

A Flying Inverted Pendulum

Markus Hehn and Raffaello D’Andrea

Abstract—We extend the classic control problem of the inverted pendulum by placing the pendulum on top of a quadrotor aerial vehicle. Both static and dynamic equilibria of the system are investigated to find nominal states of the system at standstill and on circular trajectories. Control laws are designed around these nominal trajectories. A yaw-independent description of quadrotor dynamics is introduced, using a ‘Virtual Body Frame’. This allows for the time-invariant description of curved trajectories. The balancing performance of the controller is demonstrated in the ETH Zurich Flying Machine Arena testbed. Development potential for the future is highlighted, with a focus on applying learning methodology to increase performance by eliminating systematic errors that were seen in experiments.

I. INTRODUCTION

The inverted pendulum is a classic control problem, offering one of the most intuitive, easily describable and realizable nonlinear unstable systems. It has been investigated for several decades (see, for example, [11], and references therein). It is frequently used as a demonstrator to showcase theoretical advances, e.g. in reinforcement learning [6], neural networks [14], and fuzzy control [13].

In this paper, we develop a control strategy that enables an inverted pendulum to balance on top of a quadrotor. Besides being a highly visual demonstration of the dynamic capabilities of modern quadrotors, the solution to such a complex control problem offers insight into quadrotor control strategies, and could be adapted to other tasks.

Quadrotors offer exceptional agility. Thanks to the off-center mounting of the propellers, extraordinarily fast rotational dynamics can be achieved. This is combined with typically high thrust-to-weight ratios, resulting in large achievable translational accelerations when not carrying a payload.

While most early work on quadrotors focused on near-hover operation (e.g. [5], and references therein), a growing community is working on using the full dynamical potential of these vehicles. Flips have been executed by several groups, some focusing on speed and autonomous learning [7] and some on safety guarantees [3]. Other complex maneuvers, including flight through windows and perching have been demonstrated [8].

In Section II, we introduce the dynamic models used in the controller design. Section III presents static and dynamic nominal trajectories for the quadrotor to follow. The dynamics are then linearized around these trajectories in Section IV, and linear state feedback controllers are designed in Section V. The experimental setup and results are shown

The authors are with the Institute for Dynamic Systems and Control, ETH Zurich, Switzerland. {hehnm, rdandrea}@ethz.ch




Fig. 1. The inertial coordinate system \mathbf{O} and the vehicle coordinate system \mathbf{V} .

in Section VI and conclusions are drawn in Section VII, where an outlook is also presented.

II. DYNAMICS

We derive the equations of motion of the quadrotor and the inverted pendulum for the trajectory-independent general case.

Given that the mass of the pendulum is small compared to the mass of the quadrotor, it is reasonable to assume that the pendulum’s reactive forces on the quadrotor are negligible. The dynamics of the quadrotor, then, do not depend on the pendulum, whereas the dynamics of the pendulum are influenced by the motion of the quadrotor. This assumption is justified by the experimental setup, with the weight of the pendulum being less than 5% of that of the quadrotor vehicle.

A. Quadrotor

The quadrotor is described by six degrees of freedom: The translational position (x, y, z) is measured in the inertial coordinate system \mathbf{O} as shown in Figure 1. The vehicle attitude \mathbf{V} is defined by three Euler angles. From the inertial coordinate system, we first rotate around the z -axis by the yaw angle α . The coordinate system is then rotated around the new y -axis by the pitch angle β , and finally rotated about the new x -axis by the roll angle γ :

$${}^{\mathbf{O}}R(\alpha, \beta, \gamma) = R_z(\alpha) R_y(\beta) R_x(\gamma) , \quad (1)$$

where

$$R_x(\gamma) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \gamma & -\sin \gamma \\ 0 & \sin \gamma & \cos \gamma \end{bmatrix} , \quad (2)$$

$$R_y(\beta) = \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix}, \quad (3)$$

$$R_z(\alpha) = \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

The translational acceleration of the vehicle is dictated by the attitude of the vehicle and the total thrust produced by the four propellers. With a representing the mass-normalized collective thrust, the translational acceleration in the inertial frame is

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix} = {}^O_V R(\alpha, \beta, \gamma) \begin{bmatrix} 0 \\ 0 \\ a \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -g \end{bmatrix}. \quad (5)$$

The vehicle attitude is not directly controllable, but it is subject to dynamics. The control inputs are the desired rotational rates about the vehicle body axes, $(\omega_x, \omega_y, \omega_z)$, and the mass-normalized collective thrust, a , as shown in Figure 2. High-bandwidth controllers on the vehicle track the desired rates using feedback from gyroscopes. The quadrotor has very low rotational inertia, and can produce high torques due to the outward mounting of the propellers, resulting in very high achievable rotational accelerations on the order of 200 rad/s². The vehicle has a fast response time to changes in the desired rotational rate (experimental results have shown time constants on the order of 20 ms). We will therefore assume that we can directly control the vehicle body rates and ignore rotational acceleration dynamics. As with the vehicle body rates, we assume that the thrust can be changed instantaneously. Experimental results have shown that the true thrust dynamics are about as fast as the rotational dynamics, with propeller spin-up being noticeably faster than spin-down.

The rates of the Euler angles are converted to the vehicle body coordinate system V through their respective transformations:

$$\begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix} = \begin{bmatrix} \dot{\gamma} \\ 0 \\ 0 \end{bmatrix} + R_x^{-1}(\gamma) \begin{bmatrix} 0 \\ \dot{\beta} \\ 0 \end{bmatrix} + R_x^{-1}(\gamma) R_y^{-1}(\beta) \begin{bmatrix} 0 \\ 0 \\ \dot{\alpha} \end{bmatrix}. \quad (6)$$

The above can be written more compactly by combining the Euler rates into a single vector, calculating the relevant rows of the rotation matrices, and solving for the Euler angle rates:

$$\begin{bmatrix} \dot{\gamma} \\ \dot{\beta} \\ \dot{\alpha} \end{bmatrix} = \begin{bmatrix} \cos \beta \cos \gamma & -\sin \gamma & 0 \\ \cos \beta \sin \gamma & \cos \gamma & 0 \\ -\sin \beta & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}. \quad (7)$$

B. Inverted Pendulum

The pendulum has two degrees of freedom, which we describe by the translational position of the pendulum center of mass relative to its base in \mathbf{O} (r along the x-axis, s along

the y-axis). For notational simplicity, we describe the relative position of the pendulum along the z-axis as

$$\zeta := \sqrt{L^2 - r^2 - s^2}, \quad (8)$$

where L denotes the length from the base of the pendulum to its center of mass. We model the pendulum as an inertialess point mass that is rigidly attached to the mass center of the quadrotor, such that rotations of the vehicle do not cause a motion of the pendulum base. In the experimental setup, the point that the pendulum is attached to is mounted off-center by about 10% of the length of the pendulum. While this assumption causes modeling errors, it simplifies the dynamics to such a great extent that the problem becomes much more tractable. The Lagrangian [9] of the pendulum can be written as

$$\mathcal{L} = \frac{1}{2} \left((\dot{x} + \dot{r})^2 + (\dot{y} + \dot{s})^2 + \left(\dot{z} - \frac{r\dot{r} + s\dot{s}}{\zeta} \right)^2 \right) - g(z + \zeta), \quad (9)$$

where we assume unit pendulum mass without loss of generality. The first term represents the kinetic energy of the pendulum, and the second the potential energy. The full, nonlinear dynamic equations can be derived from \mathcal{L} using conventional Lagrangian mechanics:

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{r}} \right) - \frac{\partial \mathcal{L}}{\partial r} = 0 \quad (10)$$

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{s}} \right) - \frac{\partial \mathcal{L}}{\partial s} = 0, \quad (11)$$

resulting in a system of equations of the form

$$\begin{bmatrix} \ddot{r} \\ \ddot{s} \end{bmatrix} = \mathbf{f}(r, s, \dot{r}, \dot{s}, \ddot{x}, \ddot{y}, \ddot{z}), \quad (12)$$

where \mathbf{f} are the nonlinear equations (13) and (14).

C. Combined dynamics

The full dynamics of the combined system are described entirely by Equations (5), (7), and (12). The three body rate control inputs $(\omega_x, \omega_y, \omega_z)$ control the attitude \mathbf{V} of the vehicle in a nonlinear fashion. This attitude, combined




Fig. 2. The control inputs of the quadrotor: The rotational rates ω_x , ω_y , and ω_z are tracked by an on-board controller, using gyroscope feedback.

$$\ddot{r} = \frac{1}{(L^2 - s^2)\zeta^2} \left(-r^4 \ddot{x} - (L^2 - s^2)^2 \ddot{x} - 2r^2 (s\dot{r}\dot{s} + (-L^2 + s^2) \ddot{x}) + r^3 (\dot{s}^2 + s\ddot{s} - \zeta(\mathbf{g} + \ddot{z})) + r(-L^2 s\ddot{s} + s^3 \ddot{s} + s^2 (\dot{r}^2 - \zeta(\mathbf{g} + \ddot{z})) + L^2 (-\dot{r}^2 - \dot{s}^2 + \zeta(\mathbf{g} + \ddot{z}))) \right) \quad (13)$$

$$\ddot{s} = \frac{1}{(L^2 - r^2)\zeta^2} \left(-s^4 \ddot{y} - (L^2 - r^2)^2 \ddot{y} - 2s^2 (r\dot{r}\dot{s} + (-L^2 + r^2) \ddot{y}) + s^3 (\dot{r}^2 + r\ddot{r} - \zeta(\mathbf{g} + \ddot{z})) + s(-L^2 r\ddot{r} + r^3 \ddot{r} + r^2 (\dot{s}^2 - \zeta(\mathbf{g} + \ddot{z})) + L^2 (-\dot{r}^2 - \dot{s}^2 + \zeta(\mathbf{g} + \ddot{z}))) \right) \quad (14)$$

with the thrust a , controls the translational acceleration of the vehicle. While the acceleration drives the translational motion of the vehicle linearly, it also drives the motion of the pendulum through nonlinear equations. The combined system consists of thirteen states (three rotational and six translational states of the quadrotor, and four states of the pendulum), and four control inputs (three body rates, and the thrust).

III. NOMINAL TRAJECTORIES

In this section, we find static and dynamic equilibria of the system that satisfy Equations (5), (7), and (12). These are used as nominal trajectories to be followed by the quadrotor. Corresponding nominal control inputs are also described. We denote nominal values by a zero index (x_0 , r_0 , etc.).

A. Constant position

In a first case, we require x_0 , y_0 , and z_0 to be constant. Substituting these constraints into (5), it can be seen that $\beta_0 = 0$ and $\gamma_0 = 0$ solve the equations with $a_0 = g$, while α_0 can be chosen freely. We arbitrarily choose to set $\alpha_0 = 0$.

Using the given angles $(\alpha_0, \beta_0, \gamma_0)$, equation (7) can be solved with the body rate control inputs being $\omega_{x_0} = \omega_{y_0} = \omega_{z_0} = 0$.

Inserting the nominal states (x_0, y_0, z_0) into the pendulum equations of motion (12), they simplify to

$$\ddot{r} = r \frac{g\zeta^3 - L^2 (\dot{r}^2 + \dot{s}^2) + (s\dot{r} - r\dot{s})^2}{L^2 \zeta^2}, \quad (15)$$

$$\ddot{s} = s \frac{g\zeta^3 - L^2 (\dot{r}^2 + \dot{s}^2) + (s\dot{r} - r\dot{s})^2}{L^2 \zeta^2}. \quad (16)$$

These equations are solved by the static equilibrium

$$r = r_0 = 0 \quad (17)$$

$$s = s_0 = 0, \quad (18)$$

meaning that, as expected, the inverted pendulum is exactly over the quadrotor.

B. Circular trajectory

As a second nominal trajectory, the quadrotor is required to fly a circle of a given radius R at a constant rotational rate Ω , at a constant altitude z_0 .

We seek to transform the equations of motion into different coordinate systems, such that the nominal states and the linearized dynamics about them can be described in a time-invariant manner.

To describe the vehicle position, the following coordinate system \mathbf{C} is introduced, with (u, v, w) describing the position in \mathbf{C} :

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} =: R_z(\Omega t) \begin{bmatrix} u \\ v \\ w \end{bmatrix} = \begin{bmatrix} \cos \Omega t & -\sin \Omega t & 0 \\ \sin \Omega t & \cos \Omega t & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix}. \quad (19)$$

To describe the vehicle attitude, a second set of Euler angles is introduced, describing the ‘virtual body frame’ \mathbf{W} and named η , μ , and ν :

$${}^O_W R(\eta, \mu, \nu) = R_z(\eta) R_y(\mu) R_x(\nu), \quad (20)$$

subject to the constraint that

$${}^O_W R(\alpha, \beta, \gamma) \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = {}^O_W R(\eta, \mu, \nu) \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \quad (21)$$

The above equation defines values for three elements of the rotation matrices. As every column of a rotation matrix has unit norm [12], however, this equation only defines two of the angles (η, μ, ν) .

Comparing the constraint (21) with the translational equation of motion of the vehicle (5), it is straightforward to see that the virtual body frame \mathbf{W} represents an attitude that is constrained such that it effects the same translational motion of the quadrotor as the vehicle attitude \mathbf{V} . The remaining degree of freedom represents the fact that rotations about the axis along which ω_z acts have no affect on the quadrotors translational motion.

Applying (19), its derivatives, (21), and setting the free parameter $\eta = \Omega t$, the quadrotor equation of motion (5) simplifies to

$$\begin{bmatrix} \ddot{u} \\ \ddot{v} \\ \ddot{w} \end{bmatrix} = \begin{bmatrix} a \sin \mu \cos \nu + \Omega^2 u + 2\Omega \dot{v} \\ -a \sin \nu - 2\Omega \dot{u} + \Omega^2 v \\ a \cos \mu \cos \nu - g \end{bmatrix}. \quad (22)$$

The circular trajectory is described by $u_0 = R$, $v_0 = 0$, and $\dot{w}_0 = 0$. Using these values, the nominal Euler angles μ_0 and ν_0 , and the nominal thrust a_0 can be calculated:

$$\mu_0 = \arctan\left(-\frac{\Omega^2 R}{g}\right), \quad (23)$$

$$\nu_0 = 0, \quad (24)$$

$$a_0 = \sqrt{g^2 + (\Omega^2 R)^2}. \quad (25)$$

Knowing the nominal values for (η_0, μ_0, ν_0) , we solve for $(\alpha_0, \beta_0, \gamma_0)$ using Equation (21). Analogous to the constant

position case, we set $\alpha_0 = 0$, simplifying (21) to

$$\begin{bmatrix} \sin \beta_0 \cos \gamma_0 \\ -\sin \gamma_0 \\ \cos \beta_0 \cos \gamma_0 \end{bmatrix} = \begin{bmatrix} \cos \Omega t \sin \mu_0 \cos \nu_0 + \sin \Omega t \sin \nu_0 \\ \sin \Omega t \sin \mu_0 \cos \nu_0 - \cos \Omega t \sin \nu_0 \\ \cos \mu_0 \cos \nu_0 \end{bmatrix}, \quad (26)$$

which can be solved for β_0 and γ_0 . This completes the description of the nominal states required for the translational motion (5): In the coordinate systems \mathbf{C} and \mathbf{W} , the nominal position and attitude are constant. Using Equations (19) and (26), the time-varying nominal states in \mathbf{O} and \mathbf{V} may be found.

To calculate the rotational rate control inputs in Equation (7), we take the first derivative of Equation (26). It can be shown that

$$\dot{\beta}_0 = \frac{R\Omega^3 \cos^{-1} \gamma_0 (\tan \beta_0 \tan \gamma_0 \cos(\Omega t) + \cos^{-1} \beta_0 \sin(\Omega t))}{\sqrt{g^2 + (\Omega^2 R)^2}} \quad (27)$$

$$\dot{\gamma}_0 = \frac{R\Omega^3 \cos^{-1} \gamma_0 \cos(\Omega t)}{\sqrt{g^2 + (\Omega^2 R)^2}}. \quad (28)$$

Combining these equations with the results from Equations (26) and (7), the nominal states can be solved for the nominal control inputs $(\omega_{x_0}, \omega_{y_0}, \omega_{z_0})$. The full derivation is made available online at www.idsc.ethz.ch/people/staff/hehn-m.

Identically to the vehicle, the pendulum relative coordinates r and s are rotated by Ωt :

$$\begin{bmatrix} r \\ s \end{bmatrix} =: \begin{bmatrix} \cos \Omega t & -\sin \Omega t \\ \sin \Omega t & \cos \Omega t \end{bmatrix} \begin{bmatrix} p \\ q \end{bmatrix}. \quad (29)$$

Applying this rotation to the Lagrangian derivations of the motion of the pendulum (10), (11) and setting the base motion $(\ddot{x}, \ddot{y}, \ddot{z})$ to the circular trajectory, the pendulum dynamics can be shown to be

$$p \left(\frac{p\ddot{p} + \dot{p}^2 + q\ddot{q} + \dot{q}^2}{\zeta^2} + \frac{q^2\dot{q}^2 + p^2\dot{p}^2 + 2pq\dot{p}\dot{q}}{\zeta^4} - \frac{g}{\zeta} - \Omega^2 \right) + \ddot{p} - 2\Omega\dot{q} - R\Omega^2 = 0 \quad (30)$$

$$q \left(\frac{q\ddot{q} + \dot{q}^2 + p\ddot{p} + \dot{p}^2}{\zeta^2} + \frac{p^2\dot{p}^2 + 2pq\dot{p}\dot{q} + q^2\dot{q}^2}{\zeta^4} - \frac{g}{\zeta} - \Omega^2 \right) + \ddot{q} + 2\Omega\dot{p} = 0 \quad (31)$$

We seek a solution where $\dot{p}_0 = \dot{q}_0 = \ddot{p}_0 = \ddot{q}_0 = 0$, leading to the following constraints for equilibrium points:

$$\Omega^2(q_0) + \frac{gq_0}{\zeta_0} = 0, \quad (32)$$

$$\Omega^2(R + p_0) + \frac{gp_0}{\zeta_0} = 0. \quad (33)$$

The only solution to the first equation is $q_0 = 0$. The second equation defines a nonlinear relationship between Ω , R , and p_0 , with a solution p_0 always existing in the range $-R \leq p_0 \leq 0$. This result is intuitive: The center of mass of the pendulum must lie towards the center of the circle, such

that the centripetal force acting on it is compensated for by gravity. If the center of mass were to lie further than R inwards, the centripetal force would change sign, making the pendulum fall.

IV. DYNAMICS ABOUT NOMINAL TRAJECTORIES

Linear approximate dynamics are derived through a first-order Taylor expansion of the equations of motion (5), (7), and (12) about the nominal trajectories found in Section III. We denote small deviations by a tilde (\tilde{x} , \tilde{r} , etc.). We present only the resulting linearized dynamics. The corresponding derivations are not shown in this paper due to space constraints, but are made available online.

A. Constant position

Assuming a constant nominal position and zero yaw angle, the three translational degrees of freedom of the quadrotor-pendulum system decouple entirely along the three axes of the \mathbf{O} coordinate system, resulting in the following equations:

$$\ddot{\tilde{r}} = \tilde{r} \frac{g}{L} - \tilde{\beta} g \quad (34)$$

$$\ddot{\tilde{s}} = \tilde{s} \frac{g}{L} + \tilde{\gamma} g \quad (35)$$

$$\ddot{\tilde{x}} = \tilde{\beta} g \quad (36)$$

$$\ddot{\tilde{y}} = -\tilde{\gamma} g \quad (37)$$

$$\dot{\tilde{\beta}} = \tilde{\omega}_y \quad (38)$$

$$\dot{\tilde{\gamma}} = \tilde{\omega}_x \quad (39)$$

$$\ddot{\tilde{z}} = \tilde{a} \quad (40)$$

The two horizontal degrees of freedom represent fifth-order systems, with the vehicle forming a triple integrator from the body rate to its position. The vertical motion is represented by a double integrator from thrust to position.

B. Circular trajectory

To derive linear dynamics about the circular trajectory, the Euler angle rates $\dot{\mu}$ and $\dot{\nu}$ are treated as control inputs. Solving the time-derivative of equation (21) (analogously to solutions (27) and (28) for the nominal case) allows the calculation of $(\dot{\tilde{\beta}}, \dot{\tilde{\gamma}})$. These can then be converted to the true inputs $(\tilde{\omega}_x, \tilde{\omega}_y, \tilde{\omega}_z)$ using Equation (7).

In contrast to a constant nominal position, the dynamics on a circular trajectory do not decouple. The linearized equations of motion can be shown to be

$$\ddot{\tilde{p}} = \frac{\zeta_0^2}{L^2} \left[\tilde{p}(\Omega^2 + \frac{gL^2}{\zeta_0^3}) + 2\dot{q}\dot{\Omega} + \right. \quad (41)$$

$$\left. \tilde{\mu}(-\frac{p_0}{\zeta_0} a_0 \sin \mu_0 - a_0 \cos \mu_0) + \tilde{a}(\frac{p_0}{\zeta_0} \cos \mu_0 - \sin \mu_0) \right]$$

$$\ddot{\tilde{q}} = \tilde{q}(\Omega^2 + \frac{g}{\zeta_0}) - 2\dot{p}\dot{\Omega} + \tilde{\nu} a_0 \quad (42)$$

$$\ddot{\tilde{u}} = \tilde{a} \sin \mu_0 + \tilde{\mu} a_0 \cos \mu_0 + 2\tilde{\nu}\dot{\Omega} + \tilde{u}\Omega^2 \quad (43)$$

$$\ddot{\tilde{v}} = -\tilde{\nu} a_0 - 2\tilde{u}\dot{\Omega} + \tilde{v}\Omega^2 \quad (44)$$

$$\ddot{\tilde{w}} = \tilde{a} \cos \mu_0 - \tilde{\mu} a_0 \sin \mu_0 \quad (45)$$

We observe that the linearized dynamics are indeed time-invariant in the coordinate systems \mathbf{C} and \mathbf{W} .

The above equations reduce to Equations (34) – (40) if setting $R = 0$ and $\Omega = 0$. If $R = 0$, Ω is a free parameter (see Equation (33)), allowing the description of the dynamics (34) – (40) in a rotating coordinate system.

V. CONTROLLER DESIGN

We design linear full state feedback controllers to stabilize the system about its nominal trajectories. We use an infinite-horizon linear-quadratic regulator (LQR) design [1] and determine suitable weighting matrices.

A. Constant position

Because the system is decoupled in its nominal state, the control design process can be separated. The two horizontal degrees of freedom are single-input, five-state systems. The vertical degree of freedom is a single-input, two-state system. Although simpler design methods exist for such systems, LQR is used to make the results easily transferable to the design for a circular trajectory. We design a lateral controller that is identically applied to the \tilde{x} - \tilde{r} -system and the \tilde{y} - \tilde{s} -system (except for different signs mirroring the signs in the equations of motion (34)–(37)). A controller for the vertical direction is designed separately.

For the lateral controllers, we penalize only the vehicle position (\tilde{x} or \tilde{y}) and the control effort ($\tilde{\omega}_y$ or $\tilde{\omega}_x$). There is no penalty on the pendulum state. One tuning parameter remains: The ratio of the penalties on position and control effort controls the speed at which the position set point is tracked. Values for this ratio are tuned manually until the system shows fast performance, without saturating the control inputs. This tuning is initially carried out in simulation, and then refined on the experimental setup.

The vertical controller is tuned much in the same manner as the lateral controller. Here again, we tune only the ratio between penalties on position errors and control effort until satisfying performance is achieved.

B. Circular trajectory

On the circular trajectory, the system represents a thirteen-state system with three control inputs that cannot readily be decoupled. To more easily tune the weighting matrices, we use the same approach as in the constant position case and penalize only the position errors (\tilde{u} , \tilde{v} , and \tilde{w}) and control effort ($\tilde{\mu}$, $\tilde{\nu}$, and $\tilde{\alpha}$). The relative size of weights on control inputs and states is carried over from the standstill design. Because the controller for the vertical axis was tuned separately in the standstill case, the relative size of the penalties on the horizontal positions (\tilde{u} , \tilde{v}) and the vertical position (\tilde{w}) is adapted. Again, this is first carried out in simulation and then improved upon using the experimental testbed.

VI. EXPERIMENTAL RESULTS

The algorithms presented herein were implemented in the Flying Machine Arena, an aerial vehicle development

platform at ETH Zurich [7]. We present results demonstrating the performance of the controllers designed in the previous Section.

A. Experimental setup

We use modified Ascending Technologies ‘Hummingbird’ quadrotors [4]. The vehicles are equipped with custom electronics, allowing greater control of the vehicle’s response to control inputs, a higher dynamic range, and extended interfaces [7]. A small cup-shaped pendulum mounting point is attached to the top of the vehicle, approximately 5 cm above the geometrical center of the vehicle. The pendulum can rotate freely about the mounting point up to an angle of approximately 50 degrees. At larger angles, the mounting point offers no support and the pendulum falls off the vehicle.

Commands are sent through a proprietary low-latency 2.4 GHz radio link at a frequency of 50 Hz. Command loss is in the range of 0.1%. An infrared motion tracking system provides precise vehicle position and attitude measurements at 200 Hz, using retro-reflective markers mounted to the vehicle. The total closed-loop latency is approximately 30 ms.

The inverted pendulum consists of a carbon fiber tube, measuring 1.15 m in length. The top end of the pendulum carries a retro-reflective marker, allowing the position of this point to be determined through the motion tracking system in the same manner as vehicles are located. The center of mass of the pendulum is 0.565 m away from its base. Figure 3 shows the quadrotor and the pendulum.

Conventional desktop computers are used to run all control algorithms, with one computer acting as an interface to the




Fig. 3. The quadrotor balancing the pendulum, at standstill. The mass center is about half-length of the pendulum.




Fig. 4. State errors during balancing: Pendulum position error (\tilde{r} , \tilde{s}) and quadrotor position error (\tilde{x} , \tilde{y}). The pendulum is manually placed on the vehicle at approximately $t = 4.25$ s, and the balancing controller is switched on at approximately $t = 4.75$ s.

testbed. Data is exchanged over Ethernet connections. A Luenberger observer is used to filter the sensory data and provide full state information to the controller. The observer also compensates for systematic latencies occurring in the control loop, using the known control inputs to project the system state into the future.

B. Constant position

Experiments are initialized by manually placing the pendulum on the mounting point. The vehicle holds a constant position using a separate controller, waiting for the pendulum. The balancing controller is switched on if \tilde{r} and \tilde{s} are sufficiently small for 0.5 seconds.

Figure 4 shows the pendulum position errors (\tilde{r} , \tilde{s}) and the horizontal quadrotor position errors (\tilde{x} , \tilde{y}). The pendulum is placed on the vehicle at approximately $t = 4.25$ s, and the control is switched from position holding to balancing at approximately $t = 4.75$ s. The pendulum position errors are relatively large in the beginning, but quickly converge to values close to zero. The vehicle settles at a stationary offset on the order of 5 cm from the desired position. Note that the balancing controller does not provide feedback on integrated errors. The main suspected reason for these steady-state errors are miscalibrations in the system: Errors in the vehicle attitude measurement lead to the linear controller trading off the attitude error and the position error, and biased measurements of the on-board gyroscopes result in a biased response to control inputs.

Though originally designed for a constant position, this controller has been successfully tested for set point tracking at moderate speeds. A video of this is available online at www.idsc.ethz.ch/people/staff/hehn-m.

C. Circular trajectory

We arbitrarily choose to set $p_0 = -\frac{R}{2}$, bringing the center of mass of the pendulum half-way between the vehicle and the circle center, nominally. Assuming a given R , this fixes Ω through the equilibrium constraint (33). Figure 5 shows the system performance when circling. The pendulum is first balanced at a constant position. At $t = 2$ s, a

switch to a circular nominal trajectory and a corresponding controller occurs, with $R = 0.1$ m. The controller is seen to stabilize the pendulum, with the pendulum relative position errors in the rotating coordinate system (p , q) converging to non-zero values. The vehicle errors show two distinct components: Like the pendulum errors, there is a non-zero mean error. Additionally, the error oscillates at the rotational rate Ω , representing a near-constant position error in an inertial coordinate system. Figure 6 shows a comparison of the actual and nominal trajectories of the vehicle and pendulum in the inertial coordinate system \mathbf{O} . It confirms that the oscillating errors in the rotating coordinate system are constant position errors in the inertial coordinate system, with a magnitude of approximately 0.1m. The mean errors in \mathbf{C} are represented by the circle radius being significantly larger than the nominal value R .

A video showing the experiments of both cases presented herein is available online at www.idsc.ethz.ch/people/staff/hehn-m.




Fig. 5. Errors in the rotating coordinate system: Pendulum position error (\tilde{p} , \tilde{q}) and quadrotor position error (\tilde{u} , \tilde{v}). The pendulum is balanced by the quadrotor during the entire duration shown on this plot. At $t = 2$ s, the controller is switched from a constant nominal position to a circular trajectory with $R = 0.1$ m.




Fig. 6. The trajectory of the quadrotor and pendulum in \mathbf{O} , compared to the nominal trajectory.




Fig. 7. Simulation results: Errors in the rotating coordinate system: Pendulum position error (\tilde{p} , \tilde{q}) and quadrotor position error (\tilde{u} , \tilde{v}). At $t = 2$ s, the controller is switched from a constant nominal position to a circular trajectory with $R = 0.1$ m.

D. Comparison with simulation results

The controllers have also been tested in simulation. The Flying Machine Arena software environment allows the testing of controllers by simply re-routing the controller's outputs to a simulation. The simulation reproduces the behavior of the entire system. It includes the full dynamics of the quadrotor, including the on-board control loops, rotational accelerations, and propeller dynamics. It also reproduces system latencies and the noise characteristics of sensors. As the simulation output mimics the motion system's output as closely as possible, the same state observer is employed in reality and in simulation. This simulation environment has been extended to include the inverted pendulum. The pendulum is modeled with its full nonlinear dynamics (12), neglecting the off-center mounting on the vehicle for reasons of tractability, as discussed in Section II-B. Figures 7 and 8 show the exact same circular trajectory experiment that was carried out on the testbed (Section VI-C). For this simulation, systematic errors in the gyroscopic sensors and the motion tracking system were disabled. In the initial standstill phase (the first two seconds in Figure 7), all errors are close to zero. During circling, the errors converge to nearly stationary values that are significantly smaller than in the real testbed. The close match of the vehicle's nominal trajectory and the pendulum's simulated trajectory in Figure 8 appears to be coincidental.

This result highlights the influence of biased sensory information, leading to position errors in the inertial coordinate system \mathbf{O} . These are observed as oscillating errors in \mathbf{C} .

The mean errors on the trajectory also show different characteristics in simulation and reality. This is particularly noticeable in the pendulum position error \tilde{p} and \tilde{q} . The simulation contains a detailed dynamic model of the quadrotor that has been validated in several experiments. It is therefore probable that the differences are due mainly to the simulated pendulum dynamics. The non-modeled off-center mounting of the pendulum could explain this discrepancy.

Circle radii R of up to 0.5 m have been successfully tested in simulation and reality. The vehicle and pendulum




Fig. 8. Simulation results: The trajectory of the quadrotor and pendulum in \mathbf{O} , compared to the nominal trajectory.

positioning errors increase with the circle size, but the controllers are still capable of keeping the pendulum in balance.

VII. CONCLUSION AND OUTLOOK

We have developed linear controllers for stabilizing a pendulum on a quadrotor, which can be used for both static and dynamic equilibria of the pendulum. The virtual body frame is a useful tool to describe motions in a convenient coordinate system (e.g. allowing the use of symmetries), without enforcing this rotation for the vehicle. Using its properties and a rotating coordinate system, the system description is time-invariant on circular trajectories.

This allows the straightforward application of well-established state feedback design principles. Controllers for standstill and circular motion have been validated experimentally and are shown to stabilize the pendulum. This key milestone allows us to shift our focus towards improving system performance.

