Mechatronics 24 (2014) 954–965

Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

A frequency domain iterative learning algorithm for high-performance, periodic quadrocopter maneuvers

Markus Hehn *, Raffaello D'Andrea

Institute for Dynamic Systems and Control, ETH Zurich, Switzerland

ARTICLE INFO

Article history:
Received 2 September 2013
Revised 13 September 2014
Accepted 28 September 2014
Available online 23 October 2014

Keywords:
Learning control
Quadrocopter control
Frequency domain learning
Repetitive control
Unmanned aerial vehicle control

ABSTRACT

Quadrocopters offer an attractive platform for aerial robotic applications due to, amongst others, their hovering capability and large dynamic potential. Their high-speed flight dynamics are complex, however, and the modeling thereof has proven difficult. Control algorithms typically rely on simplified models, with feedback corrections compensating for unmodeled effects. This can lead to significant tracking errors during high-performance flight, and repeated execution typically leads to a large part of the tracking errors being repeated. This paper introduces an iterative learning scheme that non-causally compensates repeatable trajectory tracking errors during the repeated execution of periodic flight maneuvers. An underlying feedback control loop is leveraged by using its set point as a learning input, increasing repeatability and simplifying the dynamics considered in the learning algorithm. The learning is carried out in the frequency domain, and is based on a Fourier series decomposition of the input and output signals. The resulting algorithm requires little computational power and memory, and its convergence properties under process and measurement noise are shown. Furthermore, a time scaling method allows the transfer of learnt maneuvers to different execution speeds through a prediction of the disturbance change. This allows the initial learning to occur at reduced speeds, and thereby extends the applicability of the algorithm for high-performance maneuvers. The presented methods are validated in experiments, with a quadrocopter flying a figure-eight maneuver at high speed. The experimental results highlight the effectiveness of the approach, with the tracking errors after learning being similar in magnitude to the repeatability of the system.

1. Introduction

Aerial robots serve as platforms for robotic applications that provide numerous benefits, including the ability to move freely in three-dimensional space, and the significantly increased ability to overcome obstacles due to not being limited to motion on the ground. For relatively small platforms that require hovering capabilities, multi-rotor vehicles such as quadrocopters are often the vehicle of choice [1]. Compared to other such platforms, quadrocopters profit from high mechanical robustness due to a minimal number of moving parts [2], safety due to comparatively small rotor size, and high thrust-to-weight ratios allowing high-performance maneuvers as well as the transport of large payloads.

While the use of quadrocopters as robotic platforms was largely confined to research institutions in the past, a growing number of industrial applications are now in the process of being developed and deployed. Examples include aerial imaging for photogrammetry, motion picture production, and journalism [3], environmental monitoring and inspection tasks of hard-to-reach objects such as pipelines, dams, and power lines [4], the creation of ad hoc antenna networks or arrays [5], as well as disaster coordination [6].

The capability of quadrocopters to perform highly dynamic, complex, and precise motions has been demonstrated repeatedly in recent years (see, for example, Mellinger et al. [7], Michael et al. [8], Muller et al. [9], Ritz et al. [10]). In order to execute such high-performance motions, the commonly used approach consists of using a first-principles model of the quadrotor dynamics to design the nominal maneuver, and a model-based feedback control law to ensure tracking of the nominal trajectory.

Such traditional feedback controllers however have important limitations in high-performance quadrotor applications. While the first-principles models used to design the controllers capture the near-hover behavior of quadrocopters well, secondary effects become increasingly important when maneuvering speed increases. Examples of such effects are the complex drag and lift behavior of rotary wings under unsteady inflow conditions [11],

* Corresponding author.
E-mail address: hehnm@alumni.ethz.ch (M. Hehn).

http://dx.doi.org/10.1016/j.mechatronics.2014.09.013
0957-4158/

the aerodynamic effects of a vehicle moving through the turbulent wake of its propellers [12], and external influences such as wind or ground and wall effects when operating in proximity to the environment [13]. Such effects are not typically accounted for in the maneuver and controller synthesis stage in order to make the design process tractable. The execution then heavily relies on the feedback controller to compensate for potentially significant effects not captured by the nominal dynamics.

In order to improve the tracking performance of quadrocopters under feedback control, a number of researchers have proposed learning schemes. Examples of such schemes include those based on reinforcement learning techniques [14,15] and neural networks [16,17], which are designed to automatically find well-performing control policies, and adaptive control methods [18–20] that adapt parameters that are based on modeled disturbances such as payloads, center of mass shifts and external disturbances.

When a motion is to be executed repeatedly, a further opportunity to improve tracking performance may arise: Many of the disturbances that degrade tracking performance will be similar each time the vehicle performs the motion. These disturbances can then be compensated for non-causally using data from past executions. Control strategies that exploit available data from past executions in order to improve tracking performance were first proposed in the late 1970s and early 1980s [21,22] for applications in motion control and power supply control. Since then, active research in this field, covering numerous applications and problem formulations (see e.g. Wang et al. [23], Bristow et al. [24], Cuiyan et al. [25], and references therein), has shown it to be a powerful approach for high-performance reference tracking. In extensions to these learning methods, several authors have shown the application of learning control methods to systems with underlying feedback control loops (e.g. [26,27]). In such scenarios, the powerful capability of learning control to non-causally compensate repeatable disturbances is combined with real-time feedback control to correct for non-repetitive noise.

While the application of learning algorithms, and specifically non-causal strategies, to stationary systems (such as rotating machinery and robotic arms [23]) is well-established, its use for the compensation of complex aerodynamic effects in flying vehicles is less mature and has been actively researched during recent years. Several high-performance maneuvers for multi-rotor vehicles have been demonstrated with the use of learning algorithms. Broadly speaking, the learning approaches used can be categorized in two groups:

The first group is characterized by its ability to learn motions that are parameterized. The motion is thus described by a (finite) set of design parameters, chosen by the user. After the execution of the motion, these parameters are adapted to compensate for disturbances. The direction and magnitude of the correction may be model-based, or based on the user's intuition. A discussion on the importance of choosing 'good' design parameters may be found in Lupashin and D'Andrea [28], where a learning algorithm for this kind of parameterized motions is demonstrated for multiple flips and fast translations with quadrocopters. A further demonstration of this class of learning algorithms is provided in Mellinger et al. [29]. The ability to shape the tracking performance strongly depends on the number of parameters that are optimized; in the above examples, the objective is to minimize the error at specific time instants ('key frames'), and a relatively small number of parameters is sufficient to do so. This makes the methods computationally lightweight.

The second group of learning approaches considers more generic motions that need not be specified by parameters. The system dynamics are considered in discrete time, and the correction consists of correction values (typically control inputs or set points) for each discrete time step. After execution of the motion, a numerical optimization over the correction values is performed in order to minimize a metric related to the tracking error. In this optimization, a model of the system dynamics provides the mapping from corrections to the tracking error. This approach is commonly known as a form of iterative learning control [24], and its application to high-performance quadrocopter flight has been demonstrated [30–33].

The delimitation between the two groups is not strict. Indeed, the second group of learning approaches could be seen as using a very large number of values to parameterize the correction.

The algorithm presented in this paper can be characterized to be a form of repetitive control [23] in that it is a technique for non-causally compensating repeated tracking errors in the execution of periodic motions. Algorithms of this form have previously shown good performance when applied to related problems where aerodynamic disturbances are considered, in particular the rejection of periodic wind disturbances on wind turbines [34,35].

Similar methods can also be found in the field of time waveform replication, as commonly applied to vibration testing systems (e.g., [36] and references therein). In such applications, the first-principles models guiding the iterative learning process are often replaced by experimentally identified frequency response functions.

Similar to the second group of learning algorithms, we do not assume a parameterized motion. However, we reduce the dimensionality of the corrections that we intend to learn by assuming that they are periodic. This allows us to parameterize the corrections as the coefficients of truncated Fourier series. The order of the Fourier series provides a means to trade off computational complexity and the ability to compensate for temporally local or high-frequency disturbances. Furthermore, the approach can be considered to be conceptually similar to the one presented by Lupashin and D'Andrea [28], which presents an adaptation strategy to correct for state errors at discrete points in time of parameterized motion primitives. However, we consider periodic errors (instead of errors at specific points in time), and do not require parameterization of the maneuver.

The contribution of this paper to the field of quadrocopter control lies in the application of methods from the fields of repetitive control and iterative learning control to quadrocopters. A general framework for arbitrary periodic motions is presented. We demonstrate how a feedback controller can be leveraged to shape the closed-loop dynamics of the quadrotor system, and show that a linear time-invariant approximation of the closed-loop dynamics suffices to guide the learning process. Using statistical properties of the disturbance, measurement noise, and the influence of nonlinearities, we derive the optimal inter-execution learning update step size. The validity of the approach and its performance is investigated through experiments in the ETH Flying Machine Arena with a quadrocopter under position control.

Furthermore, this paper introduces a novel method that extends the applicability of the repetitive control approach when the reference trajectory is too fast to be learnt directly, for example because the initial execution fails entirely. The core idea here lies in providing an improved initial guess of the disturbances degrading tracking performance. This typically enables learning of the trajectory because the errors are sufficiently small for the first-principles model to provide reliable information on how to compensate. To find the improved initial guess, we introduce a time scaling method that allows initial learning to occur at reduced maneuvering speeds and the transfer of learnt corrections from the reduced-speed execution to full speed. This method may also be applied to other complex dynamic systems where it is necessary to limit initial tracking errors in order to avoid the system failing. The time-scaling method provides an interesting alternative to methods that rely on aborting trials when the errors grow too large [31] in that