



The One-Wheel Cubli: A 3D inverted pendulum that can balance with a single reaction wheel[☆]

Matthias Hofer^{a,*}, Michael Muehlebach^b, Raffaello D'Andrea^a

^a ETH Zürich, Institute for Dynamic Systems and Control, Sonneggstrasse 3, Zürich 8092, Switzerland

^b Max Planck Institute for Intelligent Systems, Learning and Dynamical Systems, Max Planck Ring 4, Tübingen 72076, Germany

ARTICLE INFO

Keywords:

3D inverted pendulum
Reaction wheel
Controllability
Silver ratio
Kalman filter
Linear-quadratic regulator
Delay compensation
Active damping

ABSTRACT

This article presents a novel 3D inverted pendulum that can balance on one of its corners using only a single reaction wheel. This is achieved by a careful design of the mass moment of inertia in such a way that the inertia along the two principal tilt axes is significantly different. The consequence is a time-scale separation of the underlying tilt dynamics, which renders the system controllable. We show that controllability is maximized when the ratio of the two principal inertia values amounts to the square of the silver ratio and discuss a sensor placement that minimizes variance in our tilt estimates. Both of these aspects lead to a principled design of the system. A model is derived from first-principles and used for delay compensation, state estimation, and to design a linear-quadratic regulator that stabilizes the highly underactuated system in its upright equilibrium. Thereby, the modeling and compensation of cantilever deflections, which arise from the lightweight design, is crucial. The article includes experimental results, which underline the efficacy of the system design and highlight an excellent balancing performance of the proposed feedback controller.

1. Introduction

The control of underactuated systems has a long history, with cars, boats and planes being prominent examples. In many cases, the degrees of freedom that are actuated indirectly are stable. More challenging, and arguably more interesting systems arise when these degrees of freedom are unstable, such as the cart pole system [1], the Furuta pendulum [2], the double inverted pendulum [3] and the planar reaction wheel pendulum [4] (see [5] and references therein for more examples). This article pushes the boundary of what is possible yet a step further by introducing a three-dimensional inverted pendulum system that is actuated with a single reaction wheel.

The system presented in this work is referred to as the One-Wheel Cubli and is depicted in Fig. 1. The name is inspired by its predecessor system, the Cubli that is based on three reaction wheels (see [6]). The One-Wheel Cubli has two main degrees of freedom, namely the roll and pitch direction (the yaw direction is neglected), but only one control input. Naturally, the question arises of how it is possible to simultaneously stabilize both degrees of freedom? The aim of this article is to answer this question and discuss the methodology applied to make the One-Wheel Cubli balance on its pivot.

Reaction wheel-based systems operate on the following principle: According to Newton's third law, a change in the momentum of the

reaction wheel leads to a change in the angular momentum of the object on which the reaction wheel is mounted. This means that reaction wheels provide a control input in the form of a reaction torque that can be used to actuate the angular velocity and attitude of an object.

In the case of the One-Wheel Cubli, this torque simultaneously affects both the roll and pitch directions. However, when applying a torque, the system reacts differently in the two degrees of freedom. Due to the considerably larger mass moment of inertia in the pitch direction, the response is slower compared to the roll direction that exhibits a smaller inertia. As a consequence of this time-scale separation of the associated dynamics, the system is controllable.

Reaction wheels are used to stabilize systems that operate in a plane, such as the classical reaction-wheel-based inverted pendulum with one degree of freedom (see [7]) and to control the attitude of 3D systems containing multiple reaction wheels (see [6]). Different configurations of the reaction wheel (or sphere) for 3D inverted pendulum systems are investigated in [8]. In general, the control of each rotational degree of freedom requires a reaction wheel. Sometimes, such as to control the attitude of satellites, an additional reaction wheel is included for redundancy reasons [9].

An important distinction has to be made between the working principle of reaction wheels and control moment gyroscopes that are

[☆] This paper was recommended for publication by Associate Editor Weichao Sun.

* Corresponding author.

E-mail addresses: hofermat@ethz.ch (M. Hofer), michael.muehlebach@tuebingen.mpg.de (M. Muehlebach), rdandrea@ethz.ch (R. D'Andrea).




Fig. 1. The One-Wheel Cubli balancing on its pivot using a single reaction wheel mounted inside the housing. Two masses are attached to the ends of a cantilever beam, resulting in a considerably higher rotational inertia in the pitch direction compared to the roll direction. As a consequence, the associated dynamics show a time-scale separation that makes the system controllable. Sensory feedback of the attitude and body rates is provided by inertial measurement units. The system is powered from an external power supply. A video showing the system during balancing can be found under: <https://youtu.be/mjvJ6-KoqFM>.

also used for attitude control of satellites (see [10]). Control moment gyroscopes rely on high rotational speeds causing gyroscopic forces that are used to stabilize the system. Reaction wheels on the other hand typically operate at low rotational speeds with gyroscopic effects being negligible.

Beside applications in the attitude control of satellites, we see an educational value in the One-Wheel Cubli platform. The inverted pendulum system represents a simple, yet challenging testbed for research and education in the field of control (see [11]). With the One-Wheel Cubli, we introduce a highly underactuated system that showcases the concept of controllability. For this reason, this article has a tutorial character to hopefully make the applied concepts accessible.

1.1. Contribution

The contributions of this paper are as follows:

We present the first 3D inverted pendulum that can be stabilized about its upright (unstable) equilibrium with a single reaction wheel. The design of the system is based on a careful controllability analysis: It is shown that a measure of controllability (which will be defined below) is maximized if the ratio of the natural frequencies of the tilt dynamics amounts to the silver ratio. In other words, by designing the system to have a significantly different inertia in the two principal directions, a time-scale separation follows, which makes the system controllable. The realization of the pendulum is presented and we highlight important aspects related to the lightweight structures and materials employed.

Sensory feedback is provided by five inertial measurement units (IMUs). We place the IMUs in such a way that the variance of the resulting tilt estimate is approximately minimized. This leads to a design where four IMUs are placed close to the pivot point and one IMU is placed far away. The tilt estimate obtained from the accelerometers is also fused with the readings of the gyroscope, which further improves the quality of the estimate.

A non-linear model of the system is derived and its linearization forms the basis of our estimation and control approach. The resulting feedback controller stabilizes the system in the upright position, is

robust to disturbances, and avoids exciting structural resonances that stem from the lightweight design. The model is also used to compensate for communication delays in a systematic way.

Finally, an experimental evaluation shows that the One-Wheel Cubli can reliably balance on its pivot point and reject disturbances that are applied manually. We also study the sensitivity of the closed-loop balancing performance with respect to inaccuracies in our first principles model. In particular, we show how the closed-loop performance is affected by a model parameter that is related to the cantilever stiffness and find that a parameter mismatch results in structural oscillations, which impair the control performance.

1.2. Outline

The remainder of this article is organized as follows: Section 2 presents a controllability analysis, which sets the stage for the design of the One-Wheel Cubli as presented in Section 3. The dynamics are introduced in Section 4, followed by the estimation and control design in Sections 5 and 6, respectively. The results are presented in Section 7 and a conclusion is drawn in Section 8.

Notation: Real numbers are denoted by \mathbb{R} and vectors are expressed as n-tuples (x_1, x_2, \dots, x_n) with dimension and stacking clear from the context. The body fixed coordinate system is denoted by A and ${}_A x$ refers to the coordinates of the vector x expressed in frame A . A rotation matrix from the inertial frame to the body frame is given by A_{AI} (see Appendix A for more details). The Euclidean norm is referred to as $|\cdot|$ and $\|\cdot\|$ denotes the L^2 -norm.

2. Controllability analysis

This section discusses the controllability of an inverted pendulum with two degrees of freedom that is actuated by a single reaction wheel. We will not only show that such a system is controllable about its unstable upright equilibrium, but also quantify the degree of controllability. We will derive the important result that controllability is maximized when (i) the ratio between the two natural frequencies of the two degrees of freedom corresponds to the *silver ratio*, that is $1 + \sqrt{2}$, and

(ii) the reaction wheel is aligned in a 45° angle with respect to the principal axes of inertia of the housing. This will guide the design of the system, which includes the placement of the reaction wheel and the mass distribution of the overall system in Section 3. A brief summary of the controllability analysis is provided here. For an in-depth analysis that includes detailed derivations, the reader is referred to [12].

As will be derived in Section 4, the dynamics of an inverted pendulum with two degrees of freedom and a single reaction wheel are given by

$$\ddot{\alpha}(t) = \pi_\alpha^2 \alpha(t) + \sigma \pi_\alpha^2 \cos(\eta) T_m(t) \quad (1)$$

$$\ddot{\beta}(t) = \pi_\beta^2 \beta(t) + \sigma \pi_\beta^2 \sin(\eta) T_m(t), \quad (2)$$

where $\alpha(t)$ and $\beta(t)$ denote the two inclinations about the two tilt axes. The constants $\pi_\alpha, \pi_\beta > 0$ are the natural frequencies of the two degrees of freedom, T_m stands for the input torque and σ for a related proportionality constant. The angle η describes how the reaction wheel is oriented in relation to the principal axes of the housing's inertia tensor. For symmetry reasons it is sufficient to consider $\eta \in [0, \pi/2]$. For $\eta = 0$ the wheel is oriented along the first principal axis, which is related to α , and for $\eta = \pi/2$ the wheel is oriented along the second principal axis, which is related to β . We conclude immediately that the system is uncontrollable if $\eta = 0$ or if $\eta = \pi/2$. We also conclude from the symmetry between (1) and (2) that the system is uncontrollable when $\pi_\alpha = \pi_\beta$.

For the following considerations, it will be convenient to introduce the system's state $\xi(t) := (\alpha(t), \dot{\alpha}(t), \beta(t), \dot{\beta}(t))$. We use the volume of

$$X := \bigcup_{\|T_m\|^2 \leq 1} \left\{ \xi(0) \in \mathbb{R}^4 \mid \xi(t) \text{ satisfies (1) and (2),} \right. \\ \left. \lim_{t \rightarrow \infty} \xi(t) = \lim_{t \rightarrow -\infty} \xi(t) = 0 \right\}, \quad (3)$$

to quantify the *degree* of controllability. The set X describes the region of the state space that can be reached with unit control energy provided that the system starts and returns to the origin.¹ This generalizes the classical concept of controllability Gramians to unstable systems, [13]. We also note that $\text{vol}(X) = 0$ if and only if the system is uncontrollable. An alternative approach to quantify controllability in a systematic manner is to use the closed-loop performance resulting from an optimal control design (e.g. H_∞ or H_2). While this is a perfectly viable approach, it relies on an additional disturbance model that characterizes specific sensor and actuator characteristics. Our metric, $\text{vol}(X)$, however, does not rely on such an additional disturbance model and captures the transient behavior of the dynamics. Other ad-hoc measures such as the minimum singular value of the controllability matrix are problematic, since they are not invariant to a linear coordinate transformation.

The perhaps surprising fact is that an expression for the volume $\text{vol}(X)$ can be computed in closed form. As is shown in [12] we obtain

$$\text{vol}(X) \sim \left(\sigma \sin(\eta) \cos(\eta) \frac{\pi_\alpha \pi_\beta}{4} \frac{\pi_\alpha - \pi_\beta}{\pi_\alpha + \pi_\beta} \right)^2. \quad (4)$$

Without loss of generality we assume that $0 < \pi_\beta \leq \pi_\alpha$ and decompose the above expression into the following four parts

$$\text{vol}(X) \sim \sigma^2 (\sin(\eta) \cos(\eta))^2 \left(\frac{\pi_\alpha}{4} \right)^2 \left(\epsilon \frac{1 - \epsilon}{1 + \epsilon} \right)^2, \quad (5)$$

where $\epsilon = \pi_\beta/\pi_\alpha$. The first part, σ^2 , increases for larger σ and the second part, $(\sin(\eta) \cos(\eta))^2$, is maximized for $\eta = \pi/4$. The third part increases for larger π_α and the fourth part captures the tradeoff between the two time constants π_α and π_β and is maximized for $\epsilon = \sqrt{2} - 1$ (the inverse of the silver ratio). The corresponding graph is shown in Fig. 2.




Fig. 2. The figure shows the function $\epsilon(1 - \epsilon)/(1 + \epsilon)$ for $\epsilon \in (0, 1]$, where $\epsilon = \pi_\beta/\pi_\alpha$ is the ratio between the two natural frequencies. The function is concave for $\epsilon \in (0, 1]$ and attains its maximum at $1/\delta_s$, where $\delta_s = 1 + \sqrt{2}$ denotes the silver ratio.

The reason for the tradeoff can be motivated by the following two competing mechanisms: (i) When considering the dynamics (1) in isolation it turns out that controllability increases (in the above sense) for larger π_α , since π_α not only affects the natural frequency, but also the way the torque T_m enters (1). (ii) Due to the coupling between (1) and (2), controllability (in the above sense) deteriorates when π_β approaches π_α .

We therefore conclude:

- (i) The parameter σ should be as large as possible.
- (ii) The orientation of the reaction wheel should be roughly $\eta = \pi/4$.
- (iii) The larger of the two natural frequencies (π_α) should be as large as possible. This is of course subject to other constraints, such as the bandwidth of the actuation and sensors, the structural properties of the housing (as discussed in Section 4), and the sampling time of the digital control system.
- (iv) Once the larger of the two natural frequencies is fixed, the other should be chosen such that a ratio of approximately $\sqrt{2} - 1$ is achieved.

The controllability analysis therefore has important implications for the design of the system. The latter will be the subject of the next section.

3. Design

This section presents the design of the One-Wheel Cubli and is divided into three parts; the first part discusses the mechanical system, while the second part presents the electrical and electronic systems. The question of how to optimally place the inertial measurement units (IMUs) on the One-Wheel Cubli is addressed in the final part of this section.

3.1. Mechanical design

The controllability analysis presented in the previous section forms the basis for the mechanical design of the One-Wheel Cubli. The system consists of the housing, a reaction wheel, an electric motor and a cantilever beam with two masses at its ends.

The ratio of the natural frequencies is given by the silver ratio and the orientation of the reaction wheel should be 45°. However, the maximization of σ and π_α does not directly translate to design specifications. If we assume that the cantilever beam is infinitely stiff and we neglect the mass and inertia of the reaction wheel, we can derive simplified expressions for the parameters σ and π_α (see Appendix B for the detailed expressions). Based on this simplifications, the following relations hold,

$$\sigma \sim \frac{1}{lm}, \quad \pi_\alpha^2 \sim \frac{1}{l}, \quad (6)$$

¹ Strictly speaking, we are evaluating the region of the state space that can be reached at time $t = 0$. However, since the dynamics are time-invariant, the same region can be reached at any $t \neq 0$.

where l denotes a characteristic length from the pivot to the location of the mass and m denotes a characteristic mass of the system. We conclude that σ is inversely proportional to both a characteristic length and mass of the system and π_α^2 is inversely proportional to a characteristic length of the system. Note that π_α is independent of the mass of the system, because both the rotational inertia and the torque that arises from gravity depend linearly on the mass, which consequently cancels. As a consequence, the requirements that σ and π_α should be maximized, translate to the system being as compact and lightweight as possible.

Based on these specifications we will discuss the design of the One-Wheel Cubli in detail: The housing is made of several plate-shaped carbon elements that are connected by 3D printed adapter parts. The center plate of the housing is made of aluminum to improve the heat dissipation of the motor. All plate-shaped parts are thinned out to reduce their weight. The size of the housing is minimized, but chosen to be sufficiently large to contain the reaction wheel, the electric motor, an embedded platform and multiple sensors (see Fig. 1).

The reaction wheel is attached to the electric motor that is fixed to a milled motor mount made from steel and attached to the center plate of the housing. The axis of rotation of the reaction wheel forms a 45° angle with the roll and pitch directions as required. The reaction wheel is aligned with the diagonal of the center plate and the electric motor is located next to the diagonal. Therefore, the center of mass of the overall system is slightly shifted in the horizontal plane towards the motor and is not vertically located above the pivot point. A calibration procedure to compensate for this effect is discussed in Section 6.

The inertia of the reaction wheel should be as high as possible to minimize the rotational speed of the motor during operation. Its weight and size should be as small as possible according to (6). Simulation studies, as discussed in Section 7.1, are carried out to investigate this trade off and identify a suitable dimensioning of the reaction wheel.

Once the housing and reaction wheel are designed, the additional inertia in the pitch direction (i.e. the β -direction) can be determined such that the ratio of the natural frequencies of the tilt dynamics is close to the silver ratio. Therefore, two end masses are attached with a cantilever to the housing. They significantly increase the inertia in the pitch direction, but only mildly increase the inertia in the roll direction (i.e. the α -direction). Since, only the ratio of the inertia is given, the weight and location of the end masses are not yet determined. There is a trade off between the overall weight of the system and its structural integrity. While light end masses that are placed at a great distance from the housing can reduce the overall weight of the system, the long cantilever beams might compromise the structural rigidity of the system. End masses of 0.3 kg and a total length of the cantilever of 1.25 m (measured from one end mass to the other) represent a good compromise between a low overall weight and sufficient structural rigidity. The resulting weight of the overall system is 1.937 kg. A supporting triangle is attached to the cantilever and housing to prevent bending of the cantilever in the direction of gravity.

3.2. Electrical and electronic design

The reaction wheel is actuated by a brushless electric motor (EC 60 flat 200 W from Maxon). The motor weighs 360 g, is operated at 48 V and can provide a continuous output torque of 0.5 N m and a peak torque of maximum 3.4 N m. The continuously available output torque is determined by the size of the motor, where an integrated ventilator increases the continuous operation range. Exceeding the continuous operation range is possible for short amounts of time, but causes the motor and particularly the windings to overheat, which could damage the motor. Increasing the motor size and consequently the continuously available torque, comes at the cost of a greater motor weight. Hence, there is a tradeoff between the available torque and the resulting weight of the motor.

The motor includes an integrated Hall sensor which provides the angular velocity of the motor. The motor is operated by a motor

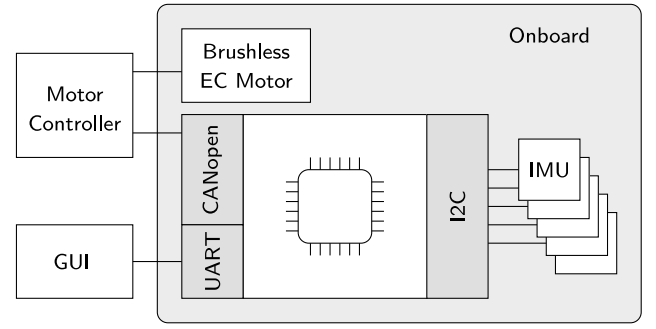


Fig. 3. A schematic overview of the electrical and electronic components of the system. The embedded platform including the microcontroller, the IMUs and the brushless EC motor are mounted onboard, while the motor controller is placed offboard. Onboard data is visualized with a graphical user interface (GUI) executed offboard on a laptop computer. Communication between the microcontroller and the motor controller is implemented via the CANopen protocol and the communication with the IMUs is provided by the I2C protocol. IMU and motor data are sent to the GUI via UART serial communication.

controller (EPOS4 50/15 from Maxon) and includes a current controller and tracks commanded torque setpoints. It also includes an overheat protection that limits the available torque when exceeding the continuous range to avoid any damage. The friction torque of the motor was observed to be negligible. A shunt regulator (DSR 70/30 from Maxon) is deployed to purge induced voltage when the motor is acting as a generator. A power supply (SM45-140 from Delta Elektronika) is used to power the motor controller.

The detailed electrical and electronic setup is shown in Fig. 3. A custom embedded platform comprising an STM32F407VGTx microcontroller meets the on-board computation requirement. Sensory feedback is provided by five 6-axis inertial measurement units (MPU-6050 from InvenSense).

3.3. IMU placement

The tilt estimation algorithm presented in [14] is used to estimate the attitude (except for yaw) of the One-Wheel Cubli from accelerometer data. It exploits the fact that there is a non-accelerated pivot point. It is convenient to represent the attitude using the gravity vector represented in the body fixed frame, ${}_A g$. All variables in the following derivation are expressed in the body frame A and the subscript is omitted for better readability. An estimate for g is given by the minimizer of the following optimization problem (see [14,15])

$$\arg \min_{g \in \mathbb{R}^3, \Phi \in \mathbb{R}^{3 \times 3}} \sum_{i=1}^M |m_i - \Phi \bar{r}_i + g|^2, \quad (7)$$

where m_i denotes the i th accelerometer measurement, \bar{r}_i denotes the position of the i th accelerometer with respect to the pivot, and M denotes the total number of IMUs. For the following derivation it is useful to express the position of each accelerometer by $\bar{r}_i = \bar{r} + r_i$ where r_i is such that $r_1 + \dots + r_M = 0$ and \bar{r} is the average position. The optimization problem (7) arises from the kinematics of a rigid body with a fixed pivot, whereby the matrix Φ contains angular and centripetal acceleration terms and is of no further interest here. The estimation via (7) can be interpreted within the maximum likelihood framework, see [15], and its solutions can be computed in closed form:

$$\hat{g} = - \sum_{i=1}^M \tau_i m_i, \quad \tau_i := \frac{1}{M} - r_i^\top \left(\sum_{j=1}^M r_j r_j^\top \right)^{-1} \bar{r}. \quad (8)$$

Hence, the estimate g is a weighted sum of the different sensor measurements. If we assume that each sensor is corrupted by isotropic

zero-mean noise with unit variance, whereby the noise is independent across the components of $\{m_1, \dots, m_M\}$, we obtain the following closed-form expression for the variance of \hat{g}

$$\text{Var}\{\hat{g}\} = \sum_{i=1}^M \tau_i^2 = \frac{1}{M} + \bar{r}^\top \left(\sum_{i=1}^M r_i r_i^\top \right)^{-1} \bar{r}, \quad (9)$$

where we have exploited the fact that $r_1 + \dots + r_M = 0$. The expression consists of two terms: The first term captures the $1/M$ reduction of the sensor noise due to the fact that there are M sensors. The second term arises from the compensation of angular and centripetal accelerations, which is implicit in (7) and strongly depends on the size of $|\bar{r}|$. For $\bar{r} = 0$ the compensation of angular and centripetal acceleration terms does not induce additional uncertainty. Expression (9) also points to the fact that an estimation via (7) is only possible if there are at least four sensors, which are not aligned in a plane, [14]; otherwise (9) is unbounded.²

We can now optimize the placement of the IMUs such that the variance in the attitude estimate is reduced. According to (9), we should choose a configuration where the “average”, \bar{r} , is small and the “variance”, $r_1 r_1^\top + \dots + r_M r_M^\top$, is large. We therefore decide to place four IMUs close to the pivot point and one IMU as far away as possible from the pivot. The four IMUs close to the pivot are aligned in a square. This not only simplifies manufacturing but symmetric placements along regular polygons should also be preferred over non-symmetric placements — an additional in-depth discussion of this fact can be found in [15]. Note that the variance of the gyroscope measurements is not affected by the placement of the IMUs.

4. Modeling

The modeling of the One-Wheel Cubli is presented in this section and results in a linear time-invariant state space model that forms the basis for the derivation of the estimation and control approaches. We consider the One-Wheel Cubli as the combination of four rigid bodies, namely the housing, the reaction wheel and the two end masses attached to cantilever beams. The system has four main degrees of freedom: These are the roll (α), pitch (β) and yaw (γ) angles and the rotational degree of freedom of the reaction wheel (φ). As a consequence of the lightweight design, the cantilever beams with the end mass introduce a limited, but unneglectable degree of freedom corresponding to a bending motion in the horizontal direction. These beam deflections are denoted by the variables, δ_1, δ_2 . The free-body diagram forms the basis of the model derivation and is shown in Fig. 4.

The nonlinear dynamics are derived using Lagrangian mechanics and presented in Appendix A. The result is a set of nonlinear equations of motion,

$$\begin{aligned} \dot{x} &= f(x, u), \quad u = T_m, \\ x &= (\alpha, \dot{\alpha}, \beta, \dot{\beta}, \gamma, \dot{\gamma}, \varphi, \dot{\varphi}, \delta_1, \dot{\delta}_1, \delta_2, \dot{\delta}_2) \in \mathbb{R}^{12}, \end{aligned} \quad (10)$$

where x denotes the state and u the control input, i.e. the motor torque acting on the reaction wheel.

The nonlinear model serves as the basis for the simulation environment as discussed in Section 7.1. For estimation and control purposes, the dynamics are linearized about $\bar{x} = 0$ and $\bar{u} = 0$ and a model reduction is performed to obtain a more compact state space representation. The motivation is twofold: Firstly, the dynamics are independent of φ for which reason it is eliminated. Secondly, the yaw dynamics are only affected by the cantilever deflection states. However, there is only an interaction between the yaw dynamics with the cantilever deflections if both end masses move in the same direction (in a clockwise or counterclockwise direction around the vertical axis)

² It is important to note that by construction, $r_1 + r_2 + \dots + r_M = 0$. Moreover, a slightly more elaborate algorithm exists, [15], that can estimate attitude even when $M = 3$.

for symmetry reasons. We assume that this interplay between the end masses and the yaw dynamics can be excluded. The validity of this assumption will be verified in the experimental investigation discussed in Section 7.2. Hence, the states related to the yaw dynamics (γ and $\dot{\gamma}$) are assumed to be zero for all times and excluded from the model. The linearized, reduced order model is given by

$$\dot{x} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 57.7 & 0 & 0.1 & 0 & 0 & * & 4.6 & * & -4.6 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 10.6 & 0 & 0 & 75.8 & 0 & -75.8 & 0 \\ -40.9 & 0 & -7.5 & 0 & 0 & * & -3.2 & * & 3.2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -4.2 & 0 & 0 & 0 & 0 & * & -19.8 & * & -12.6 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 4.2 & 0 & 0 & 0 & 0 & * & -12.6 & * & -19.8 \end{pmatrix} x + \begin{pmatrix} 0 \\ -20.7 \\ 0 \\ -2.3 \\ 1122 \\ 0 \\ 7.4 \\ 0 \\ -7.4 \end{pmatrix} u, \quad (11)$$

where certain (large) entries related to δ_1 and δ_2 are omitted (*) for better readability. Firstly, we see that if we neglect the coupling term between $\ddot{\alpha}$ and β in the second row (which is an order of magnitude smaller) and assume the states related to the cantilever deflections are zero (this is equivalent to an infinitely stiff cantilever), the dynamics are indeed of the forms (1) and (2). Secondly, we verify that

$$\epsilon = \frac{\pi_\beta}{\pi_\alpha} = \sqrt{\frac{10.6}{57.7}} = 0.43 \approx \frac{1}{\delta_s} = \sqrt{2} - 1. \quad (12)$$

Consequently, we can conclude that the targeted silver ratio of the natural frequencies of the two principal degrees of freedom as presented in Section 2 is approximately fulfilled. The lower natural frequency is 0.5 Hz and the higher natural frequency is 1.2 Hz.

The state and control input are normalized and the system is discretized with a sampling time of $T_s = 10$ ms,

$$\begin{aligned} x(k) &= Ax(k-1) + Bu(k) \\ x(k) &= (\alpha, \dot{\alpha}, \beta, \dot{\beta}, \varphi, \dot{\varphi}, \delta_1, \dot{\delta}_1, \delta_2, \dot{\delta}_2) \in \mathbb{R}^9 \\ u(k) &= u(T_s k), \end{aligned} \quad (13)$$

where $u(k)$ denotes the discrete-time control input at time index k obtained from a zero-order hold applied to the continuous-time control input. The time indexes of the entries of the state are omitted for better readability.

5. Estimation

Sensory feedback is provided by five IMUs each featuring a 3-axis accelerometer and a 3-axis gyroscope. The attitude of the One-Wheel Cubli can be estimated based on the accelerometer measurements and the angular velocity (body rate) from the gyroscope measurements. In the first part of this section, we present a calibration procedure to improve and fuse the raw sensor measurements. A complementarity filter to improve the signal to noise ratio of the attitude measurement is discussed in a second part. In a third step, we will introduce a Kalman filter for estimating the cantilever states, which are not directly measured. We will also use the Kalman filter to compensate for communication delays.

5.1. IMU calibration

The goal is to obtain an estimate of the attitude based on the raw accelerometer measurements and an estimate of the angular velocity based on the raw gyroscope measurements. This includes the compensation of bias and scale errors in the accelerometer reading and determining the orientation of the IMUs relative to the body fixed coordinate system A .




Fig. 4. The free-body diagram of the One-Wheel Cubli, where the two end masses are represented by the black rectangles: The Euler angles α, β, γ parameterize the attitude of the housing. The pivot point is denoted by O and the reaction wheel is attached to the housing in point P . The center of mass of the housing (without reaction wheel and end masses) is located at S and the cantilevers are connected to the housing at point Q . The axis of rotation of the reaction wheel is aligned with the e_z^D -axis that is obtained by a rotation of $\eta = \pi/4$ around the e_z^A -axis. The degree of freedom of the reaction wheel is denoted by φ . The vector with the motor torque acting on the reaction wheel is denoted by T^w and the reaction torque acting on the housing is denoted by $T^{h,T}$. Two end masses (of mass m_e) are located at the end points of the cantilever at E_1 and E_2 and are treated as point masses. The bending deflections of the cantilever beams are denoted by δ_1 and δ_2 . A torsional spring and damping coefficient is introduced to model the oscillatory behavior of the cantilever (identical for both cantilevers). Only the right cantilever is cut free, where T^{h,D_2} denotes the damping torque acting on the housing and T^{e_2} its reaction torque acting on the end mass e_2 . The inertia tensor of the housing (without end masses and reaction wheel) with respect to the pivot point O is denoted by $\bar{\Theta}_O^h$ and the mass of the housing is m_h . The mass of the reaction wheel is denoted by m_w and its inertia tensor with respect to P is denoted by $\bar{\Theta}_P^w$. The gravity vector points in negative e_z^l -direction.

The matrix consisting of the IMU weights τ_1, \dots, τ_M , see Section 3.3, Eq. (8), is given by

$$\Lambda = [I_3 \tau_1 \dots I_3 \tau_M] \in \mathbb{R}^{3 \times 15}. \quad (14)$$

Note that Λ depends on the IMU locations $\tilde{r}_1, \dots, \tilde{r}_M$. The raw accelerometer and gyro measurements of IMU i at time step k are referred to as $m_i(k) \in \mathbb{R}^3$ and $\omega_i(k) \in \mathbb{R}^3$, respectively. An estimate of the acceleration of the One-Wheel Cubli at time step k expressed in the body fixed coordinate system is given by

$${}_A \hat{g}(k) = \Lambda \begin{bmatrix} A_{AH_1} & & \\ & \ddots & \\ & & A_{AH_M} \end{bmatrix} \text{diag}(s_1, \dots, s_M) \begin{pmatrix} H_1 m_1(k) + b_1 \\ \vdots \\ H_M m_M(k) + b_M \end{pmatrix}, \quad (15)$$

where $b_i \in \mathbb{R}^3$ and $s_i \in \mathbb{R}^3$ denote the unknown accelerometer bias and scale error of IMU i and A_{AH_i} denotes the rotation matrix from the coordinate system of IMU i (denoted as H_i) to the body fixed coordinate system A . The bias of each accelerometer measurement is compensated and its scale error corrected. The local measurement is transformed to the body frame and fused according to Λ .

The angular velocity of the One-Wheel Cubli is obtained from

$${}_A \hat{\omega}(k) = \frac{1}{M} \sum_{i=1}^L A_{AH_i} H_i \omega_i(k), \quad (16)$$

where the measurement of each IMU is rotated into the body fixed coordinate system and combined with equal weight. The offsets in the gyroscope measurements change over time. Hence, they are identified during start up (by averaging the readings while the system is at rest) and compensated for during operation.

We identify the unknown accelerometer bias and scale error of each IMU, the rotation matrix for each IMU and the IMU locations in a first step. A procedure to address a possible drift of the accelerometer bias (see e.g. [16]) is discussed in Section 6. A supervised learning approach is applied with ground truth data of the acceleration and

angular velocity provided by means of a Vicon motion capture system. The acceleration reference is denoted as ${}_A g_{ref}$ and the angular velocity reference is denoted as ${}_A \omega_{ref}$. A data set of size N is collected with the One-Wheel Cubli being tilted in different directions and the pivot point being in contact with the ground. The raw IMU measurements and the motion capture reference are recorded.

For each IMU, the unknown accelerometer bias, b_i , and scale parameters, s_i , the rotation matrices from IMU frames to body fixed coordinate system, A_{AH_i} , and the IMU locations, \tilde{r}_i , affecting the fusion weights are determined by solving the following optimization problem

$$\arg \min_{\substack{b_i, s_i, \tilde{r}_i \in \mathbb{R}^3 \\ A_{AH_i} \in SO(3) \\ i=1, \dots, M}} \frac{1}{g_0} \sum_{k=1}^N |{}_A g_{ref}(k) - {}_A \hat{g}(k)| + \sum_{k=1}^N |\text{diag}(1, 1, 0)({}_A \omega_{ref}(k) - {}_A \hat{\omega}(k))|, \quad (17)$$

where the acceleration data is normalized by g_0 and only the x - y -components of the angular velocity are considered. The dependency on the variables being optimized is given by (15) and (16). A downhill-simplex-method is used to solve the optimization problem, where the rotation matrices are represented as rotation vectors for dealing with the constraint $A_{AH_i} \in SO(3)$. An initial guess for all optimization variables is required. The IMU locations and the rotation matrices can initially be estimated from the CAD model and all bias and scale parameters are set to zero and one, respectively.

5.2. Complementarity filter

A complementarity filter is used to reduce measurement noise of the attitude estimate, ${}_A \hat{g}$, by fusing it with angular velocity measurements, ${}_A \hat{\omega}$. The idea is simple: integrating the gyro measurements over time gives an estimate that shows little measurement noise, but introduces sensor drift. The fusion of the accelerometer and gyro-based tilt estimates therefore yields an improved estimate with no drift, and little variance.

Starting from our tilt estimate at time k , we integrate the attitude kinematics and incorporate the angular velocity, ${}_A\hat{\omega}(k)$, which yields a gyro-based attitude estimate for time $k + 1$. This estimate is then interpolated with the accelerometer-based estimate ${}_A\hat{g}(k + 1)$, see [17], and gives the final tilt estimate at time $k + 1$. The weight corresponding to the accelerometer estimate is set to 0.02 and the weight corresponding to the gyro estimate is set to 0.98. From the fused attitude estimate, the Euler angles α and β that are used in the following sections are computed. Similarly, $\dot{\alpha}$ and $\dot{\beta}$ are computed from ${}_A\hat{\omega}$. A detailed discussion of the complementarity filter algorithm can be found in [18].

5.3. Kalman filter

A Kalman filtering approach is discussed in this section. The motivation is twofold: On the one hand, the Kalman filter provides us with estimates for the cantilever deflections, which are not measured directly (see [19] for a detailed discussion). On the other hand, the Kalman filter can also be used to compensate measurement delays (see [20] for a detailed discussion).

Note that the complementarity filter discussed in the last section could in principle be incorporated in the Kalman filter. However, our approach has the advantage that each building block (Kalman filter and complementarity filter) can be tuned separately and adapted for its specific purpose.

The measurement delay is identified for both the IMU and the motor encoder and it corresponds to approximately one time step (10 ms). A compensation approach similar to the one presented in [21] is applied. Therefore, we introduce the augmented state, $\tilde{x} \in \mathbb{R}^{14}$, that includes the directly measured state variables delayed by one time step,

$$\tilde{x}(k) = (x(k), \alpha(k-1), \dot{\alpha}(k-1), \beta(k-1), \dot{\beta}(k-1), \dot{\varphi}(k-1)). \quad (18)$$

The measurement model can then be stated as,

$$z(k) = \underbrace{[0_{5 \times 9} \quad C]}_{\tilde{C}} \tilde{x}(k), \quad C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (19)$$

where $0_{5 \times 9}$ denotes a matrix of zeros with the size as specified. Similarly, the augmented \tilde{A} and \tilde{B} matrices are defined as,

$$\tilde{A} = \begin{bmatrix} A & 0_{9 \times 5} \\ I_5 & 0_{5 \times 9} \end{bmatrix} \in \mathbb{R}^{14 \times 14}, \quad \tilde{B} = \begin{bmatrix} B \\ 0_{5 \times 1} \end{bmatrix} \in \mathbb{R}^{14}, \quad (20)$$

where I_5 denotes the identity matrix of size 5.

We checked that the observability matrix of the augmented system has full rank, which confirms that the augmented system is observable. A steady-state Kalman filter is designed for the linear time-invariant system providing an estimate of the augmented state, i.e.

$$\hat{\tilde{x}}(k) = (I_{14} - L\tilde{C})\tilde{A}\hat{\tilde{x}}(k-1) + (I_{14} - L\tilde{C})\tilde{B}u(k-1) + Lz(k), \quad (21)$$

where $\hat{\tilde{x}}(k)$ denotes the state estimate at time k and L the steady state Kalman filter gain. The results of the IMU calibration procedure and the complementarity filter approach provide an initial guess for the measurement noise covariance. The identification of the process noise covariance is done in simulation and refinements of the tuning are done during the experimental evaluation as discussed in Section 7. The observability of the system and the positive definiteness of the process noise covariance (see Section 7.2) guarantee that the estimation error dynamics are stable (see [19]).

6. Control

The control approach of the One-Wheel Cubli is presented in this section. The stabilizability of the pair (\tilde{A}, \tilde{B}) can be verified by using the Popov-Belevitch-Hautus test. Note that the system is not controllable,

as the states related to the cantilever deflection can only be stabilized to zero, but not controlled to a nonzero value.

The One-Wheel Cubli is an underactuated system with one control input to stabilize fourteen states, out of which two states are unstable and one is marginally stable (five states are introduced by the state augmentation for the delay compensation). The linear-quadratic regulator (see [22] for a detailed discussion) is a suitable choice since it allows us to exploit the couplings in the dynamics through the linear description of the system. The control input is computed as a function of the state estimate given by the Kalman filter,

$$u(k) = -K \hat{x}(k), \quad (22)$$

where K results from solving the associated discrete time algebraic Riccati equation. Since the state includes the cantilever deflection angles and velocities, it permits us to penalize a potential excitation of cantilever oscillations. A small cost is assigned to the additional states $\alpha(k-1), \dot{\alpha}(k-1), \beta(k-1), \dot{\beta}(k-1), \dot{\varphi}(k-1)$ that form the augmented state. The detailed tuning of the controller is discussed in Section 7.

Note that since the system is stabilizable and both weighting matrices of the linear-quadratic regulator are positive definite (see Section 7.2), the resulting closed-loop system is guaranteed to be stable. As a consequence of the separation principle, the combination of the steady-state Kalman filter and the linear-quadratic regulator is also guaranteed to be stable (for more details see [22]).

As mentioned in Section 3.1, the center of mass is slightly off-center. Additionally, the effect of the cabling (e.g. for powering the motor) can further affect the center of mass and changes between experiments. Therefore, the center of mass is estimated by low pass filtering the angles (α, β) . The estimate of the center of mass $(\alpha_{\text{CoM}}, \beta_{\text{CoM}})$ is then subtracted from the estimated angles (α, β) . This strategy has the same effect as adding an integrator to the wheel velocity, but has the advantage of providing the estimates $\alpha_{\text{CoM}}, \beta_{\text{CoM}}$ which have a clear physical interpretation. As a consequence, we can guarantee that the wheel speed averages to zero when the One-Wheel Cubli balances in an upright position. Note that this procedure can also compensate for drifts of the accelerometer bias, which are not addressed by the calibration procedure discussed in Section 5.1.

7. Results

7.1. Simulation

A simulation environment is developed based on the nonlinear dynamical model presented in Section 4. The simulation environment also includes the thermal overheat protection of the motor, communication delays and the measurement noise. The Kalman filter and linear-quadratic regulator close the loop of the simulation environment.

The simulation environment is particularly important during the design phase of the One-Wheel Cubli, since it allows for a principled selection of the electric motor. The available motor torque depends on the weight of the motor that again affects the closed-loop behavior and consequently the required torque to stabilize the system. Hence, a simulation environment is indispensable to analyze these interactions. Furthermore, simulating the closed-loop behavior reveals to what extent the required torque exceeds the continuously available torque and if there are any overheating issues. The sizing of the reaction wheel is also investigated in simulation. Increasing the inertia of the wheel reduces the wheel speed, but also increases the weight of the system, causing a higher tilting torque. The optimal approach here is to design the reaction wheel to fully leverage the available space within the housing and place its mass as far from the center of rotation as possible. This gives the optimum of the reaction wheel's inertia compared to its mass.

Moreover, the simulation studies have revealed two important properties: On the one hand, addressing the potential cantilever oscillations with the control approach is crucial to stabilize the system. If neglected,




Fig. 5. (Left) The figure shows the rejection of a disturbance applied to the One-Wheel Cubli when balancing. The angles α and β are shown in the top two plots. The torque is shown in the middle plot, where the red envelope indicates the available torque as a function of the wheel speed. The two bottom plots show the speed of the reaction wheel and the thermal state of the motor. A disturbance in the positive β -direction is applied manually by dropping an object on one of the end masses. To stabilize the system, the controller first commands a negative torque, causing α to increase as well. Due to the higher natural frequency of the α -direction, the increase in α happens faster than in β . Subsequently, a positive torque is applied to recover both angles simultaneously and bring the system back to the upright equilibrium. The applied motor torque causes the wheel speed to first become slightly negative and then to become clearly positive, imposing an (inactive) constraint on the available motor torque. The I_2t overhear state of the motor increases as the applied torque exceeds the maximum continuous torque of 0.5 Nm. However, the increase is below 10% of the possible short term overheating of the motor. (Right) The rejected disturbance is shown in the α - β -plane, with the time being represented by the color. The diagonal gray lines indicate the level curves of the control input as a (linear) function of α and β only (all other states are assumed to be zero). The upper left region corresponds to a positive input torque and the spacing between two level curves is 1 Nm. The system primarily moves in the direction of the level curves, as a deviation from the level curves causes a large control input that regulates the system back towards the zero level curve. The depicted disturbance corresponds approximately to the maximum disturbance from which the system can recover. It gives an example of the angular range that can be stabilized.

the excited cantilever oscillations can grow without bounds and destabilize the system. On the other hand, a systematic compensation of the communication delays is important for reducing the required torque.

Finally, the simulation environment can be used to identify an initial guess for the tuning of the Kalman filter and the linear-quadratic regulator.

7.2. Experiments

The results of the experimental evaluation of the One-Wheel Cubli are presented in this section. Before performing balancing experiments with two degrees of freedom, a simplified experiment is conducted. By adding a wedge to the pivot point, one degree of freedom is locked, which simplifies the task of balancing. Nonetheless, this allows us to refine the tuning of the Kalman filter and the linear-quadratic regulator for each degree of freedom separately and also identify initial values for α_{CoM} and β_{CoM} in a separate experiment.

The following procedure is applied to identify the tuning of the linear-quadratic regulator: The weights related to the delayed states, the angular velocity of the housing, the wheel speed and the cantilever deflections are set to small values. The input weight is fixed and the weight related to the tilt angle of the housing (that is not locked) is increased until the system can successfully balance. If cantilever oscillations occur or the wheel speed is excessively high, the corresponding weights are increased. The same procedure is repeated for the other tilt angle.

After successfully stabilizing the pitch and roll directions separately, the wedge is removed and we consider the problem of controlling both axes at the same time. Thereby, the previously identified tuning parameters for the linear-quadratic regulator serve as initial guesses. They are refined for best performance and finally set to,

$$Q = \text{diag} \left(1, \frac{1}{10}, 1, \frac{1}{10}, 5, 1, 1, 1, 1, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10}, \frac{1}{10} \right)$$

$$R = 20, \quad (23)$$

with Q referring to the state cost and R to the input cost.

The following tuning is used for the Kalman filter,

$$\begin{aligned} V &= \text{diag} \left(\frac{1}{50}, \frac{1}{1000}, \frac{1}{50}, \frac{1}{500}, 2, \frac{1}{500}, \frac{1}{500}, \frac{1}{500}, \right. \\ &\quad \left. \frac{1}{500}, \frac{1}{100}, \frac{1}{100}, \frac{1}{100}, \frac{1}{100}, \frac{1}{100} \right) \\ W &= \text{diag} \left(\frac{1}{1000}, \frac{1}{1000}, \frac{1}{20000}, \frac{1}{20000}, \frac{1}{1000} \right), \end{aligned} \quad (24)$$

where V and W denote the process and measurement noise covariance, respectively.

The results of the One-Wheel Cubli reliably balancing on its pivot and rejecting manually applied disturbances are shown in Fig. 5. If left undisturbed, the One-Wheel Cubli can balance indefinitely (the longest experiment was stopped after more than an hour of uninterrupted balancing). The interested reader is referred to the video attachment of this article to gain an impression of the experiments (<https://youtu.be/mjvJ6-KoqFM>).

The behavior of the One-Wheel Cubli recovering from a disturbance reveals an inherent property of the underactuated system. From Fig. 4, we can see that a positive input torque causes the One-Wheel Cubli to move in negative α and negative β -directions. The gain matrix of the linear-quadratic regulator is,

$$K = \begin{pmatrix} -24.2 & -16.7 & 25.8 & 39.7 & 0.5 & -0.1 & -5.2 & 0.1 & 5.2 \end{pmatrix}, \quad (25)$$

where the entries related to the delayed states are all zero and not shown for the sake of brevity. Taking into account the sign convention as given by (22), we see that a positive deviation in α results in a positive input torque (driving α back to zero), while a positive deviation in β results in a negative input torque that causes both α and β to increase (as can be seen in Fig. 5). Due to the faster response of α (because of its higher natural frequency), the positive torque contribution of α eventually becomes dominant, driving both angles back




Fig. 6. The amplitude spectrum of the torque (top plot) and the α -direction (bottom plot) resulting from data collected during a balancing experiment. An enlarged version of the relevant frequency band is shown in the plots on the right hand side. The three differently colored curves correspond to different values of the natural frequency of the cantilever oscillation, f_0 , assumed by the controller. A correctly estimated value of $f_0 = 57.2$ Hz results in the red curve and no oscillations occurring in closed loop. The blue and green curves result from too low ($f_0 = 56.2$ Hz) and too high ($f_0 = 58.2$ Hz) values of f_0 and indicate considerably larger oscillations in the amplitude spectrum of both the input torque and α . Note that the frequencies at which the oscillations occur differ from the natural frequencies as we consider forced oscillations happening in closed loop. Apart from the structural oscillations, the relevant frequencies of the control input and the angle α during the balancing experiment lie below 5 Hz.

to zero. This behavior shows that to recover from a disturbance, the controller exploits the time-scale separation of the dynamics (i.e. the higher natural frequency in α -direction compared to the β -direction).

Next, we analyze the importance of systematically addressing the cantilever oscillations and their sensitivity to the associated model parameters. The dynamics of the cantilever deflections are primarily determined by the stiffness parameter k and the corresponding natural frequency (not to be confused with the natural frequencies of (1) or (2)). Three scenarios are compared: In the first scenario, the natural frequency of the cantilever oscillation assumed by the controller matches the true frequency, while in the second and third scenario, the natural frequency assumed by the controller is either too low or too high. The results of this analysis are shown in Fig. 6. The closed-loop behavior shows a high sensitivity towards the stiffness of the cantilever. A small deviation leads to a degradation of the control performance and emphasizes the importance of systematically addressing the flexibility of the cantilever as done with our control approach. Finally, we verify that the movements about yaw are small, which is in line with the assumption of Section 4 (and enables us to exclude the yaw dynamics from the model).

8. Conclusion

This article shows that a 3D inverted pendulum system can balance on its pivot with a single reaction wheel. The mass moments of inertia of the two principal degrees of freedom are chosen such that the ratio of the natural frequencies of the associated dynamics fulfills the silver ratio, or equivalently maximizes a controllability measure of the system. We derived a nonlinear model of the system, which captures structural modes that arise from the lightweight design. We observed that an accurate model of these structural modes is important for balancing in an upright position. Sensory feedback of the attitude and body rates is provided by multiple IMUs. These are placed on the system in such a way that the variance of our accelerometer-based tilt

estimate is minimal. A calibration approach for the IMU measurements is presented and compensates for alignment, scale and bias offsets in the gyroscope and accelerometer readings. We further introduced a modular filtering approach that is easily tunable, compensates for communication delays, and fully reconstructs the state of the system. The motor torque is computed with a linear-quadratic regulator that also compensates for structural vibration modes. We found that a precise identification of the natural frequency of the cantilever is important to avoid the excitation of structural oscillations.

Even though the controllability analysis is based on a simplified model of the system, it fully captures fundamental tradeoffs that arise from the coupling of the tilt dynamics. As such, the controllability analysis represents the basis for the mechanical design of the system.

The compromising effect of the potential cantilever oscillations is addressed purely in software. A state observer provides an estimate of these oscillations and they are included in the control design. Alternatively, the hardware could be improved by making the cantilever beam stiffer which would either remove the oscillations completely or at least move them to frequencies that are significantly above the bandwidth of the controller. However, such an approach is likely to increase the weight of the system, which is problematic due to the torque limits of the motor.

Furthermore, the effect of the cabling used for powering the motor and the embedded platform could be mitigated by integrating the motor controller plus a battery on board the One-Wheel Cubli. This would also increase the weight of the system, but make it self-contained. This will be investigated in future works.

With our article we hope to stimulate further exciting research with underactuated robotic systems. We also hope to highlight how powerful linear analysis tools are, and how they can be used to create new mechatronic systems with counter-intuitive properties.

Credit authorship contribution statement

Michael Muehlebach and Raffaello D'Andrea discovered that a 3D inverted pendulum system can be balanced with a single reaction wheel. They set out to build such a system, but failed twice. Matthias Hofer, with the help of Michael Muehlebach and Raffaello D'Andrea, did a careful and systematic redesign, which involved the analysis, modeling and compensation of structural vibrations, switching to a different motor and the compensation of communication delays. This led to the final design as discussed in the article. Matthias Hofer and Michael Muehlebach both wrote parts of the article, which was reviewed by Raffaello D'Andrea.

Data availability

Data will be made available on request.

Acknowledgments

The authors would like to thank Michael Egli and Matthias Müller for their significant contribution to the mechanical and electrical/electronic design of the One-Wheel Cubli. Further thank goes to Clemens Murer and the technical support team of Maxon Motors, the ETH Zurich workshop platform, Marcus Aaltonen, Atif Ali, Daniel Wagner and Marc-André Corzillius.

Michael Muehlebach receives funding from the German Research Foundation and the Branco Weiss Fellowship, administered by ETH Zurich, Switzerland.

Appendix A. Derivation of the dynamics

In this section, the detailed derivation of the non-linear dynamics of the One-Wheel Cubli using Lagrangian mechanics is presented. The derivation is based on the free-body diagram shown in Fig. 4.

First, the variables required to formulate the Lagrangian of the system are introduced. The four rigid bodies are denoted by lower case letters as h (housing), w (reaction wheel), e_1 (end mass 1) and e_2 (end mass 2). The pivot point is denoted by O and the reaction wheel is attached to the housing in P . The center of mass of the housing is denoted by S and assumed to lie between O and Q , where the cantilevers are attached to the housing. The end masses are located at E_1 and E_2 . The generalized coordinates of the system are,

$$q = (\alpha, \beta, \gamma, \varphi, \delta_1, \delta_2). \quad (26)$$

The following coordinate systems are used in the derivation:

$$I \xrightarrow{\alpha} C \xrightarrow{\beta} B \xrightarrow{\gamma} A \xrightarrow{\eta=c} D \xrightarrow{\varphi} G \quad (27)$$

$$A \xrightarrow{\delta_{1,2}} F_{1,2}. \quad (28)$$

The inertial coordinate system is denoted by I . The Euler angle coordinate frames are referred to as C (roll), B (pitch) and A (yaw). The coordinate frame D aligns the direction of rotation of the reaction wheel with the body fixed frame A . The coordinate frame G is attached to the reaction wheel and describes its rotational degree of freedom. The coordinate frames F_1 and F_2 are attached to the end masses and describe the cantilever deflections.

The Euler–Lagrange equation of the system is

$$\frac{d}{dt} \left(\frac{\partial T}{\partial \dot{q}} \right)^\top - \left(\frac{\partial T}{\partial q} \right)^\top - \left(\frac{\partial U}{\partial q} \right)^\top = f_{NP}, \quad (29)$$

where T stands for the kinetic energy, U for the potential energy and f_{NP} denotes the non-potential forces related to the motor torque and cantilever damping. Different position vectors are defined, where the left subscript denotes the coordinate system to which their coordinates refer. The right subscript (from left to right) defines from where to where the vector points,

$$A^r_{OP} = (0, 0, l_P), \quad A^r_{OS} = (0, 0, l_S) \quad (30)$$

$$A^r_{OQ} = (0, 0, l_Q), \quad F_1^r_{OE_1} = (-l_E, 0, l_Q) \quad (31)$$

$$F_2^r_{OE_2} = (l_E, 0, l_Q). \quad (32)$$

A number of rotation matrices is introduced to convert variables between the different coordinate systems,

$$\begin{aligned} A_{CI} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & \sin \alpha \\ 0 & -\sin \alpha & \cos \alpha \end{pmatrix}, & A_{BC} &= \begin{pmatrix} \cos \beta & 0 & -\sin \beta \\ 0 & 1 & 0 \\ \sin \beta & 0 & \cos \beta \end{pmatrix} \\ A_{AB} &= \begin{pmatrix} \cos \gamma & \sin \gamma & 0 \\ -\sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{pmatrix}, & A_{DA} &= \begin{pmatrix} \cos \eta & \sin \eta & 0 \\ -\sin \eta & \cos \eta & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ A_{GD} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \varphi & \sin \varphi \\ 0 & -\sin \varphi & \cos \varphi \end{pmatrix}, & A_{F_1A} &= \begin{pmatrix} \cos \delta_1 & \sin \delta_1 & 0 \\ -\sin \delta_1 & \cos \delta_1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ A_{F_2A} &= \begin{pmatrix} \cos \delta_2 & \sin \delta_2 & 0 \\ -\sin \delta_2 & \cos \delta_2 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \end{aligned} \quad (33)$$

where for example A_{CI} describes a rotation in the positive mathematical sense from I to C . We introduce the angular velocities,

$$I^{\omega}_{IC} = (\dot{\alpha}, 0, 0), \quad C^{\omega}_{CB} = (0, \dot{\beta}, 0) \quad (34)$$

$$B^{\omega}_{BA} = (0, 0, \dot{\gamma}), \quad D^{\omega}_{DG} = (\dot{\varphi}, 0, 0) \quad (35)$$

$$A^{\omega}_{AF_1} = (0, 0, \dot{\delta}_1), \quad A^{\omega}_{AF_2} = (0, 0, \dot{\delta}_2), \quad (36)$$

where for example I^{ω}_{IC} is the angular velocity obtained by rotating coordinate system C wrt. I . The velocities of P , E_1 and E_2 are obtained by applying the differentiation rule for vectors that are expressed in non-inertial coordinate systems, namely

$$A^v_P = A^{\omega}_{IA} \times A^r_{OP} \quad (37)$$

$$F_1^v_{E_1} = F_1^{\omega}_{IF_1} \times F_1^r_{OE_1} \quad (38)$$

$$F_2^v_{E_2} = F_2^{\omega}_{IF_2} \times F_2^r_{OE_2}. \quad (39)$$

Next, the angular velocities of the four rigid bodies can be expressed with the help of the previously introduced variables,

$$A^{\Omega}^h = A^{\omega}_{IA} = A^{\omega}_{IC} + A^{\omega}_{CB} + A^{\omega}_{BA} \quad (40)$$

$$= A_{AB} A_{BC} C^{\omega}_{IC} + A_{AB} B^{\omega}_{CB} + A^{\omega}_{BA} \quad (41)$$

$$D^{\Omega}^w = D^{\omega}_{IG} = D^{\omega}_{IA} + D^{\omega}_{DG} = A_{DA} A^{\omega}_{IA} + D^{\omega}_{DG} \quad (42)$$

$$A^{\Omega}^{E_1} = A^{\omega}_{IF_1} = A^{\omega}_{IA} + A^{\omega}_{AF_1} \quad (43)$$

$$A^{\Omega}^{E_2} = A^{\omega}_{IF_2} = A^{\omega}_{IA} + A^{\omega}_{AF_2}. \quad (44)$$

From these variables, the Jacobi matrices of rotation can be derived, namely

$$A^{\Omega}^h = A^{\bar{J}}^h_R \dot{q} + A^{\bar{v}}^h_R, \quad D^{\Omega}^w = D^{\bar{J}}^w_R \dot{q} + D^{\bar{v}}^w_R \quad (45)$$

$$A^{\Omega}^{E_1} = A^{\bar{J}}^{e_1}_R \dot{q} + A^{\bar{v}}^{e_1}_R, \quad A^{\Omega}^{E_2} = A^{\bar{J}}^{e_2}_R \dot{q} + A^{\bar{v}}^{e_2}_R, \quad (46)$$

where the \bar{v} account for terms that are independent of \dot{q} . The left subscript of $A^{\bar{J}}^h_R$ denotes the coordinate system in which the Jacobi matrix is expressed, the right subscript stands for rotation (as opposed to translation) and the superscript denotes the corresponding rigid body. With the help of the Jacobi matrices, the non-potential forces acting on the housing can be expressed as,

$$f_{NP}^{h,T} = A^{\bar{J}}^h_R{}^\top A_{AD} D^T T^{h,T} \quad (47)$$

$$f_{NP}^{h,D} = A^{\bar{J}}^h_R{}^\top (A^T T^{h,D_1} + A^T T^{h,D_2}), \quad (48)$$

where $f_{NP}^{h,T}$ includes the motor torque and $f_{NP}^{h,D}$ corresponds to the cantilever damping. The torques acting on the housing are given by

$$D^T T^{h,T} = (-T_m, 0, 0) \quad (49)$$

$$A^T T^{h,D_1} = (0, 0, d\dot{\delta}_1) \quad (50)$$

$$A^T T^{h,D_2} = (0, 0, d\dot{\delta}_2). \quad (51)$$

Similarly, for the reaction torque acting on the reaction wheel,

$$f_{NP}^{w,T} = D^{\bar{J}}^w_R{}^\top D^T T^w, \quad D^T T^w = (T_m, 0, 0), \quad (52)$$

and the reaction torques acting on both end masses,

$$f_{NP}^{e_1} = A^{\bar{J}}^{e_1}_R{}^\top A^T T^{e_1}, \quad A^T T^{e_1} = (0, 0, -d\dot{\delta}_1) \quad (53)$$

$$f_{NP}^{e_2} = A^{\bar{J}}^{e_2}_R{}^\top A^T T^{e_2}, \quad A^T T^{e_2} = (0, 0, -d\dot{\delta}_2). \quad (54)$$

Finally, the non-potential forces are combined to,

$$f_{NP} = f_{NP}^{h,T} + f_{NP}^{w,T} + f_{NP}^{h,D} + f_{NP}^{e_1} + f_{NP}^{e_2}. \quad (55)$$

With $A^{\bar{\Theta}}^h_O = \text{diag}(I_{h,x}, I_{h,y}, I_{h,z})$ denoting the housing inertia tensor wrt. O and expressed in frame A , and $D^{\bar{\Theta}}^w_P = \text{diag}(I_{w,x}, I_{w,y}, I_{w,z})$ referring to the reaction wheel inertia tensor wrt. P and expressed in frame D , the kinetic energy of each rigid body can be written as,

$$T^h = \frac{1}{2} A^{\Omega}^h{}^\top A^{\bar{\Theta}}^h_O A^{\Omega}^h \quad (56)$$

$$T^w = \frac{1}{2} m_w A^v_P{}^\top A^v_P + \frac{1}{2} D^{\Omega}^w{}^\top D^{\bar{\Theta}}^w_P D^{\Omega}^w \quad (57)$$

$$T^{e_1} = \frac{1}{2} m_{e_1} F_1^v_{E_1}{}^\top F_1^v_{E_1} \quad (58)$$

$$T^{e_2} = \frac{1}{2} m_{e_2} F_2^v_{E_2}{}^\top F_2^v_{E_2}, \quad (59)$$

Table A.1

The physical parameters of the One-Wheel Cubli model. The denomination of the parameter, its meaning, value, unit and identification method are summarized from left to right.

Parameter	Meaning	Value	Unit	Identification
l_P	Length from O to P	0.133	m	CAD
l_S	Length from O to S	0.155	m	CAD
l_Q	Length from O to Q	0.214	m	CAD
l_E	Length from Q to E_1, E_2	0.5975	m	CAD
η	Angle between e_x^A and e_x^D axes	$\pi/4$	rad	CAD
m_h	Mass of housing	1.101	kg	CAD
m_w	Mass of reaction wheel	0.228	kg	CAD
m_e	Mass of one end mass	0.304	kg	CAD
$I_{h,x}$	Housing inertia in e_x^A -direction wrt. O	2.992×10^{-2}	$\text{m}^2 \text{kg}$	CAD
$I_{h,y}$	Housing inertia in e_y^A -direction wrt. O	5.836×10^{-2}	$\text{m}^2 \text{kg}$	CAD
$I_{h,z}$	Housing inertia in e_z^A -direction wrt. O	3.023×10^{-2}	$\text{m}^2 \text{kg}$	CAD
$I_{w,x}$	Reaction wheel inertia ^a in e_x^D -direction wrt. P	9.044×10^{-4}	$\text{m}^2 \text{kg}$	CAD
$I_{w,y}$	Reaction wheel inertia in e_y^D -direction wrt. P	4.117×10^{-4}	$\text{m}^2 \text{kg}$	CAD
$I_{w,z}$	Reaction wheel inertia in e_z^D -direction wrt. P	4.117×10^{-4}	$\text{m}^2 \text{kg}$	CAD
k	Spring constant	1.080×10^4	N m rad^{-1}	First-principles
d	Damping constant	0.435	N m s rad^{-1}	Tuning
g_0	Gravitational acceleration	9.81	m s^{-2}	–

^aThis includes the rotational inertia of the motor.

where the end masses are modeled as point masses and off diagonal mass moments of inertia are neglected. The combined kinetic energy of the system is

$$T = T^h + T^w + T^{e_1} + T^{e_2}. \quad (60)$$

The potential energies of the four rigid bodies are comprised of the gravitational potential and the spring potentials for the cantilever beams, i.e.

$$U^h = -m_h A^T r_{OS}^T A g \quad (61)$$

$$U^w = -m_w A^T r_{OP}^T A g \quad (62)$$

$$U^{e_1} = -m_{e_1} F_1^T r_{OE_1}^T F_1 g + \frac{1}{2} k \delta_1^2 \quad (63)$$

$$U^{e_2} = -m_{e_2} F_2^T r_{OE_2}^T F_2 g + \frac{1}{2} k \delta_2^2. \quad (64)$$

Thereby, the gravity vector points in the negative z -direction when expressed in the inertial coordinate system,

$$l g = (0, 0, -g_0). \quad (65)$$

The combined potential energy of the system is,

$$U = U^h + U^w + U^{e_1} + U^{e_2}. \quad (66)$$

Finally, we can formulate the Euler–Lagrange equation of the One-Wheel Cubli that provides the six equations of motion for the system,

$$\frac{d}{dt} \left(\frac{\partial T}{\partial \dot{q}} \right)^T - \left(\frac{\partial T}{\partial q} \right)^T - \left(\frac{\partial U}{\partial q} \right)^T = f_{NP}. \quad (67)$$

The equations are derived using a symbolic toolbox. All parameters used in the above derivation are summarized in [Table A.1](#). All parameters related to mass, mass moment of inertia, and length are obtained from the CAD files. The identification of the parameters describing the cantilever oscillation is discussed in [Appendix C](#).

The structure of (67) is known to be,

$$M(q)\ddot{q} - h(q, \dot{q}, T_m) = 0, \quad (68)$$

which can be brought to a state space form,

$$\frac{d}{dt} \begin{pmatrix} q \\ \dot{q} \end{pmatrix} = \begin{pmatrix} \dot{q} \\ M^{-1}(q)h(q, \dot{q}, T_m) \end{pmatrix}. \quad (69)$$

Adjusting the ordering of the components in (69), finally gives (10).

Appendix B. Simplified dynamics for an infinitely stiff cantilever

Under the assumption of the cantilever being infinitely stiff and by neglecting the weight and inertia of the reaction wheel, the parameters of (1) and (2), simplify to

$$\pi_\alpha^2 = g \frac{2l_Q m_e + l_S m_h}{2l_Q^2 m_e + I_{h,x}} \quad (70)$$

$$\pi_\beta^2 = g \frac{2l_Q m_e + l_S m_h}{2l_E^2 m_e + 2l_Q^2 m_e + I_{h,y}} \quad (71)$$

$$\sigma = -\frac{1}{g(2l_Q m_e + l_S m_h)}. \quad (72)$$

All parameters are summarized in [Table A.1](#). In the case of a flexible cantilever, the expression for π_α changes because the intermediate deflection states, $(\delta_1, \delta_1', \delta_2, \delta_2')$ decouple the inertia of the end masses from the α -direction. The expression for σ also changes and becomes different for the α and β -direction for similar reasons.

Appendix C. Cantilever oscillations

The identification of the (torsional) spring constant characterizing the cantilever oscillations is discussed in this section. The cantilever is modeled as an Euler–Bernoulli beam that is fixed at one end and has a mass attached to the other end (see [23]). Thereby, the supporting triangle between cantilever and housing is neglected. The geometry of the hollow tube profile is defined by the inner diameter (2.2×10^{-2} m) and the outer diameter (2.4×10^{-2} m) resulting in a cross-sectional area, $A = 7.225 \times 10^{-5} \text{ m}^2$, and an area moment of inertia, $J = 4.787 \times 10^{-9} \text{ m}^4$. The density of the carbon tube is $\rho = 1.560 \times 10^3 \text{ kg/m}^3$ and the Young's modulus $E = 2.50 \times 10^2 \text{ GPa}$. The eigenfrequencies are given by the characteristic equation (see [23, p. 248]),

$$1 + \cosh(\kappa \bar{l}) \cos(\kappa \bar{l}) + \lambda \kappa \bar{l} (\sinh(\kappa \bar{l}) \cos(\kappa \bar{l}) - \cosh(\kappa \bar{l}) \sin(\kappa \bar{l})) = 0, \quad (73)$$

where $\lambda = m_e / \rho A \bar{l} = 5.16$ defines the ratio of the end mass to the cantilever mass. The length of the cantilever that is unsupported by the housing is $\bar{l} = 0.523 \text{ m}$ (note that this length is smaller than l_E as used in the dynamics derivation). Numerically solving the characteristic equation for the lowest eigenmode gives $\kappa_0 = 1.65 \text{ m}^{-1}$. This gives an initial guess for the lowest eigenfrequency of the cantilever,

$$f_0 = \frac{\kappa_0^2}{2\pi} \sqrt{\frac{EJ}{\rho A}} = 44.8 \text{ Hz}. \quad (74)$$

The torsional spring constant is then given as,

$$k = m_e \bar{l}^2 (2\pi f_0)^2. \quad (75)$$

As a consequence of the neglected supporting triangle, the true natural frequency of the system, and consequently the torsional stiffness, is likely to be higher. The natural frequency is fine-tuned during experiments performed on the real system. The excitation of oscillations is minimized for $f_0 = 57.2$ Hz with a corresponding torsional stiffness of $k = 1.08 \times 10^4$ N m rad⁻¹.

The sensitivity of the closed-loop behavior towards the normalized damping parameter of the cantilever is considerably smaller than its stiffness parameter. It is tuned for best performance in closed-loop experiments, where a value of $\delta = 2.6$ s⁻¹ is identified. The corresponding torsional damping parameter is

$$d = 2m_e \bar{l}^2 \delta = 0.43 \text{ N m s rad}^{-1}. \quad (76)$$

References

- [1] Chung Chung Choo, Hauser John. Nonlinear control of a swinging pendulum. *Automatica* 1995;31(6):851–62. [http://dx.doi.org/10.1016/0005-1098\(94\)00148-C](http://dx.doi.org/10.1016/0005-1098(94)00148-C).
- [2] Furuta K, Yamakita M, Kobayashi S, Nishimura M. A new inverted pendulum apparatus for education. *IFAC Proceedings Volumes* 1991;25(12):133–8. [http://dx.doi.org/10.1016/S1474-6670\(17\)50102-0](http://dx.doi.org/10.1016/S1474-6670(17)50102-0).
- [3] Furuta Katsuhisa, Kajiwara Hiroyuki, Kosuge Kazuhiro. Digital control of a double inverted pendulum on an inclined rail. *International Journal of Control* 1980;32(5):907–24. <http://dx.doi.org/10.1080/00207178008922898>.
- [4] Spong Mark W, Corke Peter, Lozano Rogelio. Nonlinear control of the reaction wheel pendulum. *Automatica* 2001;37(11):1845–51. [http://dx.doi.org/10.1016/S0005-1098\(01\)00145-5](http://dx.doi.org/10.1016/S0005-1098(01)00145-5).
- [5] Shen Jinglai, Sanyal AK, Chaturvedi NA, Bernstein D, McClamroch H. Dynamics and control of a 3D pendulum. In: 2004 43rd IEEE conference on decision and control (CDC) (IEEE cat. no. 04CH37601). vol. 1, 2004, p. 323–8. <http://dx.doi.org/10.1109/CDC.2004.1428650>.
- [6] Muehlebach Michael, D'Andrea Raffaello. Nonlinear analysis and control of a reaction-wheel-based 3-D inverted pendulum. *IEEE Transactions on Control Systems Technology* 2017;25(1):235–46. <http://dx.doi.org/10.1109/TCST.2016.2549266>.
- [7] Block Daniel J, Astrom Karl J, Spong Mark W. *The reaction wheel pendulum*. Morgan & Claypool; 2007.
- [8] Mayr Johannes, Spanlang Franz, Gatringer Hubert. Mechatronic design of a self-balancing three-dimensional inertia wheel pendulum. *Mechatronics* 2015;30:1–10. <http://dx.doi.org/10.1016/j.mechatronics.2015.04.019>.
- [9] Ismail Zuliana, Varatharajoo Renuganth. A study of reaction wheel configurations for a 3-axis satellite attitude control. *Advances in Space Research* 2010;45(6):750–9. <http://dx.doi.org/10.1016/j.asr.2009.11.004>.
- [10] Lappas VJ, Steyn WH, Underwood CI. Attitude control for small satellites using control moment gyros. *Acta Astronautica* 2002;51(1):101–11. [http://dx.doi.org/10.1016/S0094-5765\(02\)00089-9](http://dx.doi.org/10.1016/S0094-5765(02)00089-9).
- [11] Boubaker Olfa. The inverted pendulum: A fundamental benchmark in control theory and robotics. In: *International conference on education and E-learning innovations*. 2012, p. 1–6. <http://dx.doi.org/10.1109/ICEELI.2012.6360606>.
- [12] Muehlebach Michael. The silver ratio and its relation to controllability. 2019, [arXiv:1908.07109](https://arxiv.org/abs/1908.07109).
- [13] Zhou Kemin, Salomon Gregory, Wu Eva. Balanced realization and model reduction for unstable systems. *International Journal of Robust and Nonlinear Control* 1999;9(3):183–98. [http://dx.doi.org/10.1002/\(SICI\)1099-1239\(199903\)9:3<183::AID-RNC399>3.0.CO;2-E](http://dx.doi.org/10.1002/(SICI)1099-1239(199903)9:3<183::AID-RNC399>3.0.CO;2-E).
- [14] Trimpe Sebastian, D'Andrea Raffaello. Accelerometer-based tilt estimation of a rigid body with only rotational degrees of freedom. In: *Proceedings of the international conference on robotics and automation*. 2010, p. 2630–6. <http://dx.doi.org/10.1109/ROBOT.2010.5509756>.
- [15] Muehlebach Michael, D'Andrea Raffaello. Accelerometer-based tilt determination for rigid bodies with a nonaccelerated pivot point. *IEEE Transactions on Control Systems Technology* 2018;26(6):2106–20. <http://dx.doi.org/10.1109/TCST.2017.2753171>.
- [16] Gulmammadov Farid. Analysis, modeling and compensation of bias drift in MEMS inertial sensors. In: *2009 4th International Conference on Recent Advances in Space Technologies*. 2009, p. 591–6. <http://dx.doi.org/10.1109/RAST.2009.5158260>.
- [17] Dam Erik B, Koch Martin, Lillholm Martin. *Quaternions, interpolation and animation. Technical report*, 1998.
- [18] Gui Pengfei, Tang Liqiong, Mukhopadhyay Subhas. MEMS based IMU for tilting measurement: Comparison of complementary and kalman filter based data fusion. In: *IEEE Conference on Industrial Electronics and Applications (ICIEA)*. 2015, p. 2004–9. <http://dx.doi.org/10.1109/ICIEA.2015.7334442>.
- [19] Simon Dan. The discrete-time Kalman filter. In: *Optimal state estimation*. John Wiley & Sons, Ltd; 2006, p. 121–48. <http://dx.doi.org/10.1002/0470045345>, [Chapter 5].
- [20] Richard Jean-Pierre. Time-delay systems: an overview of some recent advances and open problems. *Automatica* 2003;39(10):1667–94. [http://dx.doi.org/10.1016/S0005-1098\(03\)00167-5](http://dx.doi.org/10.1016/S0005-1098(03)00167-5).
- [21] Sirouspour Shahin, Shahdi Ali. Discrete-time linear quadratic gaussian control for teleoperation under communication time delay. *The International Journal of Robotics Research* 2006;25(2):187–202. <http://dx.doi.org/10.1177/0278364906061081>.
- [22] Anderson Brian DO, Moore John B. *Optimal control: Linear quadratic methods*. Courier Corporation; 2007.
- [23] Gross Dietmar, Hauger Werner, Schnell † Walter, Wriggers Peter. Schwingungen kontinuierlicher systeme. In: *Technische mechanik 4: hydromechanik, elemente der höheren mechanik, numerische methoden*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2014, p. 213–72. http://dx.doi.org/10.1007/978-3-642-41000-0_4.



Matthias Hofer received the B.Sc. degree in Mechanical Engineering and M.Sc. degree in Robotics, Systems and Control from ETH Zurich in 2013 and 2015, respectively. He was awarded the ETH medal for his Master thesis on approximate Model Predictive Control of an aerial vehicle. He is a doctoral candidate at the Institute for Dynamic Systems and Control at ETH Zürich under the supervision of Prof. R. D'Andrea. His main research interests include the design and control of robotic and mechatronic systems, machine learning techniques and vision-based sensing approaches.



Michael Muehlebach studied mechanical engineering at ETH Zürich and specialized in robotics, systems, and control during his Master's degree. He received the B.Sc. and the M.Sc. in 2010 and 2013, respectively, before joining the Institute for Dynamic Systems and Control for his Ph.D. He graduated under the supervision of Prof. R. D'Andrea in 2018 and joined the group of Prof. Michael I. Jordan at the University of California, Berkeley as a postdoctoral researcher. In 2021 he started as an independent group leader at the Max Planck Institute for Intelligent Systems.

Michael received the Outstanding D-MAVT Bachelor Award for his Bachelor's degree and the Willi-Studer prize for the best Master's degree. His Ph.D. thesis was awarded with the ETH Medal and the HILTI prize for innovative research. He was also awarded a Branco Weiss Fellowship and an Emmy Noether Fellowship, which funds his research group.



Raffaello D'Andrea is a professor at ETH Zürich and the founder/CEO of Verity, a self-flying drone systems company and provider of drone failsafe technology. He cofounded Kiva Systems, which was acquired by Amazon and rebranded as Amazon Robotics. During his time as a professor at Cornell University, he cofounded the systems engineering program and led the Cornell Robot Soccer team to four world championships. He is also a founder of RoboGlobal, an organization that introduced the world's first Robotics and AI Exchange Traded Fund.

His work as a new media artist has been exhibited at the Venice Biennale and is part of the permanent collections of the National Gallery of Canada and France's FRAC Centre. He is the co-creator of the Robotic Chair, a plain-looking wooden chair that falls apart and seems to defy the second law of thermodynamics by reassembling itself. He created the drone design and choreographies for Cirque du Soleil's *Paramour* on Broadway and Metallica's *WorldWired* tour, and was executive producer for the drone light shows in Drake's *Aubrey & the Three Migos* tour, Céline Dion's *Courage* tour, and Justin Bieber's *Justice* tour.

In 2020 he was inducted into the National Inventors Hall of Fame and was elected to the U.S. National Academy of Engineering. His TED and research videos, with tens of millions of views, offer an inspiring view into the world of engineering, robotics, and computer science.