

Performing and extending aggressive maneuvers using iterative learning control

Oliver Purwin^{a,*}, Raffaello D'Andrea^b

^a Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY, USA

^b Department of Mechanical and Process Engineering, ETH Zürich, Switzerland

ARTICLE INFO

Article history:

Received 24 September 2008

Received in revised form

3 May 2010

Accepted 3 September 2010

Available online 17 September 2010

Keywords:

Iterative learning control

Motion control

Real-time

Unmanned aerial vehicle

ABSTRACT

This paper presents an algorithm to iteratively perform an aggressive maneuver, i.e. drive a system quickly from one state to another. A simple model which captures the essential features of the system is used to compute the reference trajectory as the solution of an optimal control problem. Based on a lifted domain description of that same model an iterative learning controller is synthesized by solving a linear least-squares problem. The controller adjusts a feedforward signal using the results of experiments with the system. The non-causality of the approach makes it possible to anticipate recurring disturbances. Computational requirements are modest, allowing controller update in real-time. The experience gained from successful maneuvers can be used to adjust the model, which significantly reduces transients when performing similar motions. The algorithm is successfully applied to a real quadrotor unmanned aerial vehicle. The results are presented and discussed.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

With the increasing application of autonomous systems there arises a need to take advantage of their full capabilities. Pushing the envelope of these vehicles invariably involves dealing with transients and nonlinear dynamics. One approach to improve the performance is to identify the system well and apply advanced control methods. However, the required level of accuracy could require extensive system identification efforts. Further, a model-based approach only works satisfactorily if all vehicles of a series have very similar dynamics, which could necessitate tight tolerances and expensive hardware. A different paradigm is to put the complexity in the software and take advantage of the low cost of sensors. A relatively simple model in conjunction with an adaptive algorithm and a well-chosen set of sensors allows each vehicle to experimentally determine how to perform a difficult maneuver and to compensate for individual differences in the system dynamics. This data-based approach to control has been proposed by many authors.

Moore [1] presents an algorithm to control a robotic manipulator based on sub-trees (state, action, behavior). These data structures store the experimental results of the system, to be retrieved in real-time for control. Depending on the current state of the system and the desired behavior the matching action is chosen. Schaal and Atkeson [2] pursue a similar approach. Local regression is performed on the stored data, the results of which are used to compute

a local model and controller. Hjalmarsson [3] and Jansson [4] describe iterative feedback tuning (IFT) which performs a stochastic gradient search on a performance metric of the system. The gradient of the performance is computed directly from input/output data by appropriate selection of the experimental input. Similar in spirit is the one-shot method virtual reference feedback tuning (VRFT) [5,6] which computes near-optimal controller parameters from input/output data. Michini and How [7] apply a \mathcal{L}_1 adaptive controller to a small autonomous vehicle. This version of a Model Reference Adaptive Controller (MRAC) uses a simple model in conjunction with performance and robustness metrics to adapt control parameters with the goal to make the system behave like the reference model.

Another approach is called iterative learning control (ILC). The idea behind ILC is that the performance of a system executing the same kind of motion repeatedly can be improved by learning from previous executions. Given a desired output signal ILC algorithms experimentally determine an open-loop input signal which approximately inverts the system dynamics and yields the desired output. Bristow et al. [8] provide a survey of different design techniques for ILC. The update of the input signal can be based on PD-type functions, which requires very little knowledge of the underlying plant dynamics. Plant inversion leads to fast convergence, but relies on a very accurate model. On the other hand, \mathcal{H}_∞ -based methods provide more robustness at the cost of performance. Another systematic option which requires a model is the minimization of a quadratic cost criterion.

Chen and Moore [9] present an approach based on local symmetrical double-integration of the feedback signal and apply it to a simulated omnidirectional ground vehicle. Two tuning parameters adjust the low-pass characteristics and the convergence rate

* Corresponding author. Tel.: +1 607 227 0014.

E-mail address: op24@cornell.edu (O. Purwin).

of the ILC. Ghosh and Paden [10] show an approach based on approximate inversion of the system dynamics. Chin et al. [11] merge a model predictive controller [12] with an ILC [13]. The real-time feedback component of this approach is intended to reject non-repetitive noise while the ILC adjusts to the repetitive disturbance. Cho et al. [14] put this approach in a state-space framework.

ILC allows the formulation of non-causal control laws, which means that an error occurring in the future can have an impact on the present control input. Therefore it is possible to preemptively compensate for disturbances or model uncertainties which are constant from trial to trial. Non-causality can be achieved in a natural way by formulating the problem and controller in the lifted domain, i.e., by stacking up all the inputs and outputs for the entire trial length into two large vectors and writing the input/output relationship as one large matrix operation. This exploitation of the repetitive nature of the experiments is made viable by advancements in computer processors and memory.

