

# Time-Optimal Online Replanning for Agile Quadrotor Flight

Angel Romero, Robert Penicka, and Davide Scaramuzza

**Abstract**—In this paper, we tackle the problem of flying a quadrotor using time-optimal control policies that can be replanned online when the environment changes or when encountering unknown disturbances. This problem is challenging as the time-optimal trajectories that consider the full quadrotor dynamics are computationally expensive to generate, on the order of minutes or even hours. We introduce a sampling-based method for efficient generation of time-optimal paths of a point-mass model. These paths are then tracked using a Model Predictive Contouring Control approach that considers the full quadrotor dynamics and the single rotor thrust limits. Our combined approach is able to run in real-time, being the first time-optimal method that is able to adapt to changes *on-the-fly*. We showcase our approach's adaption capabilities by flying a quadrotor at more than 60 km/h in a racing track where gates are moving. Additionally, we show that our online replanning approach can cope with strong disturbances caused by winds of up to 68 km/h.

**Index Terms**—Aerial systems: Applications, integrated planning and control, motion and path planning,

SUPPLEMENTARY MATERIAL

Video of the experiments: <https://youtu.be/zBVpx3bgI6E>

## I. INTRODUCTION

QUADROTORS are one of the most versatile flying robots. Depending on the task at hand, their design can vary from very robust to extremely fast and powerful agile machines. When it comes to speed, racing quadrotors are among the fastest and most agile flying devices ever built [1]. Winning a drone race requires a set of visual, coordination and motor skills that are only achieved by the most dexterous and experienced human pilots [2]. For this reason, finding algorithms that conquer the sport of drone racing would represent a giant step forward for the entire robotics community. Thus, drone racing research has not only been feeding from the latest advances in perception, learning, planning, and control, but it has also been a strong contributor to these fields [3], [4], [5], [6], [7], [8], [9], [10]. A backbone piece in the drone racing is time-optimal (i.e., minimum-time) quadrotor flight through a series of waypoints (i.e., gates). An extremely challenging and currently missing piece of the puzzle is solving this problem *on-the-fly*.

Manuscript received: February, 24, 2022; Revised April, 19, 2022; Accepted June, 7, 2022. This paper was recommended for publication by Editor T. Ogata upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the National Centre of Competence in Research (NCCR) Robotics through the Swiss National Science Foundation (SNSF), and the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 871479 (AERIAL-CORE) and the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).

The authors are with the Robotics and Perception Group, University of Zurich, Switzerland (<http://rpg.ifi.uzh.ch>). (email: roagui@ifi.uzh.ch)

Digital Object Identifier (DOI): see top of this page.



Fig. 1. The proposed algorithm is able to adapt *on-the-fly* when encountering unknown disturbances. In the figure we show a quadrotor platform flying at speeds of more than 60 km/h. Thanks to our online replanning method, the drone can adapt to wind disturbances of up to 68 km/h while flying as fast as possible.

Planning a minimum-time trajectory is a complex problem. Algorithms need to both come up with the exact sequence of positions, velocities, orientations, inputs, etc. of the platform at every potential situation, and their allocation in time. Additionally, for this sequence to be tracked by classical control methods, it needs to be dynamically feasible, i.e., it needs to fulfill the complex non-linear dynamics of the quadrotor platform, and its actuator constraints.

Different approaches have attempted to tackle this problem. In [5], the authors propose the first method that finds time-optimal trajectories for a given series of waypoints and show how they beat professional human pilots in a drone race. However, these trajectories take several hours to compute, and hence their use is prohibitive for any online adaption to unpredicted changes (e.g., moving gates or wind gusts). To overcome this limitation, we need an algorithm to generate time-optimal trajectories in real-time. In [6], a Model Predictive Contouring Control (MPCC) method is introduced that is able to track non-feasible paths in a near-time-optimal fashion. By solving the difficult time-allocation problem online, at every time step, the MPCC method optimally selects the states and inputs that maximize the progress along the given path, while minimizing the Euclidean distance to it. The authors compare to [5] and show that the resulting policies are very close to time-optimality, opening the possibility of using as reference simplified paths that are faster to generate.

Our contributions are as follows. We introduce an efficient sampling-based method that can plan time-optimal paths for

point-mass model in real-time. Additionally, we extend the implementation of Model Predictive Contouring Control [6] (currently the state of the art in time-optimal flight) to make it suitable for online replanning. Our approach exploits the ability of MPCC to track non-feasible paths and uses efficiently-generated, time-optimal point-mass-model (PMM) trajectories to enable online replanning, being the first minimum time control method that is able to adapt to changes *on-the-fly*. We show in real world experiments that our approach is able to fly through a series of gates in minimum time, while the gates are continuously moved. Additionally, we show that the online replanning ability offers extra robustness to unknown, external disturbances by optimally flying through a strong stream of wind blowing up to 68 km/h.

## II. RELATED WORK

Our approach combines two well known concepts in quadrotor research: online trajectory replanning and time-optimal flight.

### A. Quadrotor online trajectory replanning

The existing approaches for quadrotor trajectory planning can be mostly categorized as polynomial approaches, sampling-based approaches, and optimization methods.

The idea of representing trajectories as continuous time polynomials has been a recurrent topic of research in quadrotor planning [11], [12], [4]. This is possible due to the fact that quadrotors are differentially flat systems [13], which allows for sampling the quadcopter inputs from independent polynomials in the  $x$ ,  $y$ ,  $z$  and yaw axes. Polynomials are very computationally efficient and thus allow online replanning. However, polynomial trajectory representation cannot render time-optimal policies, since the polynomials are smooth by definition.

Other methods that also take advantage of the differential flatness property are [14], [15], [16]. These methods use polynomial interpolations (splines) for trajectory optimization, and therefore, are also efficient enough for online replanning. However, even if they pursue fast flight, they jointly optimize perception objectives, smoothness, dynamic feasibility, collisions, etc., making their performance far from being minimum time.

### B. Time-optimal quadrotor flight

Time optimal quadrotor flight has been a very active area of research in the last years. The oldest approaches that have worked on this direction are [17], [18], but they have strong limitations. Either they reduce the problem to 2 dimensions, or make assumptions under which the time-optimal input policy is a pure bang-bang acceleration. More recent approaches [19] use the arc length to frame the problem independently of the time variable. However, they also suffer from limitations, since they simplify the limits of the platform to collective thrust and body rates, instead of single rotor thrusts.

The first approach that takes into account the full quadrotor dynamics to plan a pure time-optimal trajectory through a series of waypoints is [5]. They consider the time-optimal trajectory generation problem in a discretized way and use optimization-based methods to assign a time allocation and a full state to each trajectory point. This method has shown to result in very good performance, however, it suffers from

a major drawback. This algorithm takes a considerably long amount of time to converge (from 30 minutes to several hours, depending on the track length). This cripples its ability to adapt to changes or to model mismatches or unknown disturbances on the fly, bringing up two limitations. The first being the need for a control safety margin. The second being its inability to adapt to changes in the track or to major disturbances. In [6], the first limitation is tackled. Since the time allocation problem is solved online, this approach can optimally select the states and inputs in the prediction horizon at every timestep in order to maximize progress along the track, and is therefore able to make full use of the quadrotor's actuation power even in the presence of model mismatch. Another major benefit of [6] is that it only needs a continuously differentiable 3D path to track, and not a fully feasible trajectory. This allows for the use of simplified paths as reference, such as a time-optimal point mass model trajectory.

## III. METHODOLOGY

### A. Quadrotor Dynamics

The quadrotor's state space is described from the inertial frame  $I$  to the body frame  $B$ , as  $\mathbf{x} = [\mathbf{p}_{IB}, \mathbf{q}_{IB}, \mathbf{v}_{IB}, \boldsymbol{\omega}_B]^T$  where  $\mathbf{p}_{IB} \in \mathbb{R}^3$  is the position,  $\mathbf{q}_{IB} \in \mathbb{S}\mathbb{O}(3)$  is the unit quaternion that describes the rotation of the platform,  $\mathbf{v}_{IB} \in \mathbb{R}^3$  is the linear velocity vector, and  $\boldsymbol{\omega}_B \in \mathbb{R}^3$  are the bodyrates in the body frame. The input of the system is given as the collective thrust  $\mathbf{f}_B = [0 \ 0 \ f_{Bz}]^T$  and body torques  $\boldsymbol{\tau}_B$ . For readability, we drop the frame indices as they are consistent throughout the description. The dynamic equations are

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v} & \dot{\mathbf{q}} &= \frac{1}{2} \mathbf{q} \odot \begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix} \\ \dot{\mathbf{v}} &= \mathbf{g} + \frac{1}{m} \mathbf{R}(\mathbf{q}) \mathbf{f}_T & \dot{\boldsymbol{\omega}} &= \mathbf{J}^{-1} (\boldsymbol{\tau} - \boldsymbol{\omega} \times \mathbf{J} \boldsymbol{\omega}) \end{aligned} \quad (1)$$

where  $\odot$  represents the Hamilton quaternion multiplication,  $\mathbf{R}(\mathbf{q})$  the quaternion rotation,  $m$  the quadrotor's mass, and  $\mathbf{J}$  the quadrotor's inertia.

Additionally, the input space given by  $\mathbf{f}$  and  $\boldsymbol{\tau}$  is decomposed into single rotor thrusts  $\mathbf{f} = [f_1, f_2, f_3, f_4]$  where  $f_i$  is the thrust at rotor  $i \in \{1, 2, 3, 4\}$ .

$$\mathbf{f}_T = \begin{bmatrix} 0 \\ 0 \\ \sum f_i \end{bmatrix} \quad \text{and} \quad \boldsymbol{\tau} = \begin{bmatrix} l/\sqrt{2}(f_1 + f_2 - f_3 - f_4) \\ l/\sqrt{2}(-f_1 + f_2 + f_3 - f_4) \\ c_\tau(f_1 - f_2 + f_3 - f_4) \end{bmatrix} \quad (2)$$

with the quadrotor's arm length  $l$  and the rotor's torque constant  $c_\tau$ .

Furthermore, in order to approximate the most prominent aerodynamic effects, we extend the quadrotor's dynamics to include a linear drag model [20]. Let  $\mathbf{D}$  be the diagonal matrix with drag coefficients  $d$  such that  $\mathbf{D} = \text{diag}(d_x, d_y, d_z)$ , we expand the dynamic model from (1) by:

$$\dot{\mathbf{v}} = \mathbf{g} + \frac{1}{m} \mathbf{R}(\mathbf{q}) \mathbf{f}_T - \mathbf{R}(\mathbf{q}) \mathbf{D} \mathbf{R}^T(\mathbf{q}) \cdot \mathbf{v} \quad (3)$$

### B. Model Predictive Contouring Control

Model Predictive Contouring Control (MPCC) [6] offers a change of paradigm with respect to classical Model Predictive Control approaches. Instead of minimizing the error in the

input and state space, with respect to a given trajectory, it minimizes the euclidean distance to a given 3D path, while maximizing how fast the path is travelled. One of the benefits of this control approach is that it does not need a feasible trajectory to be generated beforehand. It is sufficient to provide the controller with a continuously differentiable 3D path. It achieves this by solving the difficult time allocation problem and the control problem concurrently, and in real time. The main cost function to be solved by MPCC is the following:

$$\begin{aligned} \pi(\mathbf{x}) = \operatorname{argmin}_u \sum_{k=0}^N & \|e^l(\theta_k)\|_{q_l}^2 + \|e^c(\theta_k)\|_{q_c}^2 \\ & + \|\omega_k\|_{Q_\omega}^2 + \|\Delta v_{\theta_k}\|_{r_{\Delta v}}^2 + \|\Delta \mathbf{f}_k\|_{R_{\Delta \mathbf{f}}}^2 - \mu v_{\theta,k} \\ \text{subject to } \mathbf{x}_0 = \mathbf{x} & \\ \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) & \\ \underline{\omega} \leq \omega \leq \bar{\omega} & \\ \underline{\mathbf{f}} \leq \mathbf{f} \leq \bar{\mathbf{f}} & \\ 0 \leq v_\theta \leq \bar{v}_\theta & \\ \underline{\Delta v_\theta} \leq \Delta v_\theta \leq \bar{\Delta v_\theta} & \\ \underline{\Delta \mathbf{f}} \leq \Delta \mathbf{f} \leq \bar{\Delta \mathbf{f}} & \end{aligned} \quad (4)$$

where  $f(\mathbf{x}_k, \mathbf{u}_k)$  is the dynamic equation, described in Section III-A,  $\theta_k$  is the so-called *progress* term at timestep  $k$ , and  $v_{\theta,k}$ , it's derivative with respect to time.  $e^c(\theta_k)$  is the *contour error*, which is defined as the minimum euclidean distance from the current position to the reference 3D path, and  $e^l(\theta_k)$  is the *lag error*. These errors are penalized by the weights  $q_c$  and  $q_l$ , respectively.  $q_c$  is dynamically changed depending on the position of the waypoints to accurately pass through them. For more details on this control approach, we refer the reader to [6].

### C. Time-Optimal Point-Mass Model Trajectory

As mentioned in the previous section, one of the benefits of MPCC is that its only requirement is a continuously differentiable 3D path that does not need to be dynamically feasible. In [6], the authors leverage this advantage and generate paths using a sampling based approach that considers a point-mass model instead of the full quadrotor dynamics for the path generation. They show that the platform is able to fly in a time-optimal fashion, while significantly reducing the computational burden, since the generation of point-to-point trajectories using a point mass model has a closed-form solution. In what follows, we explain how these time-optimal point-mass model trajectories are generated, and how we have made the method more efficient in order to achieve online replanning capabilities.

Given an initial state consisting of position  $\mathbf{p}_0$  and velocity  $\mathbf{v}_0$ , and given the reduced dynamics of a point-mass model,  $\dot{\mathbf{p}} = \mathbf{u}$ , with the input acceleration being constrained  $\underline{\mathbf{u}} \leq \mathbf{u} \leq \bar{\mathbf{u}}$ , it can be shown by using Pontryagin's maximum principle [21] that the time-optimal control input results in a bang-bang policy of the form:

$$u^*(t) = \begin{cases} \underline{u}, & 0 \leq t \leq t_1^* \\ \bar{u}, & t_1^* \leq t \leq T^* \end{cases} \quad (= t_1^* + t_2^*) \quad (5)$$

or vice versa, starting with  $\bar{u}_x$ . The input policy shown in (5)

results in the following trajectory:

$$\begin{aligned} p_1 &= p_0 + v_0 t_1 + \frac{\underline{u}}{2} t_1^2 \\ v_1 &= v_0 + \underline{u} t_1 \\ p_2 &= p_1 + v_1 t_2 + \frac{\bar{u}}{2} t_2^2 \\ v_2 &= v_1 + \bar{u} t_2 \end{aligned} \quad (6)$$

where it is straightforward to verify that, if we know  $p_0, v_0, p_2, v_2, \underline{u}$  and  $\bar{u}$ , then there are exactly four unknowns:  $t_1, t_2, p_1$  and  $v_1$ , and therefore, we can solve (6) in closed form. The previous (5) and (6), depict the solution for one of the axis. We therefore do this for every axis ( $x, y$  and  $z$ ) separately.

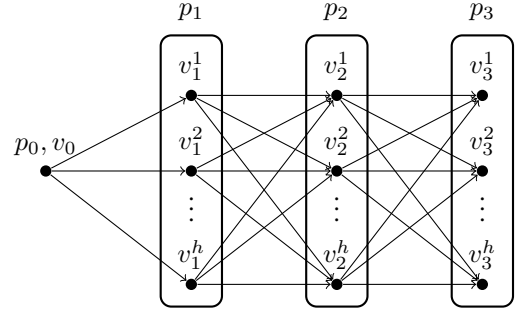


Fig. 2. Velocity search graph where every column  $p_i, i \in [1, 2, 3]$  is a gate in the receding horizon. We find the minimum time, point-mass model trajectory by doing Dijkstra's algorithm on this graph. The different ways of sampling the nodes on the graph characterize the different algorithms described in this paper.

Once we compute the per-axis minimum times ( $T_{axis}$ ), we choose the maximum of these minimum times ( $T^* = \max(T_x^*, T_y^*, T_z^*)$ ) and slow down the other two axes' policies to make all three axes' durations equal. To do this, a new parameter  $\alpha \in [0, 1]$  is introduced that scales the acceleration bounds. For example, if we need to increase the duration of the  $x$  axis, the applied control inputs are scaled to  $\alpha \underline{u}_x$  and  $\alpha \bar{u}_x$ , respectively. In order to find the value of  $\alpha$ , we solve again (6) but now with  $\alpha$  multiplying the respective accelerations, and we augment (6) with

$$t_{1,axis} + t_{2,axis} = T^* \quad (7)$$

and solve again the augmented system.

Assuming that the positions of the waypoints are known, the idea is to recursively sample the velocities at the waypoints and find the combination of velocities that make the final trajectory time optimal for a point-mass model. To solve the problem efficiently, we use a receding horizon approach where only the next  $H_g$  gates are taken into account and then treat this as a shortest path problem and solve it using Dijkstra's algorithm [21].

This is depicted in Fig. 2, where every column represents a different waypoint position and for every column, a set of velocities are sampled. There are different ways in which these velocities can be sampled, and they are discussed in the next section. In our particular application, the gate horizon has been chosen to be  $H_g = 3$  gates. To justify this choice, in [6] the authors have generated PMM trajectories for the task of completing 3 laps on the *CPC-track*, for different values of  $H_g$ . The resulting total lap times are then averaged over 10

runs and compared. In Table II of [6], one can notice how for  $H_g$  larger than 3, the improvement in total time is negligible.

#### D. Velocity search strategies

From Fig. 2 one can see how the amount of nodes inside the velocity sampling space is the main contributing factor to the amount of edges in our graph, and, therefore, to the amount of times that we need to solve (6). More generally, if there are  $h$  velocity samples per gate, and  $H_g$  gates, the total number of edges in our velocity graph  $E$  is:

$$E = h + (h^2 \cdot (H_g - 1)) \quad (8)$$

which means that the upper bound for computational complexity grows quadratically with the number of velocity samples per gate  $h$ . This is an upper bound because we use Dijkstra's algorithm [21], which means that potentially not all edges need to be computed.

Two different velocity search strategies are explored in this work:

##### 1) Random sampling

We randomly sample  $h$  different velocities such that they lie in a cone pointing towards the exit direction of the gate. The sampling distribution is uniform, and  $h = 150$  has shown to be a good value to achieve time optimal performance [6], [3]. This is the approach that is used for the generation of the fixed PMM reference in our experiments.

The worst case number of edges that we would need to compute for this approach is:

$$E_{random} = h + (h^2 \cdot (H_g - 1)) \\ = 150 + (150^2 \cdot 2) = 45150 \quad (9)$$

##### 2) Cone Refocusing via Binary Search

In order to reduce the computational burden of the method described above, instead of randomly sampling over a cone in front of the exit direction of the gate, we do it in a more directed way. First, we bin the cone in front of the gate in the following way: we sample  $s$  different velocity norms,  $s$  different values for yaw and  $s$  different values for pitch. These values are evenly distributed over the entire feasible velocity space. This means that in this case, instead of 150,  $h = s^3$ . In our case, we have chosen  $s = 3$ , which reduces significantly the number of samples per gate to  $h = 27$ . With these samples, we build the nodes of the velocity graph as shown in Fig. 2. Next, we compute the edges of this graph by using (6), (5), and use Dijkstra's algorithm to find the optimal path. Then, we repeat the process again, but this time instead of sampling the velocities on the entire feasible cone, we *refocus* the cone and sample only on a small cone around the previous minimum time trajectory. We continue this iterative process until the ratio between previous minimum time and new minimum time is less than  $\epsilon$ . In our case, this  $\epsilon$  has been chosen such that if the improvement is less than 1% in minimum time between iteration, the algorithm has converged and the iteration ends. A more detailed description of this algorithm is shown in Algorithm 1. A graphical illustration of one cone refocusing recursion step is depicted in Fig. 3. Additionally, a block diagram showing how the proposed

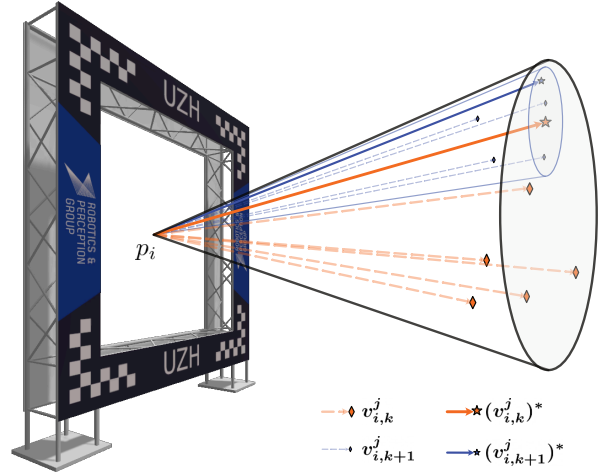


Fig. 3. One iteration of the cone refocusing recursion. The dashed orange velocity samples are equally distributed over the larger cone. The solid orange arrow represents the optimum over the large cone. Our algorithm refocuses a smaller, second cone around this optimum, and repeats the approach in the smaller cone (blue samples).

planning approach is combined with MPCC is depicted in Fig. 4

Finally, the upper bound number of edges that we need to compute is now:

$$E_{refocus} = (h + (h^2 \cdot (H_g - 1))) \cdot K \\ = (27 + (27^2 \cdot 2)) \cdot 4 = 5940 \quad (10)$$

where  $K$  is the number of refocusing iterations. Experimentally, most of the times it takes  $K = 4$  iterations for this algorithm to converge.

#### Algorithm 1: PMM generation via cone refocusing

---

**Input:**  $p[i], i \in [1, 2, \dots, H_g], x_0$  initial state,  $s$  number of bins

---

```

1 for  $k \in [0, 1, \dots]$  do
2    $V_{graph} \leftarrow []$ 
3   for  $i \in [1, 2, \dots, H_g]$  do
4      $p_i \leftarrow p[i], j \leftarrow 0$ 
5      $v_{step} \leftarrow (v_{max} - v_{min})/v_{grid}$ 
6      $\theta_{step} \leftarrow (\theta_{max} - \theta_{min})/\theta_{grid}$ 
7      $\psi_{step} \leftarrow (\psi_{max} - \psi_{min})/\psi_{grid}$ 
8     for  $l \in [0, s - 1]$  do
9        $v_l \leftarrow v_{min} + l \cdot v_{step}$ 
10      for  $m \in [0, s - 1]$  do
11         $\theta_m \leftarrow \theta_{min} + m \cdot \theta_{step}$ 
12        for  $n \in [0, s - 1]$  do
13           $\psi_n \leftarrow \psi_{min} + n \cdot \psi_{step}$ 
14           $V_{graph}[i, j] \leftarrow vec(v_l, \theta_m, \psi_n)$ 
15           $j \leftarrow j + 1$ 
16       $((v_{i,k}^j)^*, T_k^*) \leftarrow (5), (6)$  and Dijkstra's algorithm on  $V_{graph}$ 
17      Update search cone using  $(v_{i,k}^j)^*$ :
18        update  $v_{max}, v_{min}$ 
19        update  $\theta_{max}, \theta_{min}$ 
20        update  $\psi_{max}, \psi_{min}$ 
21      if  $T_k^* > \epsilon \cdot T_{k-1}^*$  then
22        break

```

---

#### E. Comparison

In this section we compare both sampling approaches introduced above. As one can tell by looking at (9) and (10), since the worst case computational complexity of Dijkstra's algorithm evolves quadratically with the number of edges of the graph, we already expect the cone refocusing method to

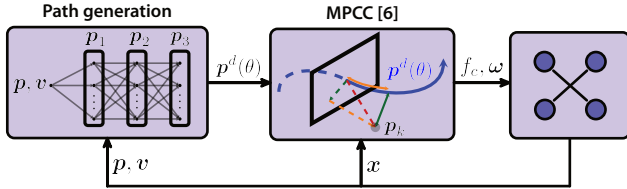


Fig. 4. Block diagram of the proposed online replanning approach. At every timestep we generate a new path that considers the current position and velocity and that is tracked by the MPCC controller architecture.

TABLE I  
COMPUTATION TIMES FOR DIFFERENT SAMPLING STRATEGIES

Approach	Solve time (Desktop)	Solve time (Jetson TX2)
Random	29.67 ms	126.5 ms
<b>Refocus (Ours)</b>	<b>3.48 ms</b>	<b>13.3 ms</b>

be computationally lighter than the random approach. This is indeed the case, as shown in Table I.

However, it still remains the question of which algorithm finds the best solution. In this context, best refers to the one that results in trajectories that achieve a lower  $T^*$ . To this end, in Fig. 5 we plot the optimal times for the running horizon, for both approaches. One can see how the random sampling approach is, at best, equal to the cone refocusing one, and most of the time is worse, since it results in higher minimum times. Additionally, since the random sampling approach cannot run in real-time, it needs to be run every 3-4 control iterations. This is clearly visible in the staircase shape of the blue line in Fig. 5. This experiment has been done in a desktop computer. The reason for a better performance of the proposed sampling-based method lies in the fact that, in some iterations, e.g., when the distance from the current position to the first gate is very small, the random approach becomes very sensitive to the choice of speed and struggles to find the optimal solution. Since we have the MPCC controller in the loop, instabilities in the online replanning lead to different MPCC control policies, which can lead to more instabilities, rendering the resulting path sub-optimal (see Fig. 5).

#### IV. SIMULATION EXPERIMENTS

##### A. Experimental Setup

The identified quadrotor parameters are shown in Table II. For

TABLE II  
QUADROTOR PARAMETERS

Parameter	Nominal Value
$m$	[kg] 0.752
$\mathbf{J}$	[ $\text{gm}^2$ ] $\text{diag}(2.5, 2.1, 4.3)$
$(u_{\min}, u_{\max})$	[N] (0, 8.5)
$(d_x, d_y, d_z)$	[kg/s] (0.26, 0.28, 0.42)

this experiment, we have identified our in-house designed race drone which we then simulated using the state-of-the-art BEM simulator [22]. We choose this simulator because it includes very complex aerodynamic effects caused by how the drone moves through air and by the interaction of the propellers with themselves. This simulator has shown excellent simulation to reality transfer. In fact, so much so, that throughout this paper all tuning parameters have been kept unchanged between simulation experiments and real world experiments. The MPCC controller runs at a frequency of 100 Hz, and the online replanning is done at every time step, when running in a

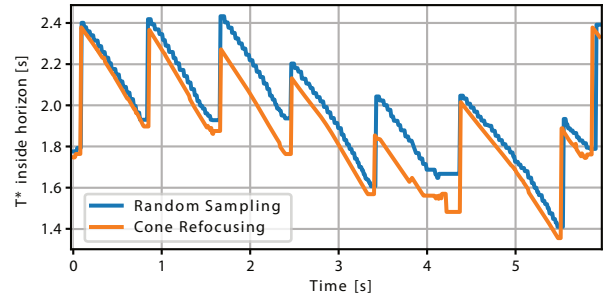


Fig. 5. Comparison of the optimal time achieved by both sampling approaches over a sequence of gates. The random sampling approach (in blue), limited to a number of samples per gate of 150, always achieves worse or equal minimum times than the proposed cone refocusing approach (in orange).

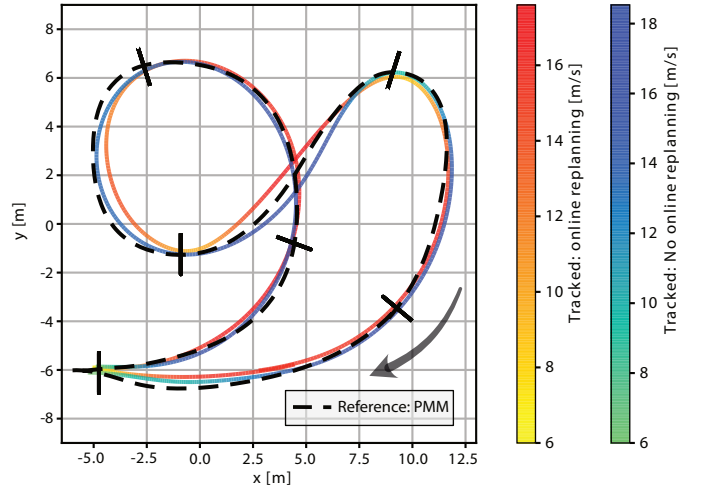


Fig. 6. Comparison in simulation with and without replanning, in the nominal case without disturbances.

desktop computer, or every 2 control iterations, when running in the NVIDIA Jetson TX2<sup>1</sup>.

The simulation experiments presented in this section will be focused on benchmarking our re-planning approach in the *SplitS* race course, the same race track that has been used in [6], [5]. We will focus our comparison in three different cases: nominal case, where we don't include any waypoint change or external disturbance; moving waypoint, task that cannot be tackled by the offline planning approach; and wind disturbance case, where we create a strong wind disturbance right before passing through one of the waypoints and compare how the different planning strategies help coping with it.

##### B. Nominal case

The aim of this experiment is to show that even under no external disturbance and under no waypoint movement, the use of replanning still presents an advantage.

In Fig. 6, we show a superposition of our online replanning approach and the one without online replanning, which is tracking the curve labelled as *Reference: PMM*. It is interesting to notice that even though both approaches present a very similar shape, our online replanning strategy presents slightly smaller maximum velocities, as one can tell by looking at the colorbars. This might seem counterintuitive when looking at the measured lap times, in Table III, since the online replanning

<sup>1</sup><https://developer.nvidia.com/EMBEDDED/jetson-tx2>



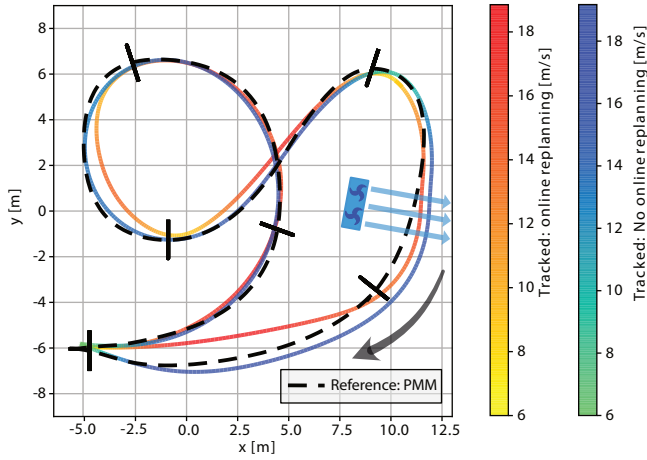


Fig. 7. Comparison in simulation of response to wind disturbance, with and without replanning. The blue arrows represent the direction of the wind disturbance. Our online replanning approach is able to steer the drone back to the gate.

