

## Efficient Bang-Bang Model Predictive Control for Quadcopters

Jelle Westenberger\*, Christophe De Wagter†, Guido C. H. E. de Croon‡

*Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, the Netherlands*

Time-optimal model predictive control is important for achieving fast racing drones but is computationally intensive and thereby rarely used onboard small quadcopters with limited computational resources. In this work, we simplify the optimal control problem (OCP) of the position loop for several maneuvers by exploiting the fact that the solution resembles a so-called ‘bang-bang’ in the critical direction, where only the switching time needs to be found. The noncritical direction uses a ‘minimum effort’ approach. The control parameters are obtained through bisection search schemes on an analytical path prediction model. The approach is compared with a classical PID controller and theoretical time-optimal trajectories in simulations. We explain the effects of the OCP simplifications and introduce a method of mitigating one of these effects. Finally, we have implemented the ‘bang-bang’ controller as a model predictive controller (MPC) onboard a Parrot Bebop and performed indoor flights to compare the controller’s performance to a PID controller. We show that the light novel controller outperforms the PID controller in waypoint-to-waypoint flight while requiring only minimal knowledge of the quadcopter’s dynamics.

**Keywords:** UAV; quadrotors; MPC; optimal control.

US

### 1. Introduction

Unmanned air vehicles (UAVs) are used in an increasing variety of applications [1]. Several applications, such as emergency response or race tasks require the drones to fly as fast as they can. Autonomous drone racing has recently emerged as a discipline to boost the development of fast-flying robots [2–4].

Traditionally, the problem of time-optimal control generation is solved offboard as available hardware lacks the computational performance to quickly solve the Optimal Control Problem (OCP) onboard a quadcopter [5]. Fast flight is achieved by tracking these trajectories with high-performance controllers [6].

Recent work demonstrated efficient trajectory optimization for snap and leveraging differential flatness to derive the corresponding control inputs [7,8]. However, the snap optimization method does not optimize for ‘minimum time’. The total flight time must be predefined and the dynamical

limits of the quadcopter are not taken into account. Including time and dynamic feasibility constraints in the optimization process increases the computational complexity of the problem [9]. On the other hand, Falanga *et al.* [10] defined a sequential quadratic programming problem to simultaneously optimize control inputs for action and perception objectives. Albeit that in this work, the reference trajectories are pre-computed. Kaufmann *et al.* [11] extended this work and demonstrated a pipeline that is fully embedded and is efficient enough to be implemented as a robust MPC. Wang *et al.* [12] even presented MPC-based trajectory generation for multiple quadrotors flying together in cluttered environments.

While the former results are great, this comes at a very high computational cost. To address this, optimal control has also been approximated with deep neural nets, which are lighter than the original optimization [13–15]. This approach is powerful but very data intensive.

Model predictive control remains very computationally expensive and few onboard implementations exist for very light drones [11,16]. In this category, classical control remains common [17].

For a lot of trajectories, the time-optimal solution simplifies to a well-timed maximal control deflection.

Received 17 January 2022; Accepted 11 May 2022; Published 6 July 2022.  
This paper was recommended for publication in its revised form by editorial board member, Jose Martinez-Carranza.  
Email Addresses: \*jellewestenberger@gmail.com, †c.dewagter@tudelft.nl, ‡g.c.h.e.decroon@tudelft.nl



Fig. 1. A comparison of a circular flight path between the proposed controller (green) and a classical PID controller (red).

This paper, therefore, proposes a light strategy to approximate time-optimal control by computing this timing onboard (see Fig. 1).

Section 2 shows that the time-optimal position control simplifies to a ‘bang-bang’ action on the attitude under well-selected conditions. In Sec. 3, we derive the differential equations that drive the proposed light MPC controller. Simulation results are presented in Sec. 4. Section 5 augments the model for the latency in attitude. Section 6 shows the results obtained onboard a Parrot Bebop before Sec. 7 gives the conclusions.

## 2. Simplified Time-Optimal Control

We have simplified the OCP by assuming a constant altitude and using the fact that for second-order systems the time-optimal solution consists of a ‘bang-bang’ motion. For quadcopter position control, this translates to a double step in either pitch or roll with maximal amplitude. The OCP is hereby reduced to a problem in which the only parameter to be optimized is the switching time, which reduces the computational complexity of the problem sufficiently to even allow implementation onboard very small quadcopters. The collective thrust is governed by the constant altitude assumption and is therefore considered to be always equal to  $W / \cos \theta \cos \phi$ , where  $W$  is the weight,  $\theta$ , and  $\phi$  are the pitch and roll angles, respectively. The lateral dynamics can then be further simplified.

### 2.1. Proving ‘bang-bang’ solution for constant angles

Hehn *et al.* [5] showed in their work that with Pontryagin’s minimum principle it can be proven that the time-optimal solution for a two-dimensional quadcopter trajectory consists of a ‘bang-bang’ input in thrust and bang-singular-bang in rotational rate. We show that this solution remains valid in our simplified OCP in which we neglect the rotational dynamics, to further reduce the computational expense. Continuing with the Hamiltonian from [5]:

$$H(\mathbf{x}, \mathbf{u}, \mathbf{p}) = 1 + p_1 \dot{x} + p_2 u_T \sin \theta + p_3 \dot{z} + p_4 (u_T \cos \theta - 1) + p_5 u_R, \quad (1)$$

where  $u_T$  is the thrust input and  $u_R$  is the rotation rate input.  $p_i$  are the co-states. Our model assumes instantaneous attitude changes and no changes in altitude. Therefore, we can discard the last three terms of Eq. (1) and change the state  $\theta$  to input  $u_\theta$ . Pontryagin’s minimum principle states that the optimal control input  $\mathbf{u}^*$  minimizes the Hamiltonian [18].

$$\mathbf{u}^* = \operatorname{argmin} p_2 u_T \sin u_\theta. \quad (2)$$

Depending on the sign of  $p_2$ ,  $u_\theta^*$  is either  $\pm 0.5\pi$  or singular when  $p_2 = 0$ . However, at these pitch angles, it would be impossible to maintain altitude. Therefore, maximum pitch and roll angles are determined based on the thrust-to-weight performance of the quadcopter while reserving a margin of available thrust for additional control.

### 2.2. Minimum-effort approach

We apply the ‘bang-bang’ solution to one critical axis. Intuitively, this axis is selected to be the direction with the largest initial position error. Control of the direction perpendicular to this axis is based on a ‘minimum effort’ approach. The intuition behind this approach is to only spend the minimum required thrust on decreasing the position error in the noncritical dimension such that a maximum available thrust can be spent on the critical dimension. This is achieved by calculating the constant attitude for which the noncritical position target is reached at the same time the critical target is reached.

## 3. Bang-Bang MPC

Based on the simplified OCP, we have created a controller that calculates the optimal roll and pitch angle from path predictions. We refer to this pipeline as the ‘bang-bang’ controller.

### 3.1. Path prediction

For the sake of computational efficiency, we have simplified the dynamics such that the quadcopter’s position and velocity can be evaluated analytically. By discarding the rotational and vertical dynamics, and partially decoupling the longitudinal and lateral dynamics, we have derived a set second-order differential equations to describe the quadcopter’s position and velocity.

Based on the aforementioned assumptions and the force diagram depicted in Fig. 2, we state that the pitch angle  $\theta$  and the thrust are constant and  $T_z$  equals the weight  $W$ . Furthermore, we assume that drag force  $D$ , consists only of flapping drag and is linearly proportional to airspeed  $\dot{x}$ , which is governed by drag coefficient  $C_d$ . So, we can write as follows:

$$\ddot{x} = g \tan \theta - \frac{C_d}{m} \dot{x}, \quad (3)$$




Fig. 2. 2D quadcopter dynamics.

where  $g$  and  $m$  are the gravitational acceleration and mass, respectively.

Equation (3) is a second-order, nonhomogeneous equation and is easily solved with the characteristic equation and method of undetermined coefficients. This yields

$$x = c_1 e^{-\frac{C_d}{m}t} + c_2 + \frac{W \tan \theta}{C_d} t, \quad (4a)$$

$$\dot{x} = c_1 \frac{-C_d}{m} e^{-\frac{C_d}{m}t} + \frac{W \tan \theta}{C_d}. \quad (4b)$$

Constants  $c_1$  and  $c_2$  are solved with the quadcopter's initial position  $x_0$  and initial velocity  $\dot{x}_0$ . This procedure can be repeated for the lateral direction, which is the direction out-of-plane in Fig. 2, taking into account the proper Euler angle rotations when deriving the lateral component of the thrust force:

$$y = c_3 e^{-\frac{C_d}{m}t} + c_4 + \frac{W}{\cos \theta} \frac{\tan \phi}{C_d} t, \quad (5a)$$

$$\dot{y} = c_3 \frac{-C_d}{m} e^{-\frac{C_d}{m}t} + \frac{W}{\cos \theta} \frac{\tan \phi}{C_d}, \quad (5b)$$

where  $c_3$  and  $c_4$  are solved with the quadcopter's initial lateral position  $y_0$  and lateral velocity  $\dot{y}_0$ .

Equations (4) and (5) can now be used to describe the quadcopter's position in time. Figure 3 illustrates an example of a single path prediction and the corresponding pitch and roll angles. The 'bang-bang' maneuver of the longitudinal path can be described by evaluating Eq. (4) for two segments; One segment up to the switching instant (the acceleration phase) and a segment up to the time of arrival (the braking phase). Only the constants and pitch angles change between the two sets. The final velocity and position of the first segments are used as initial conditions for the second segment.

The lateral path can be described by one segment because in the 'minimum effort' approach the roll angle is assumed to be constant for the entire maneuver up to the target.

### 3.2. Optimizing

In the OCP, the switching time and the noncritical angle are the two parameters to be optimized. Thanks to the




Fig. 3. Example of predicted longitudinal and lateral paths. The longitudinal direction takes on a 'bang-bang' motion that consists of two segments with opposing constant pitch angles. The lateral direction takes a minimum-effort approach which consists of a single segment and a constant roll angle.

analytical nature of the path equation a fast iterative bisection scheme can be used to find the optimal switching time and angle.

#### 3.2.1. Solving switching time

To solve the switching time, a desired velocity at the position target must be given in advance. The bisection scheme then iteratively adapts the switching time to minimize the velocity error at the target position. This procedure is described in Algorithm 1 and illustrated in Fig 4. In addition to an optimized switching time, an estimated time of arrival (ETA) is given as well. This is used in optimizing the noncritical angle.

---

#### Algorithm 1.

---

```

 $t_0 \leftarrow 0$ 
 $t_1 \leftarrow$  initial guess
 $E_t \leftarrow$  error threshold
 $y_d \leftarrow$  desired position
while  $E > E_T$  do
   $t_s \leftarrow \frac{t_0 + t_1}{2}$ 
   $t_t \leftarrow \text{get\_time\_from\_desired\_speed}(v_d)$ 
   $E \leftarrow \text{get\_position}(t_t) - y_d$ 
  if  $E > 0$  then
     $t_1 \leftarrow t_s$ 
  else
     $t_0 \leftarrow t_s$ 
  end if
end while

```

---




Fig. 4. Illustration of the switching time,  $t_s$ , optimization process. The goal is to reach the target at 30 m with zero rest speed.




Fig. 5. Trajectory comparison in simulation.

### 3.2.2. Solving minimum-effort angle

Analogously to the critical direction, a bisection scheme is used to iteratively change the angle to minimize the non-critical position error at the ETA. The goal for the quadcopter is to reach the critical and noncritical targets simultaneously.

Moreover, during the braking phase of the ‘bang-bang’ motion, the related angle will also be optimized in this fashion to correct for prediction inaccuracies in this phase.

## 4. Simulations

Simulations have been performed to compare the ‘bang-bang’ controller flight performance to a classical PID controller and to time and snap optimized trajectories, provided by the well-known ICLOCS toolbox [19]. This toolbox uses direct collocation to optimize a nonlinear OCP from an initial guess. Two maneuvers, a straight and a

Table 1. Simulated flight times.

|              | Forward | Corners |
|--------------|---------|---------|
| Bang-Bang    | 1.52 s  | 2.53 s  |
| PID          | 2.51 s  | 4.16 s  |
| Minimum snap | 1.77 s  | 2.53 s  |
| Minimum time | 1.30 s  | 1.87 s  |

cornered trajectory have been simulated to individually test the longitudinal and lateral flight behavior. The flight trajectories are shown in Fig. 5 and the resulting flight times are summarized in Table 1.

It can be seen that the ‘bang-bang’ controller outperforms the PID controller in all maneuvers and is on par with the snap-optimized solution, but at a fraction of the computational cost.

## 5. Transition Compensation

It was found in simulations that the instantaneous angle assumption of the path predictor has the largest negative effect on the performance of the ‘bang-bang’ controller. Since a quadcopter cannot achieve infinitely high rotation rates the second part (further called the braking phase) of the ‘bang-bang’ maneuver will always be initiated too late. As illustrated in Fig. 6, path predictions deviate during the rotation from rest to acceleration at 0 s, and during the transition from accelerating to braking around 1.6 s.

To mitigate this issue, we have implemented a method that approximates the elapsed time and change of speed and position during the transition. Subsequently, the initial conditions of the braking phase are augmented with these values to improve the path predictions, as illustrated in Fig. 7.

Figure 8 shows the effect of different degrees of compensation for a simulated flight. The amount of compensation in this simple simulation is determined by manually setting the expected transition duration, thereby blindly forcing a braking motion during this time period. It can be seen that increasing the amount of compensation reduces




Fig. 6. Development of the ETA of a simulated flight corrected for passed simulation time. For perfect predictions, the time of arrival would be constant.




Fig. 7. The initial conditions of the braking phase are shifted with  $\Delta t$ ,  $\Delta y$  and  $\Delta v$  to compensate for the rotation dynamics during the transition from accelerating to braking.

the velocity error at the target position. Also, it visualizes the situation of overcompensation at which point the quadcopter comes to a stop before the target is reached.

## 6. Experiments

Flight tests have been performed with a commercial Parrot Bebop quadcopter in which all software was changed. The performance of the ‘bang-bang’ MPC is compared to a traditional PID controller for different types of maneuvers.

### 6.1. Experimental setup

The ‘bang-bang’ controller has been implemented in the open-source autopilot framework Paparazzi-UAV [20] and is executed onboard a Parrot Bebop quadcopter. The flights were performed in TU Delft’s ‘CyberZoo’ indoor flight area

outfitted with an Optitrack position and attitude tracking system. The position and heading are sent to the drone via WiFi and the state estimation is executed on board, where it is merged with inertial measurements by a complementary filter. The ‘bang-bang’ MPC and PID controller give roll and pitch commands while the inner control loop, based on Incremental Nonlinear Dynamic Inversion (INDI) [21], controls the rotational rates. Figure 9 gives an overview of the control pipeline.

#### 6.1.1. PID controller

The PID controller is a high-gain cascaded position-velocity controller. That is, the position error will govern the desired speed, which in turn governs the pitch and roll commands. A single set of gain values has been selected that gives the best overall result in all tests. This set is kept constant throughout all flight experiments. Furthermore, we have defined saturation limits for the allowable speed and pitch and roll angles. For a fair comparison, the same limits have also been applied to the ‘bang-bang’ controller.

### 6.2. Transition estimators

As discussed in Sec. 5, we can compensate for the unmodeled transition dynamics by approximating the transition losses. These dynamics are difficult and costly to simulate for a real quadcopter; therefore we have derived a simple linear regression model to approximate transition losses  $\Delta t$ ,  $\Delta y$  and  $\Delta v$  from flight measurements. We assumed that these losses are a function of the speed at the switching time,  $v_i$ , and of the total angle,  $\Delta \Theta$  (roll or pitch), the




Fig. 8. Different degrees of transition compensation are illustrated here. The shaded area indicates the period at which braking is forced. The desired target speed  $v_t$  is set to 0 m/s.




Fig. 9. Control pipeline.

quadcopter needs to rotate. So that the functions have the form as shown in the following equation:

$$\Delta t = \xi_0 + \xi_{\Delta\Theta} \cdot \Delta\Theta + \xi_{v_i} \cdot v_i. \quad (6)$$

We found that the transition losses varied between forward, backward and sideways flight maneuvers. Therefore, three different sets of least-squares estimators have been derived, each corresponding with one of these directions. In the control pipeline, one set of estimators is selected based on the direction in which a ‘bang-bang’ maneuver is planned.

### 6.3. Motion primitives flights

Test flights have been performed to test the ‘bang-bang’ controller for different motion primitives. That is, four different maneuvers have been established to test the longitudinal and lateral performance in which the quadcopter starts and ends at rest. For each maneuver, a comparison is made between the ‘bang-bang’ controller, the ‘bang-bang’ controller with transition compensation, and the high-gain PID controller.

During these maneuvers, the heading is kept constant and the critical direction for which the ‘bang-bang’ maneuver is planned is based on the largest component of the initial position error.

Table 2 lists all maneuvers and their initial and target positions. The controllers are assessed on the time it takes to reach their target, the degree of overshoot, and the

velocity error while passing the target. Also, it should be noted that the ‘bang-bang’ controllers are configured to switch to a low-gain PID controller once the target position and desired speed are crossed for the first time to avoid endlessly oscillating when close to the target position.

#### 6.3.1. Results

Appendix A visualizes the results of the motion primitive flights. Figures (A.1)–(A.4) show a comparison of the three different controllers in forward, backward, sideways, and forward + sideways step response maneuvers, respectively.

For each maneuver, similar relative behavior is observed between the controllers. Due to the high gains of the PID controller, the acceleration phase consists of immediately saturated control inputs for all controllers, and therefore the paths are practically identical in this phase. Differences begin to distinguish themselves at the moment of braking. While the braking angle commands of the ‘bang-bang’ controllers are steps, the PID controller shows a steep but slanted braking command, proportional to the velocity error. The instant at which braking is initiated differs between the different controllers. As expected, the ‘bang-bang’ controller with transition compensation included brakes earlier than the ‘bang-bang’ controller without this compensation. The PID controller starts to brake even earlier. This leads to a slightly overdamped response for the PID controller, not overshooting the target.

Finally, notice the different paths in Fig. A.4. Here, it is visible that the ‘bang-bang’ controller prioritizes the longitudinal motion over the lateral motion, minimizing the roll angle which benefits stability and control. The PID controller on the other hand saturates both roll and pitch commands at the start, even though the lateral position error is smaller than the longitudinal.

Table 3 summarizes the flight results of the step responses. It can be seen that the noncompensated ‘bang-bang’ controller has the largest velocity errors and overshoot. Furthermore, it becomes obvious that the compensation

Table 2. Translated distances for each motion primitive maneuver.

| Maneuver         | Initial → Target position (x, y, z) [m] |
|------------------|-----------------------------------------|
| Forward          | (−2, 0, 1.5) → (2.5, 0, 1.5)            |
| Backward         | (−2, 0, 1.5) → (2.5, 0, 1.5)            |
| Sideways         | (0, −2, 1.5) → (0, 2, 1.5)              |
| Forward-sideways | (−2, −2, 1.5) → (2, 1, 1.5)             |
| Forward-up       | (−2, 0, 1) → (2.5, 0, 2.75)             |
| Forward-down     | (−2, 0, 2.75) → (2.5, 0, 1)             |

Table 3. Performance values of the different controllers in four maneuvers.  $n$  is the number of runs performed.

| Controller                            | Maneuver                 |                         |                         |                         |
|---------------------------------------|--------------------------|-------------------------|-------------------------|-------------------------|
|                                       | Forward                  | Backward                | Sideways                | Forward-Sideways        |
| Mean time of arrival [s]              |                          |                         |                         |                         |
| Bang-Bang                             | <b>1.38</b> ( $n = 4$ )  | <b>1.41</b> ( $n = 4$ ) | <b>1.29</b> ( $n = 5$ ) | <b>1.47</b> ( $n = 4$ ) |
| Bang-Bang Comp.                       | 1.42 ( $n = 15$ )        | 1.47 ( $n = 12$ )       | 1.37 ( $n = 11$ )       | 1.53 ( $n = 15$ )       |
| PID                                   | 1.48 ( $n = 10$ )        | 1.54 ( $n = 8$ )        | 1.43 ( $n = 8$ )        | 1.51 ( $n = 8$ )        |
| Mean overshoot [m]                    |                          |                         |                         |                         |
| Bang-Bang                             | 0.62 ( $n = 4$ )         | 0.81 ( $n = 4$ )        | 0.53 ( $n = 5$ )        | 0.27 ( $n = 4$ )        |
| Bang-Bang Comp.                       | 0.18 ( $n = 15$ )        | 0.22 ( $n = 12$ )       | 0.06 ( $n = 11$ )       | 0.11 ( $n = 15$ )       |
| PID                                   | <b>0.05</b> ( $n = 10$ ) | <b>0.04</b> ( $n = 8$ ) | <b>0.04</b> ( $n = 8$ ) | <b>0.04</b> ( $n = 8$ ) |
| Mean velocity error [ $\frac{m}{s}$ ] |                          |                         |                         |                         |
| Bang-Bang                             | 3.05 ( $n = 4$ )         | 3.51 ( $n = 4$ )        | 3.06 ( $n = 5$ )        | 1.85 ( $n = 4$ )        |
| Bang-Bang Comp.                       | 1.60 ( $n = 15$ )        | 1.88 ( $n = 12$ )       | 0.69 ( $n = 11$ )       | 0.38 ( $n = 15$ )       |
| PID                                   | <b>0.08</b> ( $n = 10$ ) | <b>0.08</b> ( $n = 8$ ) | <b>0.14</b> ( $n = 8$ ) | <b>0.06</b> ( $n = 8$ ) |

system has a positive effect on the path prediction performance. Unfortunately, the transition loss model is not accurate enough to completely mitigate the transition losses and some degree of overshoot still occurs. In these simple start-stop tests, the PID controller is marginally slower than both ‘bang-bang’ controllers but has lower overshoot and velocity errors.

#### 6.4. Consecutive waypoints flight

To test the proposed controller in a setting that more closely resembles an autonomous drone race, a flight plan with consecutive positional waypoints has been implemented. In this flight plan, the quadcopter is no longer instructed to come to a full stop at each waypoint. The desired speed at each waypoint has been set to 2 m/s as it was found iteratively that this value in combination with a position threshold of 70 cm resulted in smooth trajectories for both the PID and ‘bang-bang’ controllers (see Fig. 1). However, it is expected that the optimal threshold values are controller- and trajectory-dependent.

Because currently no heading changes were incorporated into the ‘bang-bang’ maneuver planning, the heading is kept constant. The critical direction in which a ‘bang-bang’ motion is planned is automatically adjusted based on the direction with the largest position error.

##### 6.4.1. Results

Figure 10 shows top views of flights with the two controllers. Both controllers have been assessed on the time it takes to complete one circle and the minimum position error. The results are displayed in Fig. 11. Here, the




Fig. 10. Top-view of consecutive waypoints flights.




Fig. 11. Circle completion time and minimal waypoint distance for consecutive waypoint flights.

‘bang-bang’ controller is seen to outperform the PID both in speed and target accuracy. The PID controller is unable to give priority to one direction over the other. Due to the high gain values, roll and pitch angles are quickly saturated even if the position error of one direction is much smaller than the other, which slows down the critical axis. In the various flight runs of the PID controller, large lateral oscillations can be seen. It was also found that the PID controller was more likely to reach unstable situations due to the high simultaneous pitch and roll angles compared to the ‘bang-bang’ controller.

The predicted times of arrival for a single run are illustrated in Fig. 12. From this figure, we can derive the real-time path prediction performance. As expected the time increases during the angular rotations. However, during the acceleration phases, the time is seen to decrease. We think that this is caused by inaccurate aerodynamic drag




Fig. 12. ETA corrected for passed time.

estimation and by the effect the noncritical angle has on the acceleration in the critical direction.

## 7. Conclusions

We have proposed the ‘bang-bang’ MPC which approaches time-optimal control principles while being computationally efficient enough to run onboard a commercial quadcopter. This is achieved by assuming that the solution consists of a double step control input in attitude angle for one ‘critical’ direction while the noncritical direction has a constant angle as the solution. This simplifies the OCP and drastically reduces the computational complexity. For efficient control parameter optimization, a bisection scheme in combination with an analytical path prediction model is used. We have shown both in simulation and in real-world flights that the ‘bang-bang’ controller is a feasible option for fast flight. In fact, for the consecutive waypoint flight, the ‘bang-bang’ controller is shown to be 17.5% faster than a traditional high-gain PID controller on average. The entire control pipeline easily runs at the main control loop frequency of 512 Hz on the Bebop and is sufficiently light to run on even smaller and computationally-limited quadcopters. However, a more thorough analysis is needed to quantify the required computational effort.

The ‘bang-bang’ MPC also shows promise to be an attractive easy-to-implement solution for different quadcopters. As the pipeline requires minimal knowledge of the dynamics (only  $C_d$  and mass). Despite that the transition compensation in its current state relies on measurement data, future work could mitigate this process with online transition loss estimation. Currently, the constant altitude constraint forms the largest deviation from the theoretical time-optimum solution found by ICLOCS. Finally, the pitch, roll, and thrust limits are set conservatively and flight performance could be improved if these parameters are made adaptive.

## Appendix A. Step Response Results

In this appendix, a compilation of the motion primitive flight experiments results from Sec. 6.3 are presented, showing the trajectory and attitude commands of each different controller for the corresponding maneuver. Additionally, a composition of recorded video frames is included. Where the drone colors Red, Blue and Green correspond with PID, ‘Bang-Bang’, and ‘Bang-Bang’ with transition compensation, respectively.

### Forward Maneuver

Figure A.1 displays the step maneuver in which the quadcopter is instructed to fly 4.5 m forward with a maximum pitch angle of  $35^\circ$ .




Fig. A.1. Forward maneuver.

### Backward Maneuver

The maneuver depicted in Fig. A.2 is similar to the forward maneuver except that it will fly backward.




Fig. A.2. Backward maneuver.




Fig. A.3. Sideway maneuver.




Fig. A.4. Forward-sideway maneuver.

## Sideways Maneuver

The sideways maneuver shown in Fig. A.3 tests the steps response for the roll angle only. The quadcopter is instructed to fly 4.5 m sideways with a maximum roll angle of  $35^\circ$ .

## Forward-Sideways Maneuver

The forward-sideways maneuver combines the longitudinal and lateral step performance, as can be seen in Fig. A.4. The maximum absolute roll and pitch commands are set to  $35^\circ$ .

## References

- [1] M. Hassanalian and A. Abdelkefi, Classifications, applications, and design challenges of drones: A review, *Prog. Aerosp. Sci.* **91** (2017) 99–131.
- [2] H. Moon, Y. Sun, J. Baltes and S. J. Kim, The IROS 2016 competitions [competitions], *IEEE Robot. Autom. Mag.* **24**(1) (2017) 20–29.
- [3] H. Moon, J. Martinez-Carranza, T. Cieslewski, M. Faessler, D. Falanga, A. Simovic, D. Scaramuzza, S. Li, M. Ozo, C. De Wagter, G. de Croon, S. Hwang, S. Jung, H. Shim, H. Kim, M. Park, T. C. Au and S. J. Kim, Challenges and implemented technologies used in autonomous drone racing, *Intell. Serv. Robot.* **12**(2) (2019) 137–148.
- [4] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar and D. Scaramuzza, AlphaPilot: Autonomous drone racing, *Robotics: Science and Systems 2020* (Corvallis, OR, USA, 2020), pp. 1–9.
- [5] M. Hehn, R. Ritz and R. D'Andrea, Performance benchmarking of quadrotor systems using time-optimal control, *Auton. Robots* **33** (1–2) (2012) 69–88.
- [6] E. Tal and S. Karaman, Accurate tracking of aggressive quadrotor trajectories using incremental nonlinear dynamic inversion and differential flatness, in *Proc. IEEE Conf. Decision and Control*, 2018 (2019), pp. 4282–4288.
- [7] D. Mellinger and V. Kumar, Minimum snap trajectory generation and control for quadrotors, in *2011 IEEE Int. Conf. Robotics and Automation* (IEEE, Shanghai, 2011), pp. 2520–2525.
- [8] M. Faessler, A. Franchi and D. Scaramuzza, Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories, *IEEE Robot. Autom. Lett.* **3** (2017) 620–626.
- [9] C. Richter, A. Bry and N. Roy, Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments, in *Robotics Research the 16th Int. Symp.*, Springer Tracts in Advanced Robotics, Vol. 114 (Springer, 2016), pp. 649–666.
- [10] D. Falanga, P. Foehn, P. Lu and D. Scaramuzza, PAMPC: Perception-aware model predictive control for quadrotors, in *2018 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (IEEE, Madrid, Spain, 2018), pp. 1–8.
- [11] 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, in *Proc. — IEEE Int. Conf. Robotics and Automation* (IEEE, Montreal, QC, Canada, 2019), pp. 690–696.
- [12] X. Wang, L. Xi, Y. Chen, S. Lai, F. Lin and B. M. Chen, Decentralized MPC-based trajectory generation for multiple quadrotors in cluttered environments, *Guid. Navig. Control* **1**(2) (2021) 2150007–1–20.
- [13] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun and D. Scaramuzza, Deep drone racing: From simulation to reality with domain randomization, *IEEE Trans. Robot.* **36** (2019) 1–14.
- [14] S. Li, E. Öztürk, C. De Wagter, G. C. H. E. de Croon and D. Izzo, Aggressive online control of a quadrotor via deep network representations of optimality principles, in *2020 IEEE Int. Conf. Robotics and Automation (ICRA)* (IEEE, Paris, France, 2020), pp. 6282–6287.
- [15] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun and D. Scaramuzza, Deep drone acrobatics, *RSS: Robot. Sci. Syst.* (2020), pp. 1–10.
- [16] T. Baca, G. Loianno and M. Saska, Embedded model predictive control of unmanned micro aerial vehicles, in *2016 21st Int. Conf. Methods and Models in Automation and Robotics (MMAR)* (IEEE, Miedzydroje, Poland, 2016), pp. 992–997.
- [17] S. Li, E. Horst, P. Duernay, C. D. Wagter and G. C. H. E. Croon, Visual model-predictive localization for computationally efficient autonomous racing of a 72-gram drone, *J. Field Robot.* **37** (2020) 667–692.
- [18] D. E. Kirk, *Optimal Control Theory: An Introduction* (Dover Publications, 2012).
- [19] Y. Nie, O. Faqir and E. C. Kerrigan, ICLOCS2: Try this optimal control problem solver before you try the rest, in *2018 UKACC 12th Int. Conf. Control*, Vol. 2 (2017) (IEEE, 2018), p. 336.
- [20] P. Brisset, A. Drouin, M. Gorraz, P.-S. Huard and J. Tyler, The paparazzi solution, in *MAV 2006, 2nd US-European Competition and Workshop on Micro Air Vehicles* (Sandestin, United States, 2006), HAL-01004157, pp. 1–15.
- [21] E. J. Smeur, Q. Chu and G. C. De Croon, Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles, *J. Guid. Control Dyn.* **39**(3) (2016) 450–461.



**Jelle Westenberger** received his M.Sc. in Aerospace Engineering, specializing in control and simulation, at the Delft University of Technology in 2021. He graduated on the topic of computational efficient time-optimal control for quadcopters. His interests lie in the fields of autonomous control systems, robotics, and sustainable energy technologies. He has won several awards during his studies, such as the 'André Kuipers Ruimtevaart Prijs' of the Dutch Aerospace Awards in 2017, and the Alphapilot Autonomous Drone Racing Challenge in 2019.



**Christophe De Wagter** received his Master of Science in Aerospace Engineering at the Delft University of Technology on the topic of vision-based control in 2004. In 2005 he started the Micro Air Vehicle Lab, and has worked there as a researcher since then. His areas of interest range from control theory and sensor fusion to computer vision and electronics. In parallel, he has been working as a freelance developer for local startup companies. He received awards ranging from the first prize for 'Best Fully Autonomous Indoor MAV' at the EMAV 2008 in Braunschweig, to the 'world champion AI Drone Racing 2019'.



**Guido de Croon** received his M.Sc. and Ph.D. in the field of Artificial Intelligence (AI) at the Maastricht University, the Netherlands. His research interest lies in computationally efficient, bio-inspired algorithms for robot autonomy, with an emphasis on computer vision. Since 2008 he has worked on algorithms for achieving autonomous flight with small and light-weight flying robots, such as the DelFly flapping wing MAV. In 2011–2012, he was a research fellow in the Advanced Concepts Team of the European Space Agency, where he studied topics such as optical-flow-based control algorithms for extraterrestrial landing scenarios. After his return at TU Delft, his work has included the fully autonomous flight of a 20-gram DelFly, a new theory on active distance perception with optical flow, and a swarm of tiny drones able to explore unknown environments. Currently, he is a Full Professor at TU Delft and scientific lead of the Micro Air Vehicle lab (MAVLab) of the Delft University of Technology.