Rice and Verhaegen [15] present a structured unified approach to ILC synthesis based on the lifted state-space description of the plant/controller system. The sequentially semi-separable structure of the problem is then exploited to synthesize the controller efficiently.

In practice ILC have been applied to repetitive tasks performed by stationary systems, such as wafer stages [16], chemical reactors [17], or industrial robots [18]. Applications to autonomous vehicles are more rare.

This paper presents a lifted domain ILC algorithm which enables a system to perform an aggressive motion, i.e. drive the system from one state to another. Aggressive in this context characterizes a maneuver that takes place in the nonlinear regime of the system and/or close to the state or input constraints. This maneuver would be hard to tune by hand or would require very accurate knowledge of the underlying system. Instead, the featured algorithm only requires a comparatively simple model (which captures essential system dynamics) and initial guess for the input. In case of an unstable system it is assumed that a stabilizing controller is available. The feedforward signal required to perform the maneuver is iteratively determined using experiments.

While the algorithm is applicable to a wide range of systems it is particularly intended for autonomous vehicles. Hence the requirement that the controller update can be executed online with modest computational resources, putting the emphasis on computational efficiency. In general the algorithm re-uses as much data as possible for the purposes of safety and efficiency. For example, the model used to stabilize the system initially is employed to determine the ILC update law. If a particular maneuver is performed satisfactorily the gained knowledge can be utilized to perform a maneuver which is similar to the one just learned. This way a motion can be slowly extended, eventually executing a maneuver which could not have been performed given just the initial model.

The first step of the algorithm is the computation of the reference trajectory and input. This is accomplished by solving an optimal control problem based on the given model and constraints. The nonlinear model is then linearized about this reference and discretized, resulting in a discrete time linear time-varying (LTV) system. The lifted description of this LTV system defines the input-output relationship of the system for one complete experimental run in form of a single matrix. After performing an experiment the results are stored and compared with the ideal trajectory, yielding an error vector. Solving a linear least-squares (LLS) problem based on the lifted LTV system and the error vector yields the change in the input signal for the next trial. The algorithm terminates if the norm of the error vector is sufficiently small.

The algorithm is successfully applied to an unmanned aerial vehicle (UAV). After stabilization the UAV is a marginally stable




Fig. 1. Quadrotor unmanned aerial vehicle.

system which requires accurate feedforward inputs to track a reference satisfactorily. This makes the UAV a challenging testbed to justify the chosen approach.

The contribution of this work to the field consists in the presentation of the entire unified process of how to accomplish a maneuver which pushes the dynamic envelope of the vehicle given a relatively simple model. This includes the least-squares solution of the lifted system dynamics, extension of the maneuver based on previous experiments, and application to a real UAV.

The rest of the paper is organized as follows: Section 2 presents the dynamics of the vehicle used for algorithm derivation and implementation. Section 3 describes the algorithm to perform a single maneuver, consisting of reference generation and update laws for the control input. The learned maneuver is extended in Section 4, taking advantage of the previously collected information. Section 5 shows the successful application of the algorithm to a real rotorcraft, while Section 6 provides a conclusion and future prospects.

2. Vehicle dynamics

The presented algorithm is being applied to the quadrotor UAV shown in Fig. 1. The coordinates are defined as shown in Fig. 2 (left). Hover is defined as the airborne vehicle being horizontal and motionless. The dynamics of the vehicle while hovering are unstable because there is no restoring torque in the roll/pitch directions (rotation about x or y). Once the vehicle tilts out of the xy plane the tilted thrust vector will cause translational acceleration in all three axes. To perform experiments with the airborne vehicle an initial controller is required which stabilizes the UAV in hover. The dynamics of the vehicle in motion are very complex when taking vibration or aerodynamic effects into account, such as trailing propellers moving into the downwash of the leading propellers [19,20]. A more accurate model which would require the solution of partial differential equations is not available. Therefore a PD controller has been synthesized based on the linearized six degrees of freedom (DOF) rigid body dynamics which capture the salient features of the vehicle to first order. In the following this controller will be referred to as hover controller. Note that the hover controller is only active before and after an ILC run (to bring the vehicle into position), not during a trial. For a more extensive description of the vehicle see Appendix A.

The maneuver that will be covered in the subsequent sections of the paper consists of a sideways motion in the xz plane. This reduces the dynamics to two-dimensional space and three DOF without loss of generality. The approach of the algorithm is fundamentally the same, regardless of whether the model describes the dynamics in 2D or 3D. Further, the non-linear dynamics expected when moving in 2D are still quite significant.

Fig. 2 (right) shows the free-body diagram of the equivalent 2D vehicle. The global frame of reference (FOR) is denoted by x , z ,

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات