



Calhoun: The NPS Institutional Archive
DSpace Repository

Department of Mechanical and Aerospace Engineering (MAE) Faculty and Researchers' Publications

2006-08-21

Pseudospectral Feedback Control: Foundations, Examples and Experimental Results

Ross, I. Michael; Sekhavat, Pooya; Fleming, Andrew;
Gong, Qi

The American Institute of Aeronautics and Astronautics (AIAA)

<http://hdl.handle.net/10945/29678>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Pseudospectral Feedback Control: Foundations, Examples and Experimental Results

I. Michael Ross* Pooya Sekhavat† Andrew Fleming‡ Qi Gong§

Suppose optimal open-loop controls could be computed in real time. This implies optimal feedback control. These controls are, typically, nonsmooth. Nonsmooth controls raise fundamental theoretical problems on the existence and uniqueness of feedback solutions. A simple, yet powerful, approach to address these theoretical issues is the concept of a π -solution that is closely linked to the practical implementation of zero-order-hold control sampling. In other words, even traditional feedback controls involve open-loop controls through the process of sampling. In this paper, we advance the notion of Carathéodory- π solutions that use the sampling intervals to compute optimal open-loop controls. A sampling theorem is developed which indicates that the Lipschitz constant of the dynamics is a fundamental sampling frequency. This places computation at the level of first principles in describing the foundations for achieving feedback. We obtain these controls by way of pseudospectral (PS) methods as these techniques can generate optimal open-loop controls within fractions of a second even when implemented in a MATLAB[®] environment running on legacy computer hardware. In order to facilitate an exposition of the proposed ideas to a wide audience, we introduce the core principles only while relegating the intricate details to numerous recent references. These principles are then applied to generate PS feedback controls for the slew maneuvering of NPSAT1, a spacecraft conceived, designed and built at the Naval Postgraduate School and scheduled to be launched in Fall 2007.

I. Introduction

Consider the problem of generating feedback controls for the general nonlinear system,

$$\dot{x} = f(x, u, t), \quad u(t) \in \mathbb{U}(t, x(t)), \quad x(t) \in \mathbb{X}(t) \quad (1)$$

where $t \mapsto \mathbb{X}(t) \subset \mathbb{R}^{N_x}$ and $(t, x) \mapsto \mathbb{U}(t, x) \subset \mathbb{R}^{N_u}$ are set-valued maps¹⁻³ whose ranges are compact sets denoting the state and control spaces respectively, and $f : \mathbb{R}^{N_x} \times \mathbb{R}^{N_u} \times \mathbb{R} \rightarrow \mathbb{R}^{N_x}$ is a control-parameterized vector field that satisfies C^1 -Carathéodory conditions (see Appendix) for every admissible control function, $t \mapsto u$. We defer to Sec. II for a more detailed discussion of the meaning of a solution vis-à-vis the function spaces, \mathcal{X} and \mathcal{U} , that correspond to the state, $x(\cdot)$, and control, $u(\cdot)$, trajectories respectively. The control space, $\mathbb{U}(t, x)$, is allowed to be state dependent as these are important considerations in practical flight control^{4,5} as well as theory.¹ By feedback control, we mean a map, $k : \mathbb{R} \times \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_u}$ such that,

$$k(t, x(t)) = u(t) \in \mathbb{U}(t, x(t)) \quad (2)$$

It is well recognized^{1,2} that this is one of the most difficult problems in control theory. In this paper we add two additional requirements to this problem. In the first specification, we require that the open-loop controls, $t \mapsto u$, meet some specified optimality criterion of a Bolza-type cost functional,

$$J[x(\cdot), u(\cdot), t_0, t_f] := E(x(t_0), x(t_f), t_0, t_f) + \int_{t_0}^{t_f} F(x(t), u(t), t) dt \quad (3)$$

*Professor, Department of Mechanical and Astronautical Engineering. E-mail: imross@nps.edu. Associate Fellow, AIAA.

†Postdoctoral Research Associate, Department of Mechanical and Astronautical Engineering. E-mail: psekhava@nps.edu.

‡CDR, USN. E-mail: andrew.fleming1@navy.mil. Member, AIAA.

§Postdoctoