcal{Z}, \quad \tau \in [0, T^t] \quad (5d)$$

$$\mathbf{p}_{\tau|t} \in \mathcal{F}_{\tau|t}, \quad \tau \in [0, T^t] \quad (5e)$$

$$(\mathbf{x}_{T^t|t} - \mathbf{x}_{T^t|t}^r) \in \mathcal{X}^{f,t} \quad (5f)$$

with TMPC sampling time $T^{s,t}$. This means that (5) is solved every $T^{s,t}$ seconds.

The cost penalizes the error between the prediction and the reference trajectory $\mathbf{r} = [\mathbf{x}^r \top \mathbf{u}^r \top]^\top$ using stage and terminal costs

$$\mathcal{J}^{s,t}(\mathbf{x}, \mathbf{u}, \mathbf{r}) := \|\mathbf{x} - \mathbf{x}^r\|_Q^2 + \|\mathbf{u} - \mathbf{u}^r\|_R^2 \quad (6)$$

$$\mathcal{J}^{f,t}(\mathbf{x} - \mathbf{x}^r) := \|\mathbf{x} - \mathbf{x}^r\|_P^2 \quad (7)$$

with $Q, R, P \succ 0$. The optimal input to Problem (5) is denoted by $\mathbf{u}_{\tau|t}^*$. As a result, the closed-loop system is given by

$$\mathbf{u}_{\tau+t} := \mathbf{u}_{\tau|t-T^{s,t}}^*, \quad \tau \in [0, T^{s,t}]. \quad (8)$$

B. Offline Terminal Ingredients Design

The terminal ingredients are computed offline using the semidefinite program (SDP) (9) shown at the bottom of this page, with linearized system dynamics

$$A(\mathbf{z}) = \frac{\partial f(\mathbf{x}, \mathbf{u})}{\partial \mathbf{x}} \Big|_{\mathbf{z}}, \quad B(\mathbf{z}) = \frac{\partial f(\mathbf{x}, \mathbf{u})}{\partial \mathbf{u}} \Big|_{\mathbf{z}}. \quad (10)$$

$$\min_{X, Y, c_j^{s2}} -\log \det X + \sum_{j=1}^{n^s} c_j^c c_j^{s2} \quad (9a)$$

$$\text{s. t. } \begin{bmatrix} A(\mathbf{z})X + B(\mathbf{z})Y + (A(\mathbf{z})X + B(\mathbf{z})Y)^\top & \begin{pmatrix} Q^{\frac{1}{2}}X \\ -I_{n^x} \end{pmatrix}^\top & \begin{pmatrix} R^{\frac{1}{2}}Y \\ 0 \end{pmatrix}^\top \\ * & 0 & \\ * & 0 & -I_{n^u} \end{bmatrix} \preceq 0 \quad \forall \mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} \in \mathcal{Z} \quad (9b)$$

$$\begin{bmatrix} c_j^{s2} & L_j^s \begin{bmatrix} Y \\ X \end{bmatrix} \\ \left(L_j^s \begin{bmatrix} Y \\ X \end{bmatrix} \right)^\top & X \end{bmatrix} \succeq 0, \quad j \in \mathbb{N}_{[1, n^s]}, \quad X \succeq 0. \quad (9c)$$

Based on the SDP solution, the terminal cost matrix P and feedback law gain K are computed using

$$P = X^{-1}, \quad K = YP \quad (11)$$

and the quadratic terminal cost is an incremental Lyapunov function [56] defined in (7).

The terminal set is defined as an ellipsoidal sublevel set of the incremental Lyapunov function given by

$$\mathcal{X}^{f,t} := \{\mathbf{x} \in \mathbb{R}^{n^x} \mid \|\mathbf{x}\|_P^2 \leq \alpha^2\} \quad (12)$$

with a user-chosen terminal set scaling $\alpha > 0$. $\mathcal{X}^{f,t}$ is positive invariant for control input $\mathbf{u} = \kappa^f(\mathbf{x}, \mathbf{r})$ with

$$\kappa^f(\mathbf{x}, \mathbf{r}) := \mathbf{u}^r + K(\mathbf{x} - \mathbf{x}^r) \quad (13)$$

which is crucial to ensure recursive feasibility and safety.

The objective of SDP (9) is to find a P that gives a suitable terminal cost and set for trajectory tracking. The weights c_j^c are tuning parameters that are normalized with respect to the system constraint intervals to ensure equal tightening of each constraint $j \in \mathbb{N}_{[1, n^s]}$ using tightening constant c_j^s , which will be defined later in (17a), and terminal set scaling α .

Linear matrix inequality (LMI) (9b) ensures that the incremental Lyapunov function (7) exponentially decreases when applying the feedback law (13) in the terminal set (12). LMI (9c) ensures that the system constraints are satisfied for all $\mathbf{z} \in \mathcal{Z}$, see [61] and [56] for more details.

Note that \mathcal{Z} represents the continuous state-input space, and we optimize for X , Y , and c_j^s that uniformly hold for all $\mathbf{z} \in \mathcal{Z}$. Therefore, SDP (9) is semi-infinite, and we need to grid or convexify \mathcal{Z} to solve it. If (9) is infeasible, one could leverage a more general formulation with state-dependent X and Y ; see [56] for details.

Note also that, since (9) ensures exponential stability with linear feedback K , a feasible solution for nonholonomic systems, such as cars, only exists when enforcing a minimum velocity.

C. Closed-Loop Properties

Properties 1 and 2 formalize the required properties for the reference trajectory and obstacle avoidance constraints.

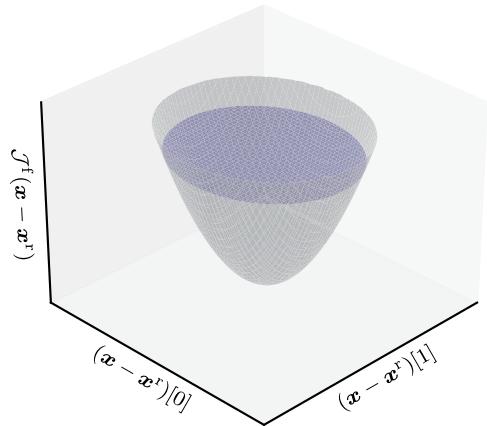


Fig. 3. Quadratic Lyapunov function $J^{f,t}(x - x^r)$ for two state tracking error dimensions. The blue ellipsoid is the terminal set $X^{f,t}$ intersecting $J^{f,t}(x - x^r)$ at value α^2 .

Property 1: The reference trajectory satisfies the following construction (a) and update (b) conditions.

1) For $\tau \in [0, T^t]$, it holds that

$$\dot{x}_{\tau|t}^r = f(x_{\tau|t}^r, u_{\tau|t}^r) \quad (14a)$$

$$r_{\tau|t} \in \bar{\mathcal{Z}} \quad (14b)$$

$$Cx_{\tau|t}^r \in \bar{\mathcal{F}}_{\tau|t} \quad (14c)$$

with

$$\begin{aligned} \bar{\mathcal{Z}} &:= \bar{\mathcal{X}} \times \bar{\mathcal{U}} \\ &= \left\{ (\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{n_x \times n_u} \mid \begin{array}{l} g_j^s(\mathbf{x}, \mathbf{u}) + c_j^s \alpha \leq 0, \\ j \in \mathbb{N}_{[1, n^s]} \end{array} \right\} \quad (15) \end{aligned}$$

and

$$\bar{\mathcal{F}}_{\tau|t} := \left\{ p_{\tau|t} \in \mathbb{R}^{n_p} \mid \begin{array}{l} g_{j,\tau|t}^o(p_{\tau|t}) + c^o \alpha \leq 0, \\ j \in \mathbb{N}_{[1, n^o]} \end{array} \right\} \quad (16)$$

with $\bar{\mathcal{Z}} \subseteq \mathcal{Z}$ and $\bar{\mathcal{F}}_{\tau|t} \subseteq \mathcal{F}_{\tau|t}$ following from tightening with $\alpha > 0$ from (12) and:

$$c_j^s := \|P^{-\frac{1}{2}}[IK^\top]L_j^s\|, \quad j \in \mathbb{N}_{[1, n^s]} \quad (17a)$$

$$c^o := \|P^{-\frac{1}{2}}C^\top\|. \quad (17b)$$

2) For $\tau \in [0, T^{s,t}]$, it holds that

$$x_{\tau+T^{s,t}}^r = x_{\tau|t+T^{s,t}}^r = x_{\tau+T^{s,t}|t}^r. \quad (18)$$

Intuitively, Property 1(a) states that the reference trajectory is dynamically feasible (14a), and it satisfies tightened system constraints (14b) and tightened obstacle avoidance constraints (14c). Property 1(b) states that the reference trajectory is continuous, i.e., the reference trajectory in the time interval $[0, T^{s,t}]$ of the current TMPC run is the same as the reference trajectory in the interval $[T^{s,t}, 2T^{s,t}]$ of the last TMPC run.

Fig. 3 illustrates an example of a 2-D Lyapunov function and corresponding ellipsoidal terminal set around the reference trajectory.

Remark 1: The terminal set scaling α can be seen as a tuning parameter to make a tradeoff between the performance of TMPC and PMPC. The smaller α , the closer PMPC can plan to the obstacle avoidance constraints, resulting in a less conservative trajectory. The bigger α , the bigger the terminal set for TMPC, making it easier for TMPC to find a solution. Although not formally proven in this article, the size of the terminal region alpha is related to the inherent robustness of the TMPC, cf. [62]. This property ensures that closed-loop properties remain valid under sufficiently small disturbances, which is crucial for application on the real robot.

Property 2: The obstacle avoidance constraints satisfy the following construction (a) and update (b) conditions.

- 1) They form a subset such that the robot does not collide with obstacles: $\mathcal{F}_{\tau|t} \cap \mathcal{O} \oplus \mathcal{R} = \emptyset$, $\tau \in [0, T^p]$.
- 2) They are updated in a consistent way, such that $p_{\tau|t}^r \in \bar{\mathcal{F}}_{\tau|t} \Rightarrow p_{\tau|t}^r \in \bar{\mathcal{F}}_{\tau-T^{s,p}|t+T^{s,p}}$, $\tau \in [T^{s,p}, T^p]$, with PMPC sampling time $T^{s,p}$ and prediction horizon T^p .

Intuitively, Property 2(b) states the reference trajectory should be contained in the updated tightened obstacle-free region, given that it is contained in the current tightened obstacle-free region at the same point in time.

Section IV shows that Properties 1 and 2 are ensured by the PMPC design.

Proposition 1 formalizes the properties of the TMPC terminal ingredients.

Proposition 1 (Terminal ingredients): Suppose the terminal ingredients are designed as above, r satisfies Property 1 and \mathcal{F} satisfies Property 2. Then, for any \mathbf{x} satisfying $(\mathbf{x} - \mathbf{x}^r) \in \mathcal{X}^{f,t}$

$$\frac{d}{dt} J^{f,t}(\mathbf{x} - \mathbf{x}^r) \leq -J^{s,t}(\mathbf{x}, \kappa^f(\mathbf{x}, \mathbf{r}), \mathbf{r}) \quad (19a)$$

$$(\mathbf{x}, \kappa^f(\mathbf{x}, \mathbf{r})) \in \mathcal{Z} \quad (19b)$$

$$C\mathbf{x} \in \mathcal{F}. \quad (19c)$$

Proof: (19a) holds given a feasible solution to SDP (9). The main steps of the proof, similar to [56], are

$$\begin{aligned} \frac{d}{dt} J^{f,t}(\delta) &= 2\delta^\top P(f(\mathbf{x}, \mathbf{u}) - f(\mathbf{x}^r, \mathbf{u}^r)) \\ &\stackrel{(a)}{=} 2\delta^\top P(A(z(s)) + B(z(s))K)\delta \\ &\stackrel{(b)}{\leq} -\delta^\top (Q + K^\top R K)\delta \\ &= -J^{s,t}(\mathbf{x}, \kappa^f(\mathbf{x}, \mathbf{r}), \mathbf{r}) \end{aligned}$$

where $\delta = \mathbf{x} - \mathbf{x}^r$. (a) uses the fact that $z(s) = \begin{bmatrix} \mathbf{x}^r \\ \mathbf{u}^r \end{bmatrix} + s \begin{bmatrix} I^{n_x} \\ K \end{bmatrix} \delta$ linearly interpolates between (\mathbf{x}, \mathbf{u}) and $(\mathbf{x}^r, \mathbf{u}^r)$ for $s \in [0, 1]$, and the mean value theorem ensuring that this equality holds for some value of s . (b) follows from LMI (9b), see [56, Lemma 3]. Furthermore, given $L \in \mathbb{R}^{1 \times n}$

satisfying $\|L\| = 1$ and δ such that $\|\delta\|_P^2 \leq \alpha^2$ it holds that

$$\begin{aligned} LC\delta &\leq \|LCP^{-\frac{1}{2}}P^{\frac{1}{2}}\delta\| \leq \|LCP^{-\frac{1}{2}}\|\|P^{\frac{1}{2}}\delta\| \\ &\leq \|CP^{-\frac{1}{2}}\|\|P^{\frac{1}{2}}\delta\| \leq \|CP^{-\frac{1}{2}}\|\alpha \end{aligned}$$

where the second and third inequalities use the triangular inequality. Applying this to L_j^o with $\|L_j^o\| = 1$, it holds that $L_j^o C \delta \leq c^o \alpha$, meaning that $L_j^o C x \leq L_j^o C x^r + c^o \alpha \leq l_j^o$, $j \in \mathbb{N}_{[1, n^o]}$ holds by (14c), given (16) and (17b), thereby proving (19c). A similar proof holds for (19b), see [56] for more details. \blacksquare

Intuitively, Proposition 1 states that, given Properties 1 and 2, the closed-loop trajectory converges to the reference trajectory and system and obstacle avoidance constraints are satisfied when applying terminal control law (13) in terminal set (12). Therefore, the terminal set is positively invariant under the terminal control law. This means that the shifted previously optimal solution appended with the terminal control law is a feasible, not necessarily optimal, candidate solution for the next TMPC run. Repeating this argument demonstrates that TMPC Problem (5) is recursively feasible. This is formally proven in Theorem 1 in Section V-B.

IV. PMPC DESIGN

The planner's goal is to optimize a dynamically feasible trajectory toward goal position p^g such that the TMPC is guaranteed to track it. Section IV-A presents the PMPC formulation to optimize the trajectory. The generation of the obstacle avoidance constraints is described in Section IV-B.

A. PMPC Formulation

The PMPC formulation is given by

$$\min_{\mathbf{x}_{\cdot|t}, \mathbf{u}_{\cdot|t}} \mathcal{J}^{f, p}(\mathbf{x}_{T^p|t}, \mathbf{u}_{T^p|t}, p^g) + \int_{\tau=0}^{T^p} \mathcal{J}^{s, p}(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}, p^g) d\tau \quad (20a)$$

$$\text{s. t. } \mathbf{x}_{\tau|t} = \mathbf{x}_{\tau+T^s, p|t-T^s, p}^*, \quad \tau \in [0, T^s, p] \quad (20b)$$

$$\dot{\mathbf{x}}_{\tau|t} = f(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}), \quad \tau \in [0, T^p] \quad (20c)$$

$$(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}) \in \bar{\mathcal{Z}}, \quad \tau \in [0, T^p] \quad (20d)$$

$$\mathbf{p}_{\tau|t} \in \bar{\mathcal{F}}_{\tau|t}, \quad \tau \in [0, T^p] \quad (20e)$$

$$f(\mathbf{x}_{T^p|t}, \mathbf{u}_{T^p|t}) = 0 \quad (20f)$$

with PMPC sampling $T^{s, p}$, and $\bar{\mathcal{Z}}$ and $\bar{\mathcal{F}}_{\tau|t}$, $\tau \in [0, T^p]$, given by (15) and (16), respectively. As mentioned, the cost is a user-chosen planning cost that can include any terms suitable for navigating in the environment (e.g., goal-oriented MPC (GO-MPC), see [4], or MPCC, see [8], [11]). Section VI-A elaborates on the specific function used to generate the experimental results. The combination of initial state constraint (20b) and system dynamics (20c), is sufficient to enforce both (14a) in Properties 1(a) and (b). Furthermore, (20d) and (20e) ensure (14b) and (14c) in Property 1(a), respectively. Finally, (20f) is a terminal set constraint to ensure the recursive feasibility of the PMPC. Note that (20b) can be efficiently implemented by

using a shorter horizon $T^p - T^{s, p}$ with initial state constraint $\mathbf{x}_{0|t} = \mathbf{x}_{T^s, p|t-T^s, p}^*$.

B. Obstacle Avoidance Constraints

This section describes the method to generate obstacle avoidance constraints of the form (3). To deal with an online observed environment with obstacles \mathcal{O} , we leverage grid map representation \mathcal{M} in which the occupied cells $\bar{\mathcal{O}}$ indicate the obstacle edges. Given these occupied cells, the method should generate constraints satisfying Property 2.

In particular, the method builds on *DecompUtil* as introduced in [57, Fig. 5–8]. Based on the grid map and a linear path segment, *DecompUtil* builds a convex obstacle-free region by appending the tangential line of a growing ellipsoid around the path segment at the first encountered occupied grid cell after removing all grid cells behind the previously generated tangential lines and clipping the result to bounding box \mathcal{B} .

In this work, the obstacle avoidance constraints are constructed following the same method, based on piecewise linear segments between optimally planned PMPC positions $\mathbf{p}_{iT^s, p|t-T^s, p}^*$ and $\mathbf{p}_{(i+1)T^s, p|t-T^s, p}^*$, $i \in \mathbb{N}_{[1, N-1]}$, $N = \frac{T^p}{T^{s, p}}$, starting with \mathcal{B} around the initial position. This approach is further referred to as *Iterative-DecompUtil* (*I-DecompUtil*) and visualized in Fig. 4.

There are two important things to note in the figure: the *selection of \mathcal{M}* and the usage of the robot radius to both tighten the obstacle-free convex regions and inflate the obstacles.

Selection of \mathcal{M} reduces the computation time by reducing the number of grid cells that need to be considered.

Since (5) and (20) optimize for the geometric robot center, the distance between obstacle avoidance constraints and occupied grid cells should be at least the robot radius. Tightening the constraints by the robot radius would result in sharp corners between subsequent obstacle-free regions around sharp obstacle corners, resulting in more conservative plans. On the other hand, inflating the obstacles by the robot radius would lead to the scenario depicted in Fig. 5, meaning that a crash may occur. To prevent both issues, half of the robot radius is used to tighten the constraints, and the other half is used to inflate the obstacles. This ensures collision avoidance, given that the robot radius is significantly larger than the grid map resolution.

The resulting obstacle-free convex regions from *I-DecompUtil* are piecewise-defined as

$$\mathcal{F}_{\tau+iT^s, p|t} = \mathcal{F}_{iT^s, p|t}, \quad \tau \in (0, T^{s, p}], \quad i \in \mathbb{N}_{[0, N-1]}. \quad (21)$$

This means that $\mathbf{p}_{(iT^s, p, (i+1)T^s, p)|t}^* \in \mathcal{F}_{iT^s, p|t}$, $i \in \mathbb{N}_{[0, N-1]}$. Note that, although the regions are constructed using piecewise linear segments, the nonlinear robot trajectory is contained in these convex regions since the constraints are imposed on the TMPC and PMPC trajectories through (5e) and (20e), respectively.

To prove recursive feasibility of PMPC formulation (20) later in Section V-B, we consider the following assumption regarding the grid map \mathcal{M} and the PMPC sampling time $T^{s, p}$.

Assumption 1: The grid map \mathcal{M} and PMPC sampling time $T^{s, p}$ are such that each trajectory segment $\mathbf{p}_{\tau+iT^s, p|t-T^s, p}^*$, $\tau \in$

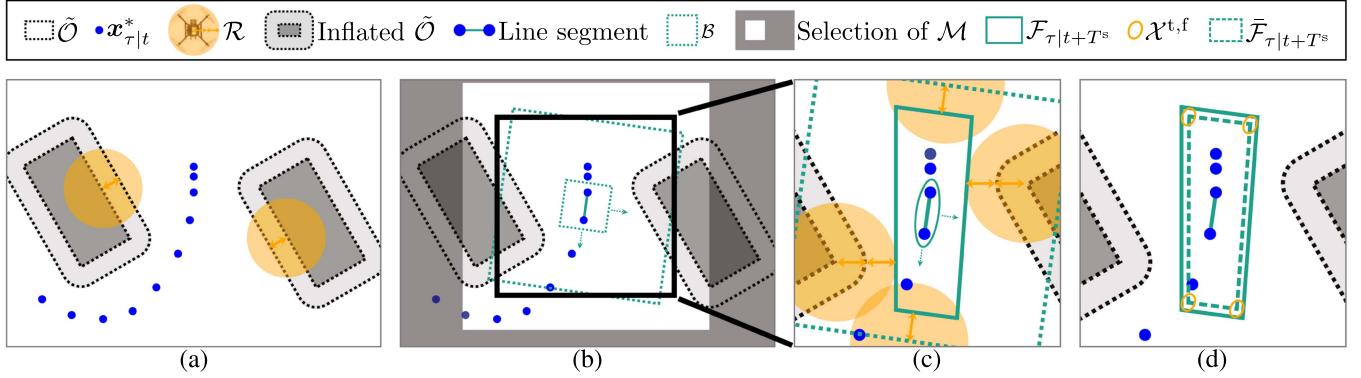


Fig. 4. 2-D visualization of map preprocessing and *I-DecompUtil*, given occupied grid cells $\tilde{\mathcal{O}}$ and the last optimized plan $\mathbf{x}_{\tau|t}^*$. (a) Obstacles are inflated by half of the robot radius (orange arrows). (b) To construct the obstacle avoidance constraints around a specific line segment, first, a subset of the grid map \mathcal{M} is selected such that the bounding box \mathcal{B} with any orientation fits in this subset. (c) Obstacle avoidance constraints $\mathcal{F}_{\tau|t+T^{s,p}}$ are constructed according to the SFC method in [57] by growing an ellipsoid around the line segment, creating the tangential lines and clipping them to \mathcal{B} . Furthermore, $\mathcal{F}_{\tau|t+T^{s,p}}$ are tightened by the other half of the robot radius, such that the robot does not collide with the obstacles if its center satisfies $\mathcal{F}_{\tau|t+T^{s,p}}$. (d) Tightened obstacle avoidance constraints $\bar{\mathcal{F}}_{\tau|t+T^{s,p}}$ are constructed by tightening $\mathcal{F}_{\tau|t+T^{s,p}}$ with $c^{\circ}\alpha$ according to (16), visualized using terminal set $\mathcal{X}^{t,f}$ in this figure. Note that (b)–(d) are repeated for all line segments.

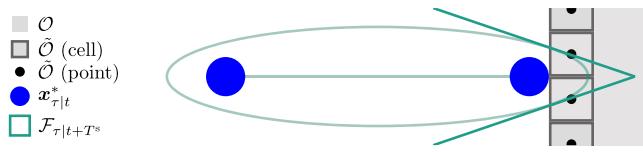


Fig. 5. Visualization of a scenario in which the robot would crash into the obstacle if the constraints are not tightened. The obstacle is inflated, here represented by \mathcal{O} , with corresponding occupied grid cells $\tilde{\mathcal{O}}$ and the point representation of $\tilde{\mathcal{O}}$ in the code. While the ellipsoid does not intersect the grid points, it overlaps with the obstacle region, i.e., $\mathcal{F} \cap \mathcal{O} \neq \emptyset$.

$[0, T^{s,p}], i \in \mathbb{N}_{[0, N-1]}$ is contained within the ellipsoid to be used to generate the first half-space constraint, $j = 1$ in (3), of the obstacle-free region $\mathcal{F}_{iT^{s,p}|t}, i \in \mathbb{N}_{[0, N-1]}$.

Remark 2: This assumption becomes a limiting factor in scenarios with a high obstacle density, a long PMPC sampling interval, and a high velocity. If this assumption is not satisfied, the PMPC sampling time or the maximum velocity constraint needs to be reduced. This was not a limiting factor in the simulations and experiments in Section VI.

Combining (21), Assumption 1, and the fact that the convex regions are generated based on the previously optimal solution $\mathbf{p}_{[0, T^p]|t-T^{s,p}}^*, \mathbf{p}_{[T^{s,p}, T^p]|t-T^{s,p}}^*$ also satisfies the new constraints

$$\mathbf{p}_{\tau|t-T^{s,p}}^* \in \mathcal{F}_{\tau-T^{s,p}|t}, \tau \in [T^{s,p}, T^p] \quad (22a)$$

$$\mathbf{p}_{T^p|t-T^{s,p}}^* \in \mathcal{F}_{T^p|t-T^{s,p}}. \quad (22b)$$

Thus, Property 2(a) is ensured by the combination of Assumption 1 and the fact that the obstacle avoidance constraints are generated by growing ellipsoids starting in free space until it touches $\tilde{\mathcal{O}}$. Moreover, Property 2(b) is satisfied by (22).

In conclusion, by formulation (20), Assumption 1, and the design of the obstacle avoidance constraints generation method, both Properties 1 and 2 are satisfied.

The following section provides an overview of how TMPC and PMPC are codesigned and details the theoretical properties of the overall HMPC framework.

V. HMPC FRAMEWORK

This section provides the design and implementation overview of the overall HMPC framework in Section V-A, followed by the theoretical analysis in Section V-B, and a summary of the framework design and properties in Section V-C.

A. Design and Implementation Overview

Fig. 2 gave an overview of the HMPC framework. Given \mathbf{p}^g from a global planner, and \mathcal{M} satisfying Assumption 1, the PMPC, formulated in (20) with sampling time $T^{s,p}$ satisfying Assumption 1 and horizon length T^p , generates obstacle avoidance constraints \mathcal{F} and optimizes a corresponding feasible reference trajectory $(\mathbf{x}^r, \mathbf{u}^r)$ toward \mathbf{p}^g . Based on the state feedback \mathbf{x} , the TMPC, formulated in (5) with sampling time $T^{s,t}$ and horizon length T^t , computes the control input \mathbf{u} required to track $(\mathbf{x}^r, \mathbf{u}^r)$.

The main motivation for this framework is its capability to solve an expensive optimization-based planning and tracking problem on an embedded computer with limited computation power in real time. Since the optimization takes a non-negligible amount of time (as shown later in Section VI), it introduces a delay in both the reference trajectory update and the control input update. To ensure the dynamic feasibility of the optimized solution, the framework is implemented with timing properties highlighted in Fig. 6.

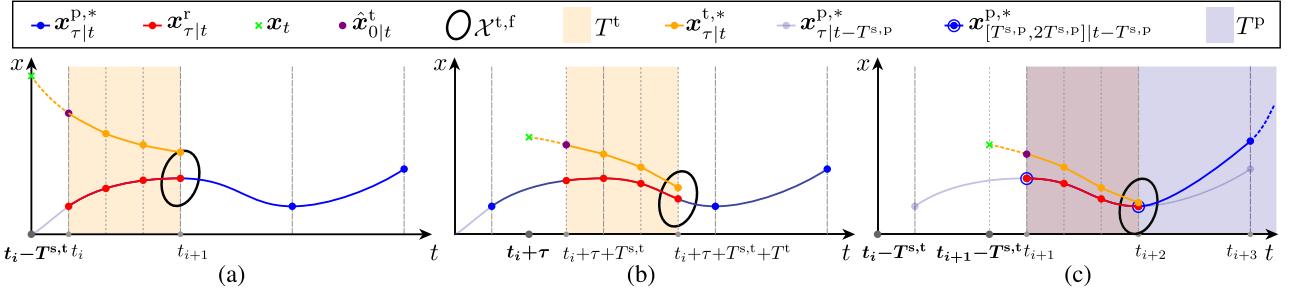


Fig. 6. Example timing diagram with 1-D state trajectories. (a) At time $t_i - T^{s,t}$, the TMPC optimizes a trajectory (orange) over horizon T^t starting from forward-simulated state $\hat{x}_{0|t_i}^t$ (purple), which is the model response when applying $u_{[0, T^{s,t}]}^{t,*}|_{t_i - T^{s,t}}$ starting from current state $x_{t_i - T^{s,t}}$ (green). The trajectory ends in terminal set $\mathcal{X}^{t,f}$ around the reference trajectory $x_{\tau|t_i}^r$ (red) that becomes valid at t_i . $x_{\tau|t_i}^r$ is the sub-sampled version of reference plan $x_{\tau|t_i}^{p,*}$ (blue). (b) In this example, the TMPC executes 2 more times ($\tau \in \{0, T^{s,t}\}$) until the next reference trajectory becomes valid, thereby getting closer to the reference. (c) At time $t_{i+1} - T^{s,t}$, the TMPC starts optimizing a trajectory based on a new reference plan $x_{\tau|t_{i+1}}^{p,*}$ that becomes valid at t_{i+1} . This reference plan is optimized by the PMPC starting at time $t_i - T^{s,t}$ and satisfies the initial state equality constraint (20b).

Algorithm 1: Offline Design.

Input: f, \mathcal{Z}

- 1: $X, Y \leftarrow$ Solve (9) using gridding or convexification
- 2: $P, K, c_j^s, c^o \leftarrow$ Compute using (11) and (17)
- 3: $\alpha \leftarrow$ Choose
- 4: $\mathcal{X}^{f,t}, \bar{\mathcal{Z}} \leftarrow$ Compute using (12) and (15)

Both PMPC and TMPC run at constant sampling times $T^{s,p}$ and $T^{s,t}$, respectively, with $T^{s,p} = \beta T^{s,t}$, $\beta \geq 2$. For convenience of visualization, the figure is constructed using $\beta = 3$. Note that the value of β used in the implementation is provided in Section VI. Furthermore, both PMPC and TMPC optimize a trajectory that becomes valid one sampling time in the future. This means that TMPC optimizes a trajectory $(u_{\tau|t_i}^{t,*}, x_{\tau|t_i}^{t,*})$ at $t_i - T^{s,t}$ before it becomes valid at t_i , based on a reference plan $(u_{\tau|t_i}^{p,*}, x_{\tau|t_i}^{p,*})$ that the PMPC started optimizing at $t_{i-1} - T^{s,t} = t_i - T^{s,p} - T^{s,t}$. At t_i , $u_{0|t_i}^{t,*}$ is sent to the robot before the TMPC starts optimizing the next trajectory $(u_{\tau|t_i+T^{s,t}}^{t,*}, x_{\tau|t_i+T^{s,t}}^{t,*})$. Thus, the TMPC optimization can take a maximum of $T^{s,t}$ and the PMPC a maximum of $T^{s,p}$ in execution time.

In conclusion, the HMPC framework gives valid reference trajectories and constraints and control inputs by the satisfaction of Properties 1 and 2 and allows sending control commands to the robot at a fixed frequency.

Remark 3: In theory, any $\beta \in \mathbb{I}_{\geq 1}$ works provided that $T^p \geq (1 + \lceil \frac{T^t}{T^{s,p}} \rceil + 2)T^{s,p}$ such that the PMPC has at least three stages free to optimize to move toward the goal and satisfy (20f). If $\lceil \frac{T^t}{T^{s,p}} \rceil > 1$ one has to adjust (20b) accordingly by constraining more PMPC stages. Given a TMPC design with fixed $T^{s,t}$ and T^t , increasing β means that $T^{s,p}$ increases, giving a coarser replanning time and, thus, reduced computational load for the PMPC, and vice versa. Hence, the user can trade off PMPC execution time and performance by choosing β .

Algorithms 1 and 2 give an overview of the relevant design choices of the HMPC framework.

Algorithm 2: Online Implementation.

Input: Results from Algorithm 1, $\mathcal{B}, \mathcal{R}, T^{s,t}, T^t$, $T^{s,p} (\beta), T^p, p^g$ and \mathcal{M}

Define: $t_i = iT^{s,p}, i \in \mathbb{N}_{\geq 0}$

Define: $t_{i,j} = t_i + jT^{s,t}, j \in \mathbb{N}_{[0, \beta-1]}$

- 1: $i = 0, j = 0$
- 2: **for** $t = t_{i,j}$ **run TMPC**
- 3: Apply $u_{\tau|t}^*, \tau \in [0, T^{s,t}]$ (8)
- 4: $x_t \leftarrow$ Measure
- 5: $x_{0|t_{i,j+1}} \leftarrow$ Forward-predict using x_t
- 6: $\mathcal{F}_{\cdot|t} \leftarrow$ Select based on $\mathcal{F}_{\cdot|t_i}$
- 7: $u_{\cdot|t}^* \leftarrow$ Solve (5) with $x_{0|t} = x_{0|t_{i,j+1}}$
- 8: $j \leftarrow j + 1$
- 9: **if** $j = \beta - 1$ **and** p^g **not reached run PMPC**
- 10: $\mathcal{F}_{\cdot|t_i}, \bar{\mathcal{F}}_{\cdot|t_i} \leftarrow$ Construct as in Section IV-B
- 11: $(x_{\cdot|t_i}^{p,*}, u_{\cdot|t_i}^{p,*}) \leftarrow$ Solve (20)
- 12: Send $(x_{\cdot|t_i}^r, x_{\cdot|t_i}^{p,*}, \mathcal{F}_{\cdot|t_i})$ to TMPC
- 13: $i \leftarrow i + 1$
- 14: $j \leftarrow 0$

B. Theoretical Analysis

First, Theorem 1 formalizes the TMPC obstacle avoidance and recursive feasibility guarantees. Second, Theorem 2 formalizes the recursive feasibility of the PMPC. By combining these results, Corollary 1 concludes the recursive feasibility and obstacle avoidance of the HMPC framework.

Theorem 1: Suppose the terminal ingredients are designed according to Section III-B, r satisfies Property 1, \mathcal{F} satisfies Property 2, and (5) is feasible at $t = 0$. Then, the resulting closed-loop system ensures that (5) is recursively feasible, satisfies system constraints $(x_t, u_t) \in \mathcal{Z}$ and avoids obstacle collisions $Cx_t \notin \mathcal{O}$ for all $t \geq 0$. Moreover, the tracking error $\|x_t - x_t^r\|$ asymptotically converges to zero.

The proof is detailed in Appendix A. It uses a candidate solution that shifts the previously optimal solution by $T^{s,t}$ and appends the terminal control law (13). Feasibility is ensured by the invariant terminal set, cf. Proposition 1,

and convergence follows by showing that the optimal cost decreases.

Theorem 2: Suppose that Assumption 1 holds and (20) is initialized with a feasible steady-state plan, i.e., $\dot{\mathbf{x}}^P = 0$, $(\mathbf{x}, \mathbf{u}) \in \bar{\mathcal{Z}}$, $\mathbf{p} \in \bar{\mathcal{F}}_{|0}$. Then, (20) is recursively feasible for $t \geq 0$, the optimal solution satisfies Property 1, and the constructed obstacle avoidance constraints satisfy Property 2.

The recursive feasibility proof is detailed in Appendix B. Similar to the proof for Theorem 1, it is based on a candidate solution equal to the previously optimal solution, appended with an input ensuring steady state. Given that the previously optimal solution ends in steady state, the candidate also ends in steady state, meaning that system and obstacle avoidance constraints are always satisfied. Satisfaction of Properties 1 and 2 follow from the reasoning explained in Section IV-B.

Corollary 1: Suppose the terminal ingredients are designed according to Section III-B, Algorithm 2 is initialized with a steady-state reference plan satisfying Property 1, and (5) and (20) are both feasible at $t = 0$. Then, (5) and (20) are recursively feasible for the closed-loop system in Algorithm 2 and the closed-loop system satisfies system constraints $(\mathbf{x}_t, \mathbf{u}_t) \in \mathcal{Z}$ and avoids obstacle collisions $\mathcal{R} \cap \mathcal{O} = \emptyset$ for $t \geq 0$. Moreover, the tracking error $\|\mathbf{x}_t - \mathbf{x}_t^r\|$ asymptotically converges to zero.

By Theorem 2, (20) is recursively feasible and ensures that the reference trajectory satisfies Property 1 and obstacle avoidance constraints satisfy Property 2. Then, given the initial feasibility of (5), applying Theorem 1 gives system and obstacle avoidance constraints satisfaction and asymptotic convergence to the reference trajectory of the closed-loop system.

C. Summary

We solve SDP (9) offline for known nonlinear system dynamics (1). A feasible solution to (9) gives matrices X and Y , and tightening constants c_j^s . X and Y are used to compute matrices P and K , which are subsequently used to calculate the tightening constants c^o and to construct the TMPC terminal cost, set, and control law.

Note that LMI (9c) ensures that c_j^s are Lipschitz constants, so the terminal set scaling α does not have to be known before solving (9) and can instead be computed as $\alpha = \frac{d}{c^o}$ with d the minimum distance between obstacles and reference trajectory.

Online, we compute the reference trajectory by solving PMPC Problem (20) using tightening constants c_j^s and c^o and obstacle avoidance constraints generated as described in Section IV-B. Theorem 2 states that the reference trajectory and corresponding constraints satisfy Properties 1 and 2, respectively, and that the PMPC is recursively feasible by including PMPC terminal constraint (20f).

Given the offline-computed terminal ingredients as proposed in Proposition 1 and online-computed reference trajectory satisfying Property 1 with corresponding obstacle avoidance constraints satisfying Property 2, Theorem 1 states that the online execution of TMPC Problem 5 guarantees asymptotic convergence to the reference trajectory, obstacle avoidance, and recursive feasibility.

Consequently, by the proposed codesign of TMPC and PMPC, Corollary 1 concludes that the HMPC framework is recursively feasible, satisfies system constraints, and avoids obstacle collisions at all times. In addition, the timing of PMPC and TMPC as described in Section V-A allows sending control commands to the robot at a fixed frequency while maintaining the properties described above.

Important to note is that the optimized trajectory of the PMPC evolves based on the previously optimal PMPC solution, so there is no state feedback from the real system. In contrast, the state feedback is incorporated in the TMPC formulation as an initial state constraint (5b).

VI. RESULTS

This section demonstrates an efficient way to generate obstacle avoidance constraints. Furthermore, it shows the performance and advantages of the proposed HMPC framework compared to SMPC, which is commonly used in local motion planning schemes [10], [12], [13]. To make a fair comparison, SMPC inherits the cost function and system dynamics from PMPC, and initial state, system, and obstacle avoidance constraints from TMPC. This means, it solves (20), where (20b), (20d), and (20e) are replaced by (5b), (5d), and (5e), respectively.

The results include simple simulations without model mismatch, Gazebo simulations with model mismatch, and lab experiments. All results will focus on the lab experiments and highlight the differences with simulations when relevant. The simulations and lab experiments use the same PMPC, TMPC, and SMPC settings, which are chosen to ensure real-time performance on the hardware of the real quadrotor. These settings include sampling times, horizon lengths, and weights.

We first detail the different MPC implementations and test setup in Section VI-A. After that, we provide the results on both obstacle avoidance constraints generation and comparison between HMPC and SMPC, in Section VI-B.

The open-source implementation to reproduce these results is available at <https://github.com/dbenders1/hmmpc>.

A. Implementation and Setup Details

The implementation and setup consider several aspects, including the considered grid map to the quadrotor model, the MPC details, the software setup, and the hardware setup, as described next.

1) *Grid Map:* The considered grid map \mathcal{M} is 2-D and generated offline with size 12×12 m and resolution 0.01 m. The rectangular obstacles are described by their center, width, and length. The obstacle contours, formed by thin lines of occupied grid cells, are efficiently created using OpenCV.¹ As described in Section IV-B, an extra inflation layer is designed around the obstacles to allow smoother planning of the PMPC around the obstacles. The same holds for the lines representing the environmental boundaries. \mathcal{M} is the same for the simulation and lab results below for consistency.

¹ [Online]. Available: <https://github.com/opencv/opencv>

2) *Quadrotor Model*: The following quadrotor model is used in the MPC schemes:

$$\begin{aligned} \begin{bmatrix} \dot{p}^x \\ \dot{p}^y \\ \dot{p}^z \end{bmatrix} &= \begin{bmatrix} v^x \\ v^y \\ v^z \end{bmatrix} \\ \begin{bmatrix} \dot{v}^x \\ \dot{v}^y \\ \dot{v}^z \end{bmatrix} &= \begin{bmatrix} s\phi s\psi + c\phi s\theta c\psi \\ -s\phi c\psi + c\phi s\theta s\psi \\ c\phi c\theta \end{bmatrix} a - \begin{bmatrix} 0 \\ 0 \\ g \end{bmatrix} \\ \begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \\ \dot{a} \end{bmatrix} &= \begin{bmatrix} -\frac{1}{\tau\phi} & 0 & 0 & 0 \\ 0 & -\frac{1}{\tau\theta} & 0 & 0 \\ 0 & 0 & -\frac{1}{\tau\psi} & 0 \\ 0 & 0 & 0 & -\frac{1}{\tau^a} \end{bmatrix} \begin{bmatrix} \phi \\ \theta \\ \psi \\ a \end{bmatrix} + \begin{bmatrix} \frac{k\phi}{\tau\phi} \\ \frac{k\theta}{\tau\theta} \\ \frac{k\psi}{\tau\psi} \\ \frac{k^a}{\tau^a} \end{bmatrix} \begin{bmatrix} \dot{\phi}^c \\ \dot{\theta}^c \\ \dot{\psi}^c \\ \dot{a}^c \end{bmatrix} \end{aligned} \quad (23)$$

with states including 3-D positions, velocities and attitudes, and acceleration given by $\mathbf{x} = [p^x \ p^y \ p^z \ v^x \ v^y \ v^z \ \phi \ \theta \ \psi \ a]^\top$ and inputs including 3-D attitude commands and collective mass-normalized thrust, or acceleration command, given by $\mathbf{u} = [\phi^c \ \theta^c \ \psi^c \ a^c]^\top$, sine (s) and cosine (c) expressions, gravitational constant g , and time and gain constants $\tau^\phi = \tau^\theta = 0.18$, $\tau^\psi = 0.56$, $\tau^a = 0.050$, $k^\phi = k^\theta = k^\psi = k^a = 1$ for roll, pitch, yaw and acceleration, respectively.

The system constraints are given by

$$\begin{aligned} -15 \text{ m} &\leq p^x, \ p^y & \leq 15 \text{ m} \\ 0 \text{ m} &\leq p^z & \leq 4 \text{ m} \\ -2 \text{ m/s} &\leq v^x, \ v^y, \ v^z & \leq 2 \text{ m/s} \\ -30^\circ &\leq \phi, \ \theta, \ \psi, \ \phi^c, \ \theta^c, \ \psi^c & \leq 30^\circ \\ 5 \text{ m/s}^2 &\leq a, \ a^c & \leq 15 \text{ m/s}^2. \end{aligned} \quad (24)$$

The PMPC and SMPC leverage an extended model in which the inputs are given by $[\dot{\phi}^c \ \dot{\theta}^c \ \dot{\psi}^c \ \dot{a}^c]$, with $-60^\circ/\text{s} \leq \dot{\phi}^c, \ \dot{\theta}^c, \ \dot{\psi}^c \leq 60^\circ/\text{s}$, to allow for penalizing angle rates in the cost function, and thus creating a smoother trajectory. The absolute angle commands $[\phi^c, \ \theta^c, \ \psi^c]$ are included in the states in that case.

Both PMPC and TMPC implement the discretization of the quadrotor model (23) using the fourth-order Runge–Kutta (RK4) integration method with a 50 ms time interval. The TMPC has a sampling time of 50 ms, requiring 1 RK4 step, whereas the PMPC has a sampling time of 500 ms, requiring $\beta = 10$ RK4 steps. Thus, although the sampling times differ between PMPC and TMPC, the model discretization is the same.

3) *TMPC Offline Computations*: The offline computations described in Section III are implemented using Yalmip. The attitude states occur nonlinearly, and the acceleration state occurs linearly in the linearized system dynamics (10). Therefore, the LMIs are generated using a grid of five points equally divided over the attitude state constraint intervals and the two endpoints of the acceleration constraint interval. The solution to SDP (9) is checked for 21 points per attitude state and two points for the acceleration state constraint interval. All LMIs are satisfied at the checked points in the state space, meaning that the theoretical

guarantees hold for those points. The total time to generate the LMIs, solve the SDP, and check the results is 38 s.

4) *PMPC and SMPC Cost Function*: The PMPC and SMPC aim to reach the goal with a cost similar to GO-MPC [4]. On a quadrotor, one can control the yaw angle independently from the position, meaning that the goal is described by a 3-D position with yaw angle $(\mathbf{p}^g, \ \psi^g) \in \mathbb{R}^4$. Consequently, the cost function is given by

$$\begin{aligned} \mathcal{J}^{s,p}(\mathbf{x}_{\tau|t}, \mathbf{u}_{\tau|t}, \mathbf{p}^g, \psi^g) &:= w^{xy,g} \mathcal{H}(\mathbf{p}_{\tau|t}^{xy}, \mathbf{p}^{xy,g}) + w^{z,g} (p_{\tau|t}^z - p^{z,g})^2 \\ &+ w^{\psi,c} (\psi_{\tau|t} - \psi^g)^2 + w^a (a_{\tau|t}^2 - g)^2 \\ &+ \|\mathbf{u}_{\tau|t}\|_U^2 + w^{\phi\theta,c} (\phi_{\tau|t}^c)^2 + w^{\theta\tau,c} (\theta_{\tau|t}^c)^2 + w^{\psi,c} (\psi_{\tau|t}^c)^2 \end{aligned} \quad (25)$$

with goal position weights $w^{xy,g}$ and $w^{z,g}$, goal yaw weight $w^{\psi,g}$, acceleration weight w^a , input weights matrix $U = \text{diag}(u_1, \dots, u_{n^u})$ and, and input attitude memory state weights $w^{\phi\theta,c}$ and $w^{\psi,c}$. Terminal cost $\mathcal{J}^{f,p}$ has the same expression as stage cost $\mathcal{J}^{s,p}$, but with different weight values. $\mathcal{H}(\mathbf{p}_{\tau|t}^{xy}, \mathbf{p}^{xy,g})$ denotes the Huber loss [63] of the 2-D Euclidean distance from position $\mathbf{p}_{\tau|t}^{xy}$ toward goal position $\mathbf{p}^{xy,g}$. The linear relation further away from the origin makes the Huber loss useful to ensure moving toward the goal with approximately constant velocity without accelerating quickly and reaching the goal slowly.

5) *PMPC and SMPC Terminal Steady State*: The terminal state used to implement (20f) in both PMPC and SMPC is given by $\{v^x = 0, v^y = 0, v^z = 0, \phi = 0, \theta = 0, \psi = 0, a = g\}$.

6) *TMPC Terminal Set*: As Section III shows, the terminal set scaling α is a tuning parameter to tradeoff PMPC and TMPC performance. In this case, α is computed using $\alpha = \frac{d}{c^o}$, in which $d = 0.1$ m is the minimum allowed distance from obstacles to reference trajectory and c^o is given by (17b).

7) *SMPC Slack*: Both the Gazebo simulations and lab experiments in Section VI-B3 include model mismatch. HMPC indirectly accounts for this by tightening the obstacle avoidance constraints in the PMPC. However, SMPC would become infeasible because the system would exceed the obstacle avoidance constraints. Therefore, the obstacle avoidance constraints in SMPC are implemented using a slack variable, e.g., see [64], allowing the system to exceed the constraints but with high penalization to steer the system back into the constraints set. Similar to the tightening of constraints in HMPC, SMPC uses a safety distance of 0.1 m to avoid collisions.

8) *PMPC, TMPC, and SMPC Settings*: Table I summarizes the PMPC, TMPC, and SMPC settings, including the sampling times, horizon lengths, and weights. The TMPC and SMPC sampling times are chosen sufficiently low to ensure accurate control but long enough to solve the MPC optimization problems in real time. Based on the TMPC sampling time $T^{s,t}$, the PMPC sampling time is chosen using $\beta = 10$. Given $T^{s,t}$, this factor is selected sufficiently low for the PMPC to optimize a reasonable plan and sufficiently high for the TMPC to satisfy the terminal set constraint.

The PMPC horizon is chosen long enough to be able to find a reasonable plan through the environments described below.

TABLE I
PMPC, TMPC, AND SMPC SETTINGS IN SIMULATIONS AND LAB EXPERIMENTS

Method	T^s (s)	T (s)	Weighting matrices
PMPC	0.5	2.5	$w^{xy,g^a} = 40, w^{z,g^a} = 40, w^{\psi,g^a} = 40, w^{xy,g^t} = 200, w^{z,g^t} = 200, w^{\psi,g^t} = 200, w^a = 40, w^{\phi\theta,c} = 16, w^{\psi,c} = 16, U = \text{diag}(16, 16, 16, 16)$
TMPC	0.05	0.5	$Q = \text{diag}(2e3, 2e3, 2e3, 20, 20, 20, 100, 100, 100, 100), R = \text{diag}(2e3, 2e3, 2e3, 100)$
SMPC	0.05	0.4	$w^{xy,g^a} = 200, w^{z,g^a} = 200, w^{\psi,g^a} = 200, w^{xy,g^t} = 2e3, w^{z,g^t} = 2e3, w^{\psi,g^t} = 2e3, w^a = 200, w^{\phi\theta,c} = 160, w^{\psi,c} = 160, U = \text{diag}(160, 160, 160, 160)$

The TMPC horizon is chosen to equal the PMPC sampling time. Furthermore, the SMPC horizon is reduced until the MPC scheme is runtime feasible on the hardware described as follows.

The GO objective weights used in PMPC and SMPC are chosen such that the terminal stage and other stages contribute equally to the cost function. Note that PMPC and SMPC have a similar weight ratio between the terminal stage and the rest of the horizon for a fair comparison. One of the differences is that the PMPC GO objective stage weights are increased to gain aggressiveness and improve goal-reaching performance, which is not possible for SMPC for stability reasons. Stability is ensured in both PMPC and SMPC by increasing the input weights. Note that this increase is more significant for SMPC than for PMPC.

The TMPC position state and attitude input weights are chosen large enough to ensure that the reference positions are accurately tracked to avoid collisions, and the system will not aggressively compensate for any reference tracking error, thereby increasing model mismatch. To further enhance stable flight, the acceleration input and state weights are increased. Since the reference acceleration is given in the body frame, its value is only valid for attitude values resulting in the corresponding body frame. Therefore, the attitude weights are increased as well.

Since we compute the terminal cost offline, we hard-code the TMPC weighting matrices to save computation time during runtime. In contrast, the PMPC and SMPC weights are not hard-coded to allow the user to change them without generating a new solver.

Note that the costs are given in continuous time and discretized using Euler.

9) *Software Setup*: In all simulations and lab experiments, the ForcesPro NLP solver [58], [59] is leveraged, using x86 compilation for simulations and ARM compilation for lab experiments. Furthermore, in both Gazebo simulations and lab experiments, PX4 stable release v1.12.3 is leveraged, using Posix compilation for simulations and NuttX compilation for lab experiments. The advantage of this setup is that both the solver and low-level controllers are the same in simulations and lab experiments to reduce sim-to-real transfer.

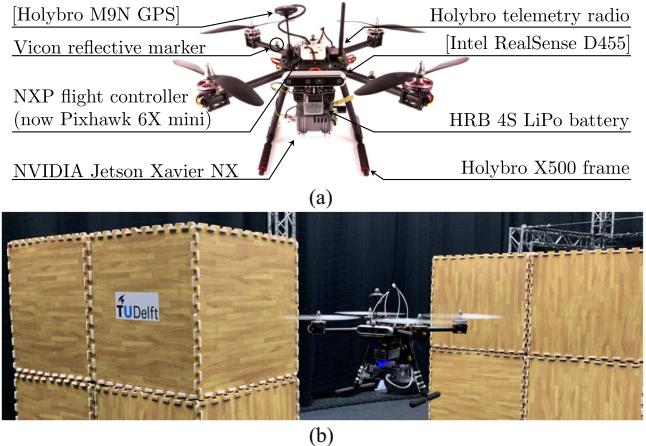


Fig. 7. Lab experiment setup. (a) Quadrotor with important components is highlighted. Components between square brackets are not used in this work's experiments. (b) Quadrotor is flying near the obstacles in the lab.

10) *Hardware Setup*: The quadrotor used to generate the lab experiment results is the HoverGames drone², with the flight controller being replaced by a Pixhawk 6X mini and an added NVIDIA Jetson Xavier NX embedded computer, which has a 6-core 1400 MHz NVIDIA Carmel ARMv8.2 processor. The position and orientation of the quadrotor are tracked using the Vicon Vantage V5 camera system and sent via Wi-Fi to the embedded computer, where all code runs. Fig. 7 shows the lab experiment setup.

The simulations are run on a Dell XPS 15 laptop with a 12-core 2.60 GHz Intel i7-10750H CPU.

B. Results

This section presents simple simulation results focusing on the computational efficiency and size of the obstacle avoidance constraint sets in Section VI-B1 and the effect of the tuning parameter β in Section VI-B2. Thereafter, it describes the performance of the HMPC framework in lab experiments and simulations in Section VI-B3.

A video of the lab experiments and simulations is available at <https://youtu.be/0RnrKk6830I>.

1) *I-DecompUtil*: The purpose of this section is twofold. First, it shows how to make the runtime code for obstacle avoidance constraint generation efficient. Second, it demonstrates the impact of the bounding box width on the computation time and the time required to reach the goal.

It is important to note that *DecompUtil* iterates through grid map \mathcal{M} to find all occupied grid cells $\tilde{\mathcal{O}}$ that are contained in the bounding box \mathcal{B} around the line segment. After finding the occupied grid cells, *DecompUtil* iterates through them to compute the constraints for a specific line segment. *I-DecompUtil* repeats this process for all line segments between the stages in the horizon.

The time to compute the obstacle avoidance constraints varies significantly depending on the implementation. Table II shows

² [Online]. Available: <https://nxp.gitbook.io/hovergames/>

TABLE II
MEAN (STANDARD DEVIATION) OF CONSTRAINT COMPUTATION TIMES (MS)
OVER A SINGLE RUN FROM START TO GOAL POSITION

		Map pre-processing	
		No	Yes
Compiler optimization	No	2312.6 (165.0)	147.7 (69.2)
	Yes	6.1 (0.8)	5.5 (1.1)

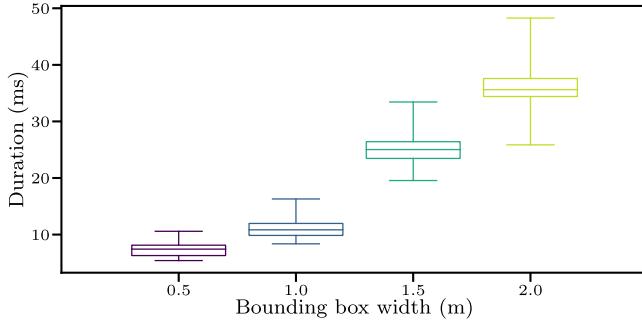


Fig. 8. Obstacle avoidance constraints computation times for different bounding box widths when moving from start to goal with an MPC with 100 stages and 0.05 s sampling time.

the mean and standard deviation of the timing results for the different combinations of two options to reduce computation time: C++ compiler optimization, i.e., setting compiler flag `-O3` to allow for loop unrolling, and map preprocessing, i.e., selecting only the relevant part corresponding to the line segment to reduce the number of occupied grid cells to iterate through. The results are obtained by moving from start to goal using an MPC with a horizon of 100 stages and a sampling time of 50 ms. A long horizon is chosen so that the impact of other, high-priority processor tasks on the constraints computation time is averaged out.

The results clearly illustrate that compiler optimization is the main factor in reducing constraints computation time, allowing for real-time feasibility. The larger the map, the more significant the impact of map preprocessing becomes, up to the point that map preprocessing becomes necessary to ensure real-time feasibility.

Another factor to consider is the bounding box width tuning parameter. This parameter affects the size of the bounding box \mathcal{B} and, correspondingly, the size of the preprocessed map. The primary purpose of \mathcal{B} is to limit the number of grid cells to consider for constraints computation. As a result, this parameter trades off the speed at which the MPC scheme converges to the goal, or can find a way to the goal at all, and the constraints computation time.

Fig. 8 shows the computation times for four different bounding box widths, given the scenario with two obstacles described in Section VI-B3. The figure illustrates that the computation time grows with the bounding box width but saturates after a certain width, depending on the obstacle shapes and density in the environment. The variation in timing depends on the environment as well.

Table III shows the time required to move from start to goal for SMPC and HMPC using different bounding box widths. The

TABLE III
TIME TO REACH GOAL WITH SMPC AND HMPC FOR DIFFERENT BOUNDING
BOX WIDTHS

		\mathcal{B} width (m)			
		0.5	1	1.5	2
Time to reach goal (s)	SMPC	52.3	52.3	52.3	52.3
	HMPC	14.8	6.4	7.3	6.9

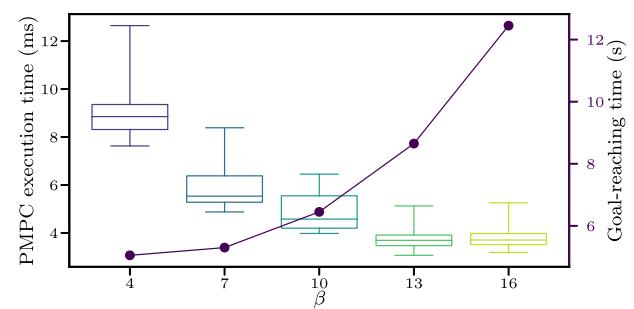


Fig. 9. PMPC execution times (boxplots) and goal-reaching times (line) for different values of β when moving from start to goal with HMPC. The TMPC is implemented as given in Table I. Given constant $T^{s,t}$, the PMPC is implemented using $T^{s,p} = \beta T^{s,t}$ with $\beta \in \{4, 7, 10, 13, 16\}$. Except for $\beta = 16$, T^p was chosen to be the closest to the value in Table I. For $\beta = 16$, the horizon has to be $T^{s,p}$ longer according to Remark 3, which causes a slight increase in computation time compared to $\beta = 13$.

SMPC has such a short horizon that the size of the obstacle avoidance constraints does not influence the goal-reaching time. On the other hand, for HMPC, a short bounding box width of 0.5 m limits the goal-reaching time. This effect vanishes in our environment setup for bounding box widths $\{1, 1.5, 2\}$ m. Instead, the specific shape of the constraint regions and the ability of the PMPC to plan more aggressive maneuvers that cause a slight goal overshooting play a more significant role in the goal-reaching times for these bounding box widths.

The combination of the results in Fig. 8 and Table III demonstrates the tradeoff above up to 1 m: the smaller the bounding box, the less computation time but the more limited the plan is, and vice versa. Since a bounding box width of 1 m gives a similar goal reaching time to the larger bounding box widths but with lower computation times, this value is used to generate the rest of the results in this section.

It is important to note that the above results are generated using the grid map as presented in Section VI-B3. This grid map is pregenerated with thin obstacle edge lines to make SMPC feasible in real time. However, in experiments with onboard perception, the grid map will contain thicker obstacle edges due to noise in sensor measurements (e.g., from a depth camera). Furthermore, these grid maps are usually constructed in 3-D instead of 2-D, which will significantly increase the number of occupied grid cells and enlarge the impact of map-preprocessing.

2) *Effect of β :* Next, we demonstrate the impact of β on the PMPC. Recall from Remark 3 that β trades off PMPC execution time and performance. This result is visible in Fig. 9: the larger β , the shorter the PMPC execution time and the longer the goal-reaching time. The longer goal-reaching time is caused by the coarser replanning time and the fact that the obstacle avoidance

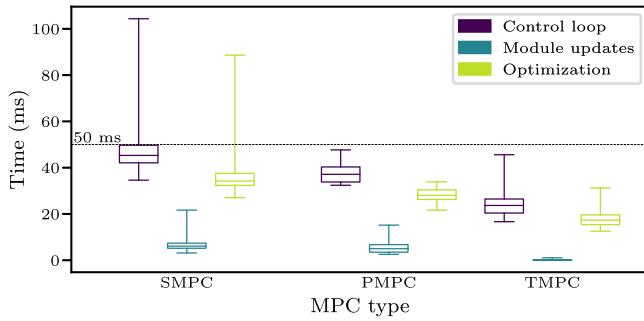


Fig. 10. Boxplot with overall execution times for SMPC and HMPC (*control loop*) and its most important parts: *Module updates* and *optimization*. The maximum control loop execution time is 50 ms for real-time feasibility of SMPC and TMPC and 500 ms for PMPC. Note that *module updates* includes constraints generation and loading new cost function terms, and is low compared to the optimization time. Since the TMPC receives the constraints from the PMPC, its *module updates* computation time is even lower. *Optimization* is the time between calling the solver and obtaining the solution. SMPC contains several control loop cycles exceeding its real-time limit, whereas PMPC and TMPC always finish on time.

constraints are more conservative for a coarser plan, meaning that the PMPC can plan less far ahead in each optimization step.

Based on Fig. 9, we have set $\beta = 10$ in the next section. Note that $\beta = 10$ allows for convenient results interpretation since $T^{s,p} = T^t$ given the numbers in Table I.

3) *HMPC*: This section describes the performance of HMPC compared to SMPC in lab experiments and highlights differences with simulation results when relevant. Two types of simulations are run: a simple simulation with perfect system knowledge, i.e., the simulated system is also (23) integrated using RK4 with sampling time $T^{s,t}$, and one in Gazebo, i.e., including model mismatch.

First, we need to know whether all MPC schemes are feasible on the embedded computer in real time. Fig. 10 shows the execution times of the control loop and its main components for SMPC, PMPC, and TMPC. PMPC and TMPC comply with execution time limits, whereas SMPC is around its maximum of 50 ms and sometimes exceeds the maximum. This problem cannot easily be avoided on the embedded computer since it has a limited number of threads. The lower the number of threads, the greater the risk of the control loop being interrupted by other operating system tasks. Since most of the SMPC control loops finish in time, the frequency at which control loops are scheduled catches up with the desired frequency after exceeding the limit. Consequently, the results still provide valuable insights.

To show the impact of the embedded hardware on the execution times, Table IV shows the mean and standard deviation of the control loop execution times for SMPC, PMPC, and TMPC in both simulations and lab experiments. The numbers clearly illustrate a drastic increase in mean and standard deviation of the control loop execution time when moving from simulations to lab experiments with embedded hardware.

The main limitation of SMPC is the computation time. Therefore, an SMPC with increasing discretization times was also tested in an attempt to reduce the number of stages while keeping the same planning horizon length. However, this scheme is not

TABLE IV
MEAN (STANDARD DEVIATION) OF CONTROL LOOP EXECUTION TIMES (MS) OF SMPC, PMPC, AND TMPC IN SIMPLE SIMULATIONS (SIM), GAZEBO SIMULATIONS (GAZ), AND LAB EXPERIMENTS USING EMBEDDED HARDWARE (LAB)

Method		System		
		sim	gaz	lab
SMPC	8.0 (1.2)	8.4 (1.3)	46.8 (7.3)	
PMPC	4.9 (0.9)	5.2 (0.9)	37.7 (4.2)	
TMPC	3.2 (0.6)	3.4 (0.7)	24.1 (4.9)	

TABLE V
TIME (S) TO MOVE FROM START TO GOAL FOR SMPC AND HMPC IN SIMPLE SIMULATIONS (SIM), GAZEBO SIMULATIONS (GAZ), AND LAB EXPERIMENTS (LAB)

Method		System		
		sim	gaz	lab
SMPC	52.3	107.5	35.1	
HMPC	6.4	8.4	8.1	

able to plan a trajectory around an obstacle as the line segments of two stages separated by a relatively large discretization time intersect with the obstacles when moving around them, i.e., Assumption 1 does not hold.

Note that since the PMPC has a sampling time of 500 ms and only needs 50 ms to solve the considered scenario, it has time left to account for processing a more realistic grid map and optimizing a plan in a more challenging environment.

Given the real-time feasibility of SMPC and HMPC, we can now analyze the schemes' closed-loop properties. Table V summarizes the goal-reaching times of SMPC and HMPC in all simulations and lab experiments. In this case, the goal is reached if the system enters the circle with a radius of 5 cm around the goal in (p^x, p^y) -plane. We can draw three conclusions based on the results. First, HMPC can move more efficiently from start to goal. Second, the presence of model mismatch significantly affects the goal-reaching times of SMPC, while HMPC is less sensitive to the presence of model mismatch. Third, the goal-reaching time of SMPC in lab experiments is significantly shorter than in simulations. Compared to simple simulations, the quadrotor moves quicker around the obstacle edges since the combination of model mismatch and slack results in higher accelerations, also in the direction of moving alongside the obstacle. Compared to the Gazebo simulations, the real quadrotor responds more quickly to control commands. This effect becomes more significant when the cost function gives less incentive to increase attitude, and consequently, acceleration. Therefore, the real quadrotor reaches the circle with a radius of 5 cm around the goal faster.

Fig. 11 shows the 2-D closed-loop trajectories of both SMPC and HMPC in the lab experiments. Furthermore, Fig. 12 displays the corresponding planning costs of SMPC and PMPC.

Related to the SMPC performance, we would like to highlight three aspects.

First of all, the *SMPC start and goal positions are closer* to each other compared to HMPC. With the original start and goal positions, SMPC gets stuck behind the first encountered obstacle. This is due to the short horizon that ends in

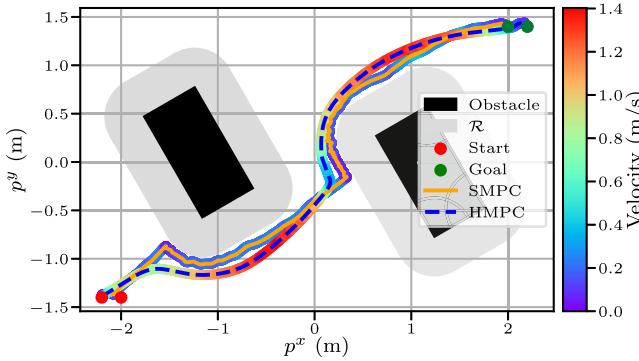


Fig. 11. 2-D closed-loop SMPC and HMPC trajectories with associated velocities in lab experiments. HMPC has a longer planning horizon and results in a smoother trajectory reaching the goal in 8.1 s versus 35.1 s for SMPC. Note that the start and goal positions for SMPC are closer. Otherwise, SMPC cannot find a way around the first encountered obstacle.

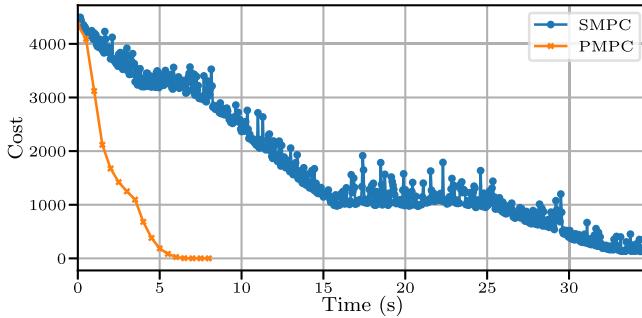


Fig. 12. Optimal costs for SMPC and PMPC runs. Due to its longer horizon, PMPC reduces its cost quickly, whereas SMPC takes longer to move around obstacles and is sensitive to model mismatch.

steady state and the fact that SMPC is a local method. After moving the start and goal positions, SMPC still struggles to find a way around the obstacles because of its short horizon. As a result, the trajectory slowly rotates around the obstacle corners, and it takes a long time to reach the goal. This is reflected in Fig. 12 by the “plateaus” in the cost function.

Second, SMPC is *sensitive to model mismatch*. Without constraints softening, a slight disturbance effect causes the system to exceed the constraints, making the optimization problem infeasible. With constraints softening, the system gets aggressively pushed back due to the high slack penalization term in the cost function. However, this increases the distance to the goal, causing the robot to move back toward the constraints, repeating the same process. This is visible in Fig. 11 by the oscillating trajectory near the obstacles and in Fig. 12 by the noisy cost values.

To prevent collisions despite this oscillating behavior, the safety margin, as described in Section VI-A7, is tuned to 0.1 m. Fig. 13 shows the maximum slack value that occurs in the prediction horizon of the different SMPC runs over time. The maximum value is 0.15 m, meaning that the predicted position in at least one of the stages exceeds the safety margin of 0.1 m. This happens in the last stage, not in the first stage,

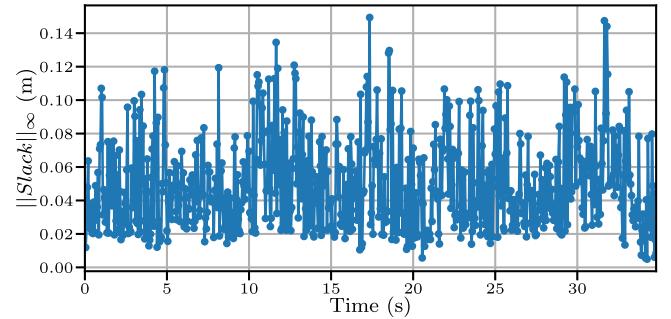


Fig. 13. Maximum SMPC slack in the horizon over time. The maximum slack value is around 0.15 m.

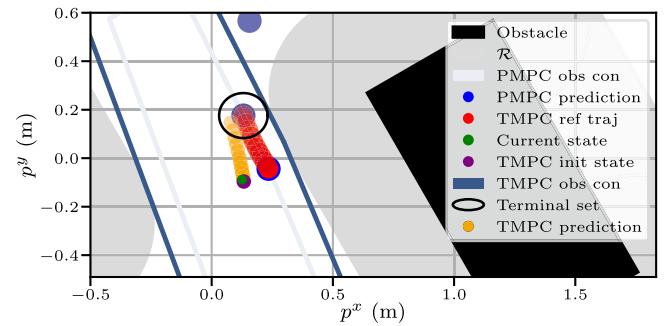


Fig. 14. Specific snapshot of the TMPC prediction starting from a forward-simulated state (TMPC init state) and ending in terminal set (12), given the obstacles, their collision region indicated by the inflation with the radius of the robot region R , PMPC prediction as reference plan and corresponding reference trajectory (TMPC ref traj), and obstacle avoidance constraints (PMPC obs con). Note that the PMPC obstacle avoidance constraints are tightened with respect to the TMPC obstacle avoidance constraints (TMPC obs con). Furthermore, the last measured state (current state) is plotted. The last measured state and forward-simulated state do not overlap, showing the presence of model mismatch.

meaning the robot does not collide with the obstacles. However, the SMPC scheme cannot provide strict safety guarantees in general.

Finally, SMPC is *not able to maintain altitude* during flight. This is caused by the fact that the SMPC needs to balance both the planning task, reaching the goal without collisions, and the tracking task, sending control commands to achieve stable flight, in a single formulation with a short horizon and in the presence of model mismatch, especially in thrust dynamics. In simple simulation, the z position varies between 1.39 and 1.41 m, in Gazebo between 1.37 and 1.67 m, and in lab experiments between 0.20 and 1.56 m.

Next, we would like to highlight six aspects related to the performance of HMPC.

First of all, to verify the implementation of the HMPC scheme, we checked the simple simulation results and concluded that the TMPC tracking error was zero at all times. This proves the *validity of Theorems 1 and 2*.

Second, in contrast to SMPC, HMPC gives a similar closed-loop trajectory and similar goal-reaching times in simple simulations, Gazebo simulations, and lab experiments. Therefore, it is *less sensitive to model mismatch compared to SMPC*.

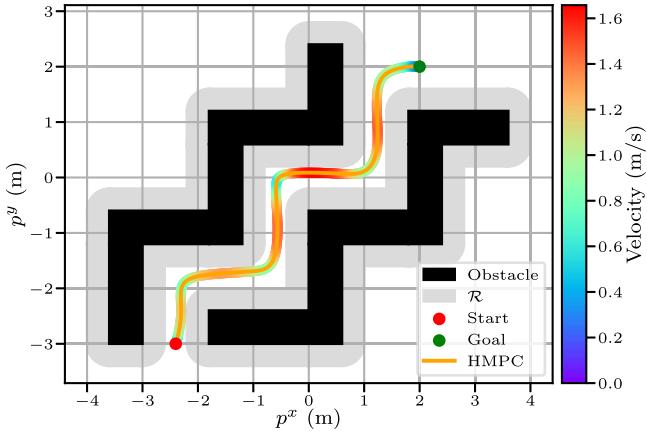


Fig. 15. 2-D closed-loop HMPC trajectory with associated velocity in a more complex Gazebo environment. HMPC can deal with more complex environments and finds its way to the goal in 8.4 s without collisions.

Third, the PMPC horizon is long enough to *plan a reasonable trajectory around the obstacles*. This is clearly shown by the quickly decreasing PMPC cost compared to the SMPC cost in Fig. 12. Due to its longer horizon, PMPC quickly finds a way toward the goal. On the other hand, SMPC needs time to move around the obstacles. The horizontal parts of the SMPC cost illustrate this effect. Moreover, the SMPC cost is noisier, which is caused by the model mismatch.

Fourth, the TMPC has a sufficiently long horizon to *satisfy the terminal set constraint in every run*, as is clearly illustrated for a particular run between the obstacles in Fig. 14.

Fifth, HMPC can *maintain altitude* during flight: in simple simulations, the z position varies between 1.37 and 1.40 m, in Gazebo between 1.39 and 1.42 m, and in lab experiments 1.32 and 1.43 m.

Finally, to show that HMPC *scales to more complex environments*, Fig. 15 illustrates the closed-loop trajectory of HMPC in a corridor-like Gazebo environment. Even though HMPC is a local method, it finds its way from start to goal in 8.4 s without colliding with the obstacles. The maximum PMPC and TMPC control loop execution times are 7.0 ms and 5.8 ms, respectively, meaning that HMPC is *runtime feasible*.

VII. CONCLUSION

Optimizing a trajectory complying with nonlinear system dynamics and avoiding collisions is a computationally expensive task. Consequently, it is not feasible to run an SMPC scheme in real time that generates obstacle avoidance constraints and optimizes the corresponding trajectory with a reasonable horizon length on an embedded computer.

To address this problem, we propose a novel HMPC scheme, which includes the codesign of a PMPC and TMPC. In the offline phase, we compute the terminal ingredients of the TMPC. These are also utilized to adjust the constraints in the PMPC during runtime. By ensuring the continuity of the reference plan and consistency in updating the obstacle avoidance constraints, the PMPC can run independently from the current system state and

construct a plan with a long time horizon. This saves constraints generation and optimization time on the side of the TMPC. Consequently, the TMPC runs at a higher frequency, resulting in accurate tracking of the reference plan.

We compared HMPC to SMPC in simulations without model mismatch, Gazebo simulations, and lab experiments. In general, the prediction horizon of SMPC needs to be reduced significantly to achieve real-time feasibility in the lab experiments. Therefore, SMPC has difficulty finding a way to the goal, needs constraints softening to remain feasible, and deviates significantly from the desired altitude. On the other hand, HMPC optimizes a smoother plan, does not require constraints softening, and maintains altitude more accurately.

To summarize, HMPC is more computationally efficient and less sensitive to model mismatch than SMPC. HMPC also provides tracking and recursive feasibility guarantees. It simplifies the task for robotic practitioners so they can leverage off-the-shelf nonlinear solvers for motion planning and tracking without having to optimize code, linearize models, or take care of safety themselves. Furthermore, due to its computational efficiency, HMPC can easily be deployed on different robotic platforms.

While the experiments and simulations demonstrate collision avoidance and recursive feasibility, the derived theoretical guarantees are only valid in the absence of model mismatch. Our next step is to extend the TMPC to a *robust* formulation that can explicitly leverage predetermined bounds on the model mismatch, similar to the ideas in [65]. Lastly, we would like to move away from motion capture cameras to onboard sensors, such as LiDAR or depth cameras, for obstacle detection and 3-D constraint generation, e.g., in outdoor scenarios.

APPENDIX A PROOF THEOREM 1

Proof: The proof consists of three parts. *Part I* proposes the previously optimal solution shifted by $T^s = T^{s,t}$, appended with the state after applying κ^f , as candidate solution with horizon $T = T^t$ for the next TMPC run. *Part II* shows that all system and obstacle constraints are satisfied for the candidate solution. Finally, *Part III* shows asymptotic convergence of the state to the reference trajectory.

Part I. Candidate solution

For convenience, define for $\tau \in [T, T + T^s]$

$$\mathbf{u}_{\tau|t}^* := \kappa^f(\mathbf{x}_{\tau|t}^*, \mathbf{r}_{\tau|t}^*) \quad (26)$$

and $\mathbf{x}_{\tau|t}^*$ according to (5c) by applying (26).

Consider the following candidate solution:

$$\tilde{\mathbf{u}}_{\tau|t+T^s} = \mathbf{u}_{\tau+T^s|t}^*, \quad \tilde{\mathbf{x}}_{\tau|t+T^s} = \mathbf{x}_{\tau+T^s|t}^*, \quad \tau \in [0, T] \quad (27)$$

with the reference trajectory defined by (18) in Property 1.

Part II: Recursive feasibility

Part II-I. System constraints satisfaction: For $\tau \in [0, T - T^s], j \in \mathbb{N}_{[1, n^s]}$, the system constraints given the candidate solution satisfy

$$g_j^s(\tilde{\mathbf{x}}_{\tau|t+T^s}, \tilde{\mathbf{u}}_{\tau|t+T^s}) \stackrel{(27)}{=} g_j^s(\mathbf{x}_{\tau+T^s|t}^*, \mathbf{u}_{\tau+T^s|t}^*) \stackrel{(5d)}{\leq} 0. \quad (28)$$

Part II-II. Obstacle avoidance: Similarly, by Property 2(b), for $\tau \in (0, T - T^s]$, $j \in \mathbb{N}_{[1, n^o]}$, the obstacle avoidance constraints given the candidate satisfy

$$\begin{aligned} g_{j, \tau+T^s|t}^o(\mathbf{p}_{\tau+T^s|t}^*) &\stackrel{(5e)(22)}{\leq} 0 \\ \stackrel{(27)}{\Rightarrow} g_{j, \tau|t+T^s}^o(\tilde{\mathbf{p}}_{\tau|t+T^s}) &\leq 0 \end{aligned} \quad (29)$$

thus ensuring obstacle avoidance by Property 2(a).

Part II-III. Terminal constraints satisfaction: Given candidate (27) and reference (18), the following holds for $\tau \in [T - T^s, T]$:

$$\begin{aligned} \mathcal{J}^{f,t}(\tilde{\mathbf{x}}_{\tau|t+T^s}, \mathbf{x}_{\tau|t+T^s}^r) &\stackrel{(27)}{=} \mathcal{J}^{f,t}(\mathbf{x}_{\tau+T^s|t}^*, \mathbf{x}_{\tau|t+T^s}^r) \\ &\stackrel{(19a)}{\leq} \mathcal{J}^{f,t}(\mathbf{x}_{T|t}^*, \mathbf{x}_{T|t}^r) \stackrel{(5f)}{\leq} \alpha^2 \end{aligned} \quad (30)$$

i.e., the terminal set $(\tilde{\mathbf{x}}_{\tau|t+T^s} - \mathbf{x}_{\tau|t+T^s}^r) \in \mathcal{X}^{t,f}$ is positive invariant. Therefore, by Proposition 1, the following holds for $\tau \in [T - T^s, T]$:

$$g_j^s(\tilde{\mathbf{x}}_{\tau|t+T^s}, \tilde{\mathbf{u}}_{\tau|t+T^s}) \stackrel{(19b)}{\leq} 0, \quad j \in \mathbb{N}_{[1, n^s]} \quad (31)$$

$$g_{j, \tau|t+T^s}^o(\tilde{\mathbf{p}}_{\tau|t+T^s}) \stackrel{(19c)}{\leq} 0, \quad j \in \mathbb{N}_{[1, n^o]}. \quad (32)$$

Part III. Asymptotic convergence

Asymptotic convergence is shown using Barbalat's Lemma. By (19a), the following holds:

$$\begin{aligned} &\mathcal{J}^{f,t}(\tilde{\mathbf{x}}_{T|t+T^s} - \mathbf{x}_{T|t+T^s}^r) - \mathcal{J}^{f,t}(\mathbf{x}_{T|t}^* - \mathbf{x}_{T|t}^r) \\ &+ \int_{\tau=T-T^s}^T \mathcal{J}^{s,t}(\mathbf{x}_{\tau+T^s|t}^*, \mathbf{x}_{\tau+T^s|t}^r, \mathbf{r}_{\tau+T^s|t}) d\tau \leq 0. \end{aligned} \quad (33)$$

Therefore, the following inequality is established:

$$\begin{aligned} &\mathcal{J}^{*,t}(\mathbf{x}_{t+T^s}, \mathbf{x}_{\cdot|t+T^s}^r) \\ &\leq \int_{\tau=0}^T \mathcal{J}^{s,t}(\tilde{\mathbf{x}}_{\tau|t+T^s}, \tilde{\mathbf{u}}_{\tau|t+T^s}, \mathbf{r}_{\tau|t+T^s}) d\tau \\ &+ \mathcal{J}^{f,t}(\tilde{\mathbf{x}}_{T|t+T^s} - \mathbf{x}_{T|t+T^s}^r) \\ &\stackrel{(33)}{\leq} \mathcal{J}^{*,t}(\mathbf{x}_t, \mathbf{x}_{\cdot|t}^r) - \int_{\tau=0}^{T^s} \mathcal{J}^{s,t}(\mathbf{x}_{\tau|t}^*, \mathbf{x}_{\tau|t}^r, \mathbf{r}_{\tau|t}) d\tau \end{aligned} \quad (34)$$

which proves

$$\begin{aligned} &\mathcal{J}^{*,t}(\mathbf{x}_{t+T^s}, \mathbf{x}_{\cdot|t+T^s}^r) - \mathcal{J}^{*,t}(\mathbf{x}_t, \mathbf{x}_{\cdot|t}^r) \\ &\leq -c^{\mathcal{J},d} \int_{\tau=t}^{t+T^s} \|\mathbf{x}_\tau - \mathbf{x}_\tau^r\|_Q^2 d\tau \end{aligned} \quad (35)$$

using $Q, R \succ 0$ with $c^{\mathcal{J},d} > 0$. Iterating this inequality and using the fact that $\mathcal{J}^{*,t}(\mathbf{x}_t, \mathbf{x}_{\cdot|t}^r)$ is uniformly bounded yields

$$\lim_{t \rightarrow \infty} \int_0^t \|\mathbf{x}_\tau - \mathbf{x}_\tau^r\|_Q^2 d\tau < \infty. \quad (36)$$

Asymptotic convergence follows by invoking Barbalat's Lemma. \blacksquare

APPENDIX B PROOF THEOREM 2

Proof: The proof consists of two parts. *Part I* proposes the previously optimal solution shifted by $T^s = T^{s,p}$, appended with an input ensuring steady state, as candidate solution with horizon $T = T^p$ for the next PMPC run. *Part II* shows feasibility of the candidate solution given the previously optimal solution.

Part I: Candidate solution

Consider the following candidate solution, valid by (20f):

$$\tilde{\mathbf{u}}_{\tau|t+T^s} = \begin{cases} \mathbf{u}_{\tau+T^s|t}^*, & \tau \in [0, T - T^s] \\ \mathbf{u}_{T|t}^*, & \tau \in [T - T^s, T] \end{cases} \quad (37a)$$

$$\tilde{\mathbf{x}}_{\tau|t+T^s} = \begin{cases} \mathbf{x}_{\tau+T^s|t}^*, & \tau \in [0, T - T^s] \\ \mathbf{x}_{T|t}^*, & \tau \in [T - T^s, T]. \end{cases} \quad (37b)$$

Part II: Recursive feasibility

Part II-I. System constraints satisfaction: For $\tau \in [0, T - T^s]$, $j \in \mathbb{N}_{[1, n^s]}$, the system constraints given the candidate solution are

$$g_j^s(\tilde{\mathbf{x}}_{\tau|t+T^s}, \tilde{\mathbf{u}}_{\tau|t+T^s}) \stackrel{(37)}{=} g_j^s(\mathbf{x}_{\tau+T^s|t}^*, \mathbf{u}_{\tau+T^s|t}^*) \stackrel{(20d)}{\leq} 0. \quad (38)$$

Part II-II. Obstacle avoidance constraints satisfaction: Similarly, by Property 2(b) for $\tau \in (0, T - T^s]$, $j \in \mathbb{N}_{[1, n^o]}$, the obstacle avoidance constraints given the candidate satisfy

$$\begin{aligned} g_{j, \tau+T^s|t}^o(\mathbf{p}_{\tau+T^s|t}^*) &\stackrel{(20e)(22)(16)}{\leq} 0 \\ \stackrel{(37b)}{\Rightarrow} g_{j, \tau|t+T^s}^o(\tilde{\mathbf{p}}_{\tau|t+T^s}) &\leq 0. \end{aligned} \quad (39)$$

Part II-III. Terminal constraints satisfaction: Given the steady-state condition in candidate (37), (38), and (39) hold for $\tau \in [T - T^s, T]$ as well. \blacksquare

ACKNOWLEDGMENT

The authors would like to thank Chris Pek for helping to create Fig. 1, and Clarence Chen, Johanna Probst, and Simone Carena for helping to take the pictures in Fig. 7. All content represents the opinion of the author(s), which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*. Cambridge, MA, USA: MIT Press, 2011.
- [2] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, “A survey of motion planning and control techniques for self-driving urban vehicles,” *IEEE Trans. Intell. Veh.*, vol. 1, no. 1, pp. 33–55, Mar. 2016.
- [3] J. Delmerico et al., “The current state and future outlook of rescue robotics,” *J. Field Robot.*, vol. 36, no. 7, pp. 1171–1191, 2019.
- [4] B. Brito, M. Everett, J. P. How, and J. Alonso-Mora, “Where to go next: Learning a subgoal recommendation policy for navigation in dynamic environments,” *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 4616–4623, Jul. 2021.
- [5] M. Lodel, B. Brito, A. Serra-Gómez, L. Ferranti, R. Babuška, and J. Alonso-Mora, “Where to look next: Learning viewpoint recommendations for informative trajectory planning,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 4466–4472.
- [6] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, vol. 2. Madison, WI, USA: Nob Hill Publishing, 2017.

[7] M. Kamel, M. Burri, and R. Siegwart, "Linear vs nonlinear MPC for trajectory tracking applied to rotary wing micro aerial vehicles," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3463–3469, 2017.

[8] 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.

[9] M. A. Santos, A. Ferramosca, and G. V. Raffo, "Nonlinear model predictive control schemes for obstacle avoidance," *J. Control, Autom. Elect. Syst.*, vol. 34, no. 5, pp. 891–906, 2023.

[10] M. Neunert et al., "Fast nonlinear model predictive control for unified trajectory optimization and tracking," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2016, pp. 1398–1404.

[11] B. Brito, B. Floor, L. Ferranti, and J. Alonso-Mora, "Model predictive contouring control for collision avoidance in unstructured dynamic environments," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 4459–4466, Oct. 2019.

[12] O. de Groot, B. Brito, L. Ferranti, D. Gavrila, and J. Alonso-Mora, "Scenario-based trajectory optimization in uncertain dynamic environments," *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 5389–5396, Jul. 2021.

[13] D. Saccani, L. Cecchin, and L. Fagiano, "Multitrajectory model predictive control for safe UAV navigation in an unknown environment," *IEEE Trans. Control Syst. Technol.*, vol. 31, no. 5, pp. 1982–1997, Sep. 2023.

[14] P. Falcone, F. Borrelli, H. E. Tseng, J. Asgari, and D. Hrovat, "A hierarchical model predictive control framework for autonomous ground vehicles," in *Proc. IEEE Amer. Control Conf.*, 2008, pp. 3719–3724.

[15] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, H. E. Tseng, and F. Borrelli, "Predictive control for agile semi-autonomous ground vehicles using motion primitives," in *Proc. IEEE Amer. Control Conf.*, 2012, pp. 4239–4244.

[16] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1:43 scale RC cars," *Optimal Control Appl. Methods*, vol. 36, no. 5, pp. 628–647, 2015.

[17] B. Elsayed and R. Findeisen, "Generic motion primitives-based safe motion planner under uncertainty for autonomous navigation in cluttered environments," in *Proc. IEEE 29th Int. Conf. Inf., Commun. Autom. Technol.*, 2023, pp. 1–6.

[18] M. Ibrahim, J. Matschek, B. Morabito, and R. Findeisen, "Hierarchical model predictive control for autonomous vehicle area coverage," *IFAC-PapersOnLine*, vol. 52, no. 12, pp. 79–84, 2019.

[19] M. Kögel, M. Ibrahim, C. Kallies, and R. Findeisen, "Safe hierarchical model predictive control and planning for autonomous systems," *Int. J. Robust Nonlinear Control*, vol. 35, no. 7, pp. 2658–2676, 2023.

[20] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. Robot. Res.*, vol. 5, no. 1, pp. 90–98, 1986.

[21] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Autom. Mag.*, vol. 4, no. 1, pp. 23–33, Mar. 1997.

[22] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robot. Res.*, vol. 30, no. 7, pp. 846–894, 2011.

[23] 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, Apr. 2022.

[24] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Proc. Robot. Research: 16th Int. Symp.*, 2016, pp. 649–666.

[25] L. Quan, Z. Zhang, X. Zhong, C. Xu, and F. Gao, "EVA-planner: Environmental adaptive quadrotor planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 398–404.

[26] B. Zhou, F. Gao, L. Wang, C. Liu, and S. Shen, "Robust and efficient quadrotor trajectory generation for fast autonomous flight," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3529–3536, Oct. 2019.

[27] C. Zhou, B. Huang, and P. Fränti, "A review of motion planning algorithms for intelligent robots," *J. Intell. Manuf.*, vol. 33, no. 2, pp. 387–424, 2022.

[28] A. Orthey, C. Chamzas, and L. E. Kavraki, "Sampling-based motion planning: A comparative review," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 7, pp. 285–310.

[29] L. Li, J. Li, and S. Zhang, "State-of-the-art trajectory tracking of autonomous vehicles," *Mech. Sci.*, vol. 12, no. 1, pp. 419–432, 2021.

[30] C. Papachristos et al., "Modeling, control, state estimation and path planning methods for autonomous multirotor aerial robots," *Found. Trends Robot.*, vol. 7, no. 3, pp. 180–250, 2018.

[31] S. Bouabdallah, A. Noth, and R. Siegwart, "PID vs LQ control techniques applied to an indoor micro quadrotor," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 2451–2456.

[32] P. Foehn and D. Scaramuzza, "Onboard state dependent LQR for agile quadrotors," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 6566–6572.

[33] M. M. Tobenkin, I. R. Manchester, and R. Tedrake, "Invariant funnels around trajectories using sum-of-squares programming," *IFAC Proc. Volumes*, vol. 44, no. 1, pp. 9218–9223, 2011.

[34] J. L. Vázquez, M. Brühlmeier, A. Liniger, A. Rupenyan, and J. Lygeros, "Optimization-based hierarchical motion planning for autonomous racing," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 2397–2403.

[35] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Sci. Robot.*, vol. 6, no. 56, 2021, Art. no. eabh1221.

[36] J. Li, M. Ran, and L. Xie, "Design and experimental evaluation of a hierarchical controller for an autonomous ground vehicle with large uncertainties," *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 3, pp. 1215–1227, May 2022.

[37] L. Numerow, A. Zanelli, A. Carron, and M. N. Zeilinger, "Inherently robust suboptimal mpc for autonomous racing with anytime feasible sqp," *IEEE J. Robot. Autom. Lett.*, IEEE, 2024, *arXiv:2401.02194*.

[38] M. Krinner, A. Romero, L. Bauersfeld, M. Zeilinger, A. Carron, and D. Scaramuzza, "Model Predictive Contouring Control for Time-Optimal Flight with Safety Constraints," in *Proc. Robot.: Sci. Syst.*, Delft, The Netherlands, 2024.

[39] R. Tedrake, I. R. Manchester, M. Tobenkin, and J. W. Roberts, "LQR-trees: Feedback motion planning via sums-of-squares verification," *Int. J. Robot. Res.*, vol. 29, no. 8, pp. 1038–1052, 2010.

[40] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *Int. J. Robot. Res.*, vol. 36, no. 8, pp. 947–982, 2017.

[41] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *Int. J. Robot. Res.*, vol. 18, no. 6, pp. 534–555, 1999.

[42] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan, "Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots," *Int. J. Robot. Res.*, vol. 39, no. 12, pp. 1419–1469, 2020.

[43] M. K. M. Jaffar and M. Otte, "PiP-X: Online feedback motion planning/replanning in dynamic environments using invariant funnels," *Int. J. Robot. Res.*, 2022, Art. no. *arXiv:2202.00772*.

[44] M. M. Nicotra and E. Garone, "The explicit reference governor: A general framework for the closed-form control of constrained nonlinear systems," *IEEE Control Syst. Mag.*, vol. 38, no. 4, pp. 89–107, Aug. 2018.

[45] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.

[46] U. Rosolia, A. Singletary, and A. D. Ames, "Unified multirate control: From low-level actuation to high-level planning," *IEEE Trans. Autom. Control*, vol. 67, no. 12, pp. 6627–6640, Dec. 2022.

[47] N. Csomay-Shanklin, A. J. Taylor, U. Rosolia, and A. D. Ames, "Multi-rate planning and control of uncertain nonlinear systems: Model predictive control and control Lyapunov functions," in *Proc. IEEE 61st Conf. Decis. Control*, 2022, pp. 3732–3739.

[48] M. Chen et al., "FaSTrack: A modular framework for real-time motion planning and guaranteed safe tracking," *IEEE Trans. Autom. Control*, vol. 66, no. 12, pp. 5861–5876, Dec. 2021.

[49] D. Althoff, M. Althoff, and S. Scherer, "Online safety verification of trajectories for unmanned flight with offline computed robust invariant sets," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 3470–3477.

[50] D. Fridovich-Keil, S. L. Herbert, J. F. Fisac, S. Deglurkar, and C. J. Tomlin, "Planning, fast and slow: A framework for adaptive real-time safe trajectory planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 387–394.

[51] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for non-linear systems," *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.

[52] S. Singh, B. Landry, A. Majumdar, J.-J. Slotine, and M. Pavone, "Robust feedback motion planning via contraction theory," *Int. J. Robot. Res.*, vol. 42, no. 9, pp. 655–688, 2023.

[53] S. Singh, M. Chen, S. L. Herbert, C. J. Tomlin, and M. Pavone, "Robust tracking with model mismatch for fast and safe planning: An SOS optimization approach," in *Proc. Algorithmic Found. Robot. XIII: Proc. 13th Workshop Algorithmic Found. Robot.*, 2020, pp. 545–564.

- [54] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Trans. Control Syst. Technol.*, vol. 15, no. 3, pp. 566–580, May 2007.
- [55] J. Köhler, M. A. Müller, and F. Allgöwer, "Analysis and design of model predictive control frameworks for dynamic operation—An overview," *Annu. Rev. Control*, vol. 57, 2024, Art. no. 100929.
- [56] J. Köhler, M. A. Müller, and F. Allgöwer, "A nonlinear model predictive control framework using reference generic terminal ingredients," *IEEE Trans. Autom. Control*, vol. 65, no. 8, pp. 3576–3583, Aug. 2020.
- [57] S. Liu et al., "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, Jul. 2017.
- [58] A. Domahidi and J. Jerez, "Forcespro," 2014. [Online]. Available: <https://forces.embotech.com/>
- [59] A. Zanelli, A. Domahidi, J. Jerez, and M. Morari, "Forces NLP: An efficient implementation of interior-point methods for multistage nonlinear nonconvex programs," *Int. J. Control.*, vol. 93, pp. 13–29, 2017.
- [60] E. Cervera, "Try to start it! the challenge of reusing code in robotics research," *IEEE Robot. Autom. Lett.*, vol. 4, no. 1, pp. 49–56, Jan. 2019.
- [61] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Philadelphia, PA, USA: SIAM, 1994.
- [62] S. Yu, M. Reble, H. Chen, and F. Allgöwer, "Inherent robustness properties of quasi-infinite horizon nonlinear model predictive control," *Automatica*, vol. 50, no. 9, pp. 2269–2280, 2014.
- [63] P. J. Huber, "Robust estimation of a location parameter," in *Breakthroughs in Statistics: Methodology and Distribution*. Berlin, Germany: Springer, 1992, pp. 492–518.
- [64] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge Univ. Press, 2004.
- [65] J. Köhler, R. Soloperto, M. A. Müller, and F. Allgöwer, "A computationally efficient robust model predictive control framework for uncertain nonlinear systems," *IEEE Trans. Autom. Control*, vol. 66, no. 2, pp. 794–801, Feb. 2021.



Dennis Benders (Graduate Student Member, IEEE) received the B.Sc. degree in electrical engineering and the M.Sc. degree in embedded systems from the Delft University of Technology, Delft, The Netherlands, in 2018 and 2020, respectively. He is currently working toward the Ph.D. degree in robotics from the Delft University of Technology.

He is currently with the Reliable Robot Control Lab and Autonomous Multi-Robots Lab, Delft University of Technology. His research interests include motion planning and control for autonomous mobile robots, with special emphasis on real-time and open-source code.



Johannes Köhler (Member, IEEE) received the Ph.D. degree in mechanical engineering from the University of Stuttgart, Stuttgart, Germany, in 2021.

He is currently a Postdoctoral Researcher with ETH Zürich, Zürich, Switzerland. His research interests are in the area of model predictive control and control and estimation for nonlinear uncertain systems.

Dr. Köhler was the recipient of the 2021 European Systems and Control Ph.D. Thesis Award, the IEEE CSS George S. Axelby Outstanding Paper Award 2022, and the Journal of Process Control Paper Award 2023.



Thijs Niesten received the B.Sc. and M.Sc. degrees in mechanical engineering from the Delft University of Technology, Delft, The Netherlands, in 2018 and 2022, respectively.

He is currently a Research Engineer in Robotics. He is with the Reliable Robot Control Lab and Autonomous Multi-Robots Lab, Delft University of Technology. His research interests include computer vision, motion planning, and control for autonomous mobile robots.



Robert Babuška received the M.Sc. (Hons.) degree in electrical engineering from the Czech Technical University in Prague, Prague, Czechia, in 1990 and the Ph.D. (*cum laude*) in automatic control degree from the Delft University of Technology, (TU Delft), Delft, The Netherlands, in 1997.

He is currently a Full Professor with Faculty of Mechanical Engineering, Department of Cognitive Robotics, TU Delft, and a distinguished Researcher with CIIRC, Department of AI, Czech Technical University in Prague, Prague, Czechia. His research interests include reinforcement learning, adaptive and learning robot control, and nonlinear system identification. He has been involved in the applications of these techniques in various fields, ranging from process control to robotics and aerospace.



Javier Alonso-Mora (Senior Member, IEEE) received the Ph.D. degree in robotics from ETH Zürich, Zürich, Switzerland in 2014.

He is currently an Associate Professor with the Cognitive Robotics Department, Delft University of Technology, Delft, The Netherlands, where he leads the Autonomous Multi-Robots Lab. He was a Postdoctoral Associate with the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. His research focuses on navigation, motion planning, learning, and control of autonomous mobile robots

and teams, interacting with humans in dynamic, uncertain environments.

Dr. Alonso-Mora was the recipient of multiple awards, including the IEEE T-ASE Best Paper Award (2024), an ERC Starting Grant (2021), and the ICRA Best Paper Award on Multi-Robot Systems (2019).



Laura Ferranti (Senior Member, IEEE) received the Ph.D. degree in systems and control (alternatively: control engineering) from the Delft University of Technology, Delft, The Netherlands, in 2017.

She is currently an Assistant Professor with the Cognitive Robotics Department, Delft University of Technology. Her research interests include optimization and optimal control, model predictive control, reinforcement learning, and embedded optimization-based control with applications in flight control, maritime transportation, robotics, and automotive.

Dr. Ferranti was the recipient of an NWO Veni Grant from the Netherlands Organization for Scientific Research (2020) and of the Best Paper Award in Multi-robot Systems at ICRA 2019.