approach produces lower lap times. This is due to the fact that, because we are planning always from the current drone state, the contour error in the first few control horizon steps is very low, enabling for a better maximization of the progress overall. On the contrary, when the reference is fixed, the contour error will generally be non-zero. This can also be seen when measuring the average contour error  $C_{ec} = \frac{\sum_{i=0}^M \|e^c(\theta_i)\|_{q_c}^2}{M}$  and average progress velocity  $C_{v_\theta} = \frac{\sum_{i=0}^M \mu v_{\theta,i}}{M}$  from the tracks in Fig. 6. For fixed planning,  $C_{ec} = 26120$  and  $C_{v_\theta} = 30750$ , whereas for replanning,  $C_{ec} = 0$  and  $C_{v_\theta} = 33932$ . As predicted, in this case the replanning approach results in higher objectives when maximizing the progress.

### C. Response to wind disturbance

In this section we analyze how the online replanning strategy used in this paper helps when reacting to unknown disturbances.

Fig. 7 shows the trajectory tracking performance for the different cases: with and without online replanning. The blue square shows where the wind disturbance has been placed for this experiment, and the arrows show its direction. In the simulator, the wind has been modeled as a constant force of 25 N in front of the area depicted in blue in Fig. 7. We can see how, given this disturbance, the blue path does not pass through the gate located in the affected area. However, because of the replanning, the proposed strategy is able to consider its current position and velocity in the new plan and successfully pass through the gate. In Table III, we show the minimum lap times of our replanning approach and for the fixed reference approach, for both nominal case and with strong wind disturbances. For completeness, a row containing the random sampling approach lap times has been added. It is important to note that for this case, the planning is not run at every timestep, but every 30 ms, as shown in Table I. From these results different details can be highlighted. First, the lap time of the fixed planning approach when presented with wind disturbance is smaller than for our approach. This is because the disturbance pushes the radius of the curve to be larger, and therefore less normal acceleration is needed and higher velocities can be achieved. However, the lap is invalid

because the platform does not pass through the gate.

Second, the lap time of the random sampling approach in the nominal case, is slightly slower than for the fixed planning in the nominal case. This is attributed to, as shown in Fig. 5, the fact that the random sampling generally generates worse optimal times than the cone refocusing approach. Another reason for this is that the MPCC controller tolerates better small changes in the reference trajectory, and this is only the case when we either replan at every time step or if we don't replan at all.

Summarizing, even with wind disturbances affecting very close to the waypoint, our approach is able to promptly correct for the external force and pass through the waypoint accurately, while still retaining time optimality.

TABLE III  
SIMULATED MINIMUM LAP TIMES

Approach	Nominal case	External disturbance case
Fixed planning [6]	5.67 s	5.51 s (Invalid)
Random Sampling	5.74 s	5.71 s
<b>Cone refocusing (Ours)</b>	<b>5.57 s</b>	<b>5.63 s</b>

## V. REAL WORLD EXPERIMENTS

To further display the capabilities of the proposed online replanning method, in this section, we deploy the algorithm on the physical platform. This platform has been built in-house from off-the-shelf components and it carries an NVIDIA Jetson TX2 board. A more detailed description of our platform can be found in [23]. Our method can run in both, an offboard desktop computer equipped with an Intel(R) Core(TM) i7-8550 CPU @ 1.80GHz, and on the Jetson TX2. A Radix FC board that contains the Betaflight<sup>2</sup> firmware is used as a low level controller. This low-level controller takes as inputs body rates and collective thrusts. In this case, the body rates are commanded directly from the MPCC controller, and the collective thrust is computed as the sum of the single rotor thrusts. It is, however, still advantageous that the MPCC considers the full state dynamics. This way, the body rates, and collective thrust are generated taking into account possible current and predicted saturations at the single rotor thrust level.

For state estimation, we use a VICON<sup>3</sup> system with 36 cameras in one of the world's largest drone flying arenas (30 × 30 × 8 m) that provide the platform with down to millimeter accuracy measurements of position and orientation.

### A. Implementation Details

In order to deploy our MPCC controller, (4) needs to be solved in real-time. To this end, we have implemented our optimization problem using acados [24] as a code generation tool, in contrast to [6], where its previous version, ACADO [25] was used. It is important to note that for consistency, the optimization problem that is solved online is written in (4) and is exactly the same as in [6]. The main benefit of using acados is that it provides an interface to HPIPM (High Performance Interior Point Method) solver [26]. HPIPM solves optimization problems using BLASFEO [27], a linear algebra library specifically designed for embedded optimization. The usage of

<sup>2</sup><https://betaflight.com/>

<sup>3</sup><https://www.vicon.com/>

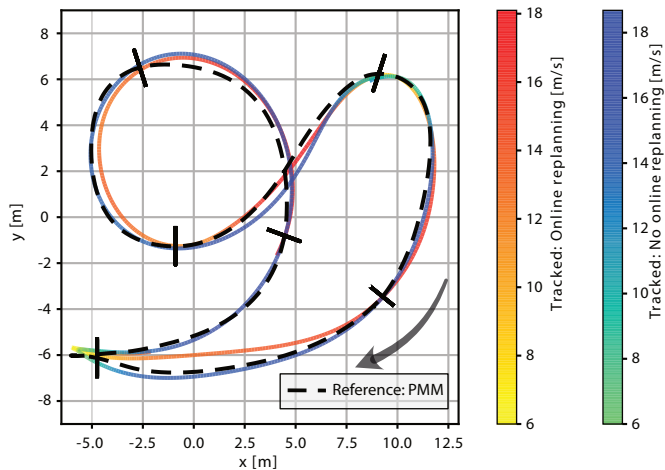


Fig. 8. Comparison in real flight with and without replanning in the nominal case without disturbances.

this new tools and new solver shows a large decrease in solve times that allows the deployment of our MPCC algorithm in an embedded platform, in real time. The solve times and their comparison with the previous implementation are shown in Table IV. Our MPCC is run at a 100 Hz frequency, and the prediction steps are of 60 ms, with a prediction horizon length of 20 steps.

TABLE IV  
SOLVE TIMES MPCC WITH DIFFERENT IMPLEMENTATIONS

Implementation	Desktop	NVIDIA Jetson TX2
ACADO+QPOASES <sup>4</sup>	4.8 ms	Doesn't converge
<b>acados+HPIPM (Ours)</b>	<b>2.21 ms</b>	<b>6.6 ms</b>

Thanks to this re-implementation, we can run the MPCC together with the online replanning proposed in this paper in the Jetson TX2. To achieve this, the planning is running in a different thread, and triggered once every 2 control iterations.

### B. Nominal case

In this section we compare the tracking performance in the *SplitS* race track, in real flight. The tracking performance is shown in Fig. 9. Both approaches achieve similar speed of around 18 m/s, and the behaviour is very similar to the one shown for the same case in the simulation experiments.

In Table V we see how the lap times for the proposed approach are also lower than for the fixed reference case. This follows the same explanation we mentioned in the simulation experiments section: the MPCC controller has less contour error because of the replanning, and this allows for a better maximization of the progress term.

### C. Moving waypoint

With this experiment we aim to showcase the ability of our approach to steer the platform through the waypoints even in the case when they move. Fixed time-optimal reference generating methods cannot adapt to changing conditions, since they cannot run in real-time. Our method can run at every control iteration, therefore, assuming perfect knowledge of the waypoints, we can directly plan a new trajectory that passes through them. In Fig. 9 we show how the drone perfectly flies through all the waypoints, even when the gate that is in the center of the frame is constantly moving.

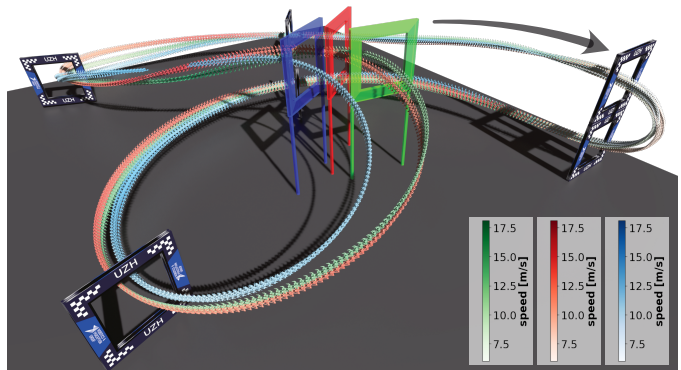


Fig. 9. Real world experiment when waypoints are moving. The red, blue and green gates represent the same gate at different times. The different colored trajectories represent different real-time replanning. Our approach is able to adapt to the movement of the gate while flying as fast as possible.

### D. Response to wind disturbance

In this section we analyze how the online replanning strategy used in this paper helps when reacting to unknown disturbances in the real world. To this end, with the purpose of generating a very strong stream of wind we adapt our experimental setup.

We have placed 4 propellers fixed to a point that is located at 0.5 m to where our platform flies. These propellers are constantly at full thrust, in order to generate a strong, constant stream of wind with velocities of up to 68 km/h.

In Fig. 10, we show the tracking performance of the proposed approach compared to the tracking performance of the fixed reference approach. On the zoomed plot, on the right of Fig. 10, one can see how the proposed approach flies exactly through the center of the gate, while the fixed reference approach gets deviated from the reference path by 0.24 m.

Similarly to the simulation results, the lap times for the wind disturbance case, shown in Table V, are only slightly worse than for the non wind disturbance case. Additionally, for the fixed planning approach, in the wind disturbance case the lap times are smaller than for the nominal case, due to, as explained in IV-C, the slightly larger curvature radius.

TABLE V  
REAL FLIGHT MINIMUM LAP TIMES

Approach	Nominal case	External disturbance case
Fixed planning	5.93 s	5.86 s
<b>Online replanning (Ours)</b>	<b>5.73 s</b>	<b>5.83 s</b>

## VI. CONCLUSION

This paper introduces a novel way of coping with dynamic environments online by replanning at every timestep and flying minimum time policies. We showed that our approach can adapt online to changing position of a gate during drone racing with speeds of more than 60 km/h. Moreover, we showed that the online replanning strategy can cope with strong wind disturbances. However, the state estimate has completely been provided by a motion capture system. Therefore, a very interesting future direction would be the adaptation of the current approach such that it relies only on vision sensors. For that, perception-awareness of the control approach would be a first necessary step to take, such that the image sensor would always be pointing towards the next waypoint without

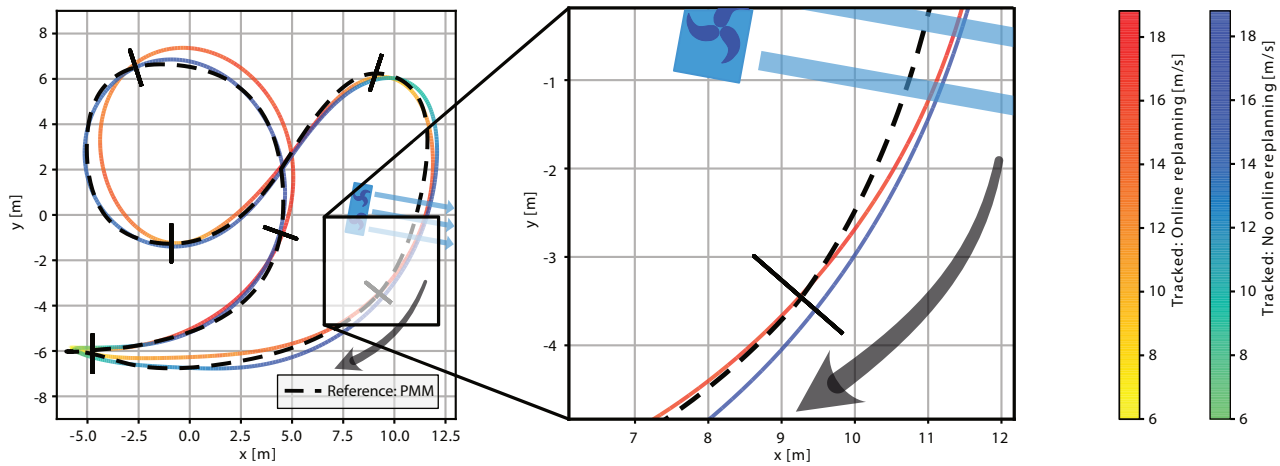


Fig. 10. Comparison in the real world of response to wind disturbance of 18.8 m/s, with and without replanning. Our online replanning approach (red line) is able to quickly correct for the disturbance and steers the drone back to the center of the gate.

sacrificing flying performance. Secondly, vision-based state estimation of both, the platform state and the position of the gates.

Another connection to our work goes in the direction of sampling-based MPC (SBMPC) [28], [29], where also a receding horizon problem is addressed by sampling, but in this case to directly generate feasible control inputs instead a nominal path to track. We believe it might be an interesting future direction to apply sampling techniques from SBMPC to our planning approach.

## VII. ACKNOWLEDGMENTS

The authors want to thank Thomas Laengle and Cafer Mertcan Akcay for their support with the experimental setup.

## REFERENCES

- [1] E. Ackermann, "Ai-powered drone learns extreme acrobatics," Oct. 2020, accessed:Feb.24,2022. [Online]. Available: <https://spectrum.ieee.org/ai-powered-drone-extreme-acrobatics>
- [2] C. Pfeiffer and D. Scaramuzza, "Human-piloted drone racing: Visual processing and control," *IEEE RA-L*, 2021.
- [3] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, "Alphapilot: Autonomous drone racing," in *RSS*, Jul 2020.
- [4] Z. Han, Z. Wang, N. Pan, Y. Lin, C. Xu, and F. Gao, "Fast-racing: An open-source strong baseline for SE(3) planning in autonomous drone racing," *IEEE RA-L*, vol. 6, no. 4, pp. 8631–8638, 2021.
- [5] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, 2021.
- [6] A. Romero, S. Sun, P. Foehn, and D. Scaramuzza, "Model predictive contouring control for time-optimal quadrotor flight," *IEEE Transactions on Robotics*, pp. 1–17, 2022.
- [7] E. Kaufmann, M. Gehrig, P. Foehn, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Beauty and the beast: Optimal methods meet learning for drone racing," *ICRA*, 2019.
- [8] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: From simulation to reality with domain randomization," *IEEE TRO*, vol. 36, no. 1, p. 1–14, Feb 2020.
- [9] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "Pampc: Perception-aware model predictive control for quadrotors," *IROS*, 2018.
- [10] G. Ryou, E. Tal, and S. Karaman, "Multi-fidelity black-box optimization for time-optimal quadrotor maneuvers," *IJRR*, Jul 2021.
- [11] E. Mueller, A. Censi, and E. Frazzoli, "Low-latency heading feedback control with neuromorphic vision sensors using efficient approximated incremental inference," in *IEEE Conf. Decision Control (CDC)*, 2015.
- [12] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *ISRR*, 2013, pp. 1–8.
- [13] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *IJRR*, vol. 31, no. 5, pp. 664–674, 2012.
- [14] B. Zhou, J. Pan, F. Gao, and S. Shen, "Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight," *IEEE TRO*, 2021.
- [15] B. Zhou, F. Gao, J. Pan, and S. Shen, "Robust real-time uav replanning using guided gradient-based optimization and topological paths," *ICRA*, 2020.
- [16] B. Zhou, F. Gao, L. Wang, C. Liu, and S. Shen, "Robust and efficient quadrotor trajectory generation for fast autonomous flight," *IEEE RA-L*, 2019.
- [17] W. V. Loock, G. Pipeleers, and J. Swevers, "Time-optimal quadrotor flight," in *ECC*, 2013.
- [18] M. Hehn, R. Ritz, and R. D'Andrea, "Performance benchmarking of quadrotor systems using time-optimal control," *AURO*, vol. 33, no. 1-2, pp. 69–88, Mar. 2012.
- [19] S. Spedicato and G. Notarstefano, "Minimum-time trajectory generation for quadrotors in constrained environments," *IEEE Transactions on Control Systems Technology*, 2018.
- [20] M. Faessler, A. Franchi, and D. Scaramuzza, "Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories," *IEEE RA-L*, vol. 3, no. 2, pp. 620–626, Apr. 2018.
- [21] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Athena Scientific, 2000.
- [22] L. Bauersfeld, E. Kaufmann, P. Foehn, S. Sun, and D. Scaramuzza, "Neurobem: Hybrid aerodynamic quadrotor model," *RSS*, 2021.
- [23] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, and D. Scaramuzza, "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *AAAS Science Robotics*, 2022.
- [24] R. Verschueren, G. Frison, D. Kouzoupis, J. Frey, N. v. Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados—a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, Oct 2021.
- [25] B. Houska, H. Ferreau, and M. Diehl, "ACADO Toolkit – An Open Source Framework for Automatic Control and Dynamic Optimization," *Optimal Control Applications and Methods*, vol. 32, no. 3, pp. 298–312, 2011.
- [26] G. Frison and M. Diehl, "Hpipm: a high-performance quadratic programming framework for model predictive control," *IFAC-PapersOnLine*, vol. 53, no. 2, p. 6563–6569, 2020.
- [27] G. Frison, D. Kouzoupis, T. Sartor, A. Zanelli, and M. Diehl, "Blasfeo: Basic linear algebra subroutines for embedded optimization," *ACM Transactions on Mathematical Software*, vol. 44, no. 4, pp. 42:1–42:30, Jul 2018.
- [28] D. D. Dunlap, E. G. Collins Jr, and C. V. Caldwell, "Sampling based model predictive control with application to autonomous vehicle guidance," in *Florida Conference on Recent Advances in Robotics*, 2008.
- [29] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. Rehg, and E. Theodorou, "Robust sampling based model predictive control with sparse objective information," in *Robotics: Science and Systems XIV*. Robotics: Science and Systems Foundation, Jun 2018. [Online]. Available: <http://www.roboticsproceedings.org/rss14/p42.pdf>