Experimental results revealed systematic errors when applying the control laws. There appear to be different sources of these errors:

- Miscalibrations of sensors cause biases in the experimental setup. These errors are observed in the attitude information from the motion capture system, and in the vehicle on-board control loops using gyroscope feedback.
- The simplifying assumption that the pendulum is mounted at the center of mass of the vehicle is violated in the experimental setup. Rotations of the vehicle therefore cause a motion of the pendulum base point.
- The equations of motion used to derive nominal trajectories and linear models neglect many real-world effects, such as drag and underlying dynamics of the control inputs.
- The control laws are designed assuming continuous-time control, while the vehicle is controlled at only 50Hz.

We have identified two approaches for extending the controller design presented in this paper. The first, and most

straightforward approach, is to include states that represent the integrated errors, and to weigh them appropriately in the controller design. This would permit compensation for some of the systematic errors. For instance, one would expect this to drive the vehicle position errors in the standstill case to zero.

Alternatively, a machine learning approach could be applied. The measurement data indicates that systematic errors greatly dominate stochastic errors. During the circular trajectory in particular, there are systematic, repeated errors that could well be learned and compensated for in a feed-forward fashion. The system could therefore ‘learn’ better nominal trajectories, resulting in a correction of the nominal control inputs. This could, for instance, be accomplished with iterative learning control [10], [2]. The present problem is especially well suited to this type of approach due to its repetitive nature. We are planning to use this experimental setup as a testbed and benchmark for learning methods.

The concept of the virtual body frame is applicable to a wide range of quadrotor control problems that goes beyond balancing a pendulum. It allows the time-invariant description of general circular trajectories if the circle size and rate are constant. We are investigating extensions of this concept to allow its application to more general problems.

ACKNOWLEDGMENTS

The authors wish to thank Sergei Lupashin for his innumerable contributions to the Flying Machine Arena hardware and software, Michael Sherback for his work on the quadrotor control and simulation algorithms, and Angela Schöllig for her work on the Flying Machine Arena testbed.

REFERENCES

- [1] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control, Volume 2*. Athena Scientific, 2007.
- [2] Insik Chin, S. Joe Qin, Kwang S. Lee, and Moonki Cho. A two-stage iterative learning control technique combined with real-time feedback for independent disturbance rejection. *Automatica*, 40(11):1913–1922, November 2004.
- [3] JH Gillula, Haomiao Huang, MP Vitus, and CJ Tomlin. Design of Guaranteed Safe Maneuvers Using Reachable Sets: Autonomous Quadrotor Aerobatics in Theory and Practice. In *IEEE International Conference on Robotics and Automation*, 2010.
- [4] Daniel Gurdan, Jan Stumpf, Michael Achtelik, Klaus-Michael Doth, Gerd Hirzinger, and Daniela Rus. Energy-efficient Autonomous Four-rotor Flying Robot Controlled at 1 kHz. In *IEEE International Conference on Robotics and Automation*, 2007.
- [5] Jonathan P. How, Brett Bethke, Adrian Frank, Daniel Dale, and John Vian. Real-Time Indoor Autonomous Vehicle Test Environment. *IEEE Control Systems Magazine*, (April):51–64, 2008.
- [6] Michail G. Lagoudakis and Ronald Parr. Least-Squares Policy Iteration. *Journal of Machine Learning Research*, 4(6):1107–1149, January 2003.
- [7] Sergei Lupashin, Angela Schöllig, Michael Sherback, and Raffaello D’Andrea. A Simple Learning Strategy for High-Speed Quadcopter Multi-Flips. In *IEEE International Conference on Robotics and Automation*, 2010.
- [8] Nathan Michael, Daniel Mellinger, Quentin Lindsey, and Vijay Kumar. The GRASP Multiple Micro UAV Testbed. *IEEE Robotics and Automation Magazine*, 17(3):56 – 65, 2010.
- [9] Francis C. Moon. *Applied Dynamics: With Applications to Multibody and Mechatronic Systems*. Wiley, 1998.
- [10] Angela Schöllig and Raffaello D’Andrea. Optimization-Based Iterative Learning Control for Trajectory Tracking. In *European Control Conference*, pages 1505–1510, 2009.
- [11] Jinglai Shen, Amit K. Sanyal, Nalin A. Chaturvedi, Dennis Bernstein, and Harris McClamroch. Dynamics and control of a 3D pendulum. In *43rd IEEE Conference on Decision and Control*, pages 323–328 Vol.1, 2004.
- [12] Gilbert Strang. *Linear Algebra and its Applications*. Thomson Brooks/Cole, 2006.
- [13] H.O. Wang, K. Tanaka, and M.F. Griffin. An approach to fuzzy control of nonlinear systems: Stability and design issues. *IEEE Transactions on Fuzzy Systems*, 4(1):14–23, 1996.
- [14] Victor Williams and Kiyotoshi Matsuoka. Learning to balance the inverted pendulum using neural networks. In *1991 IEEE International Joint Conference on Neural Networks, 1991*, pages 214–219, 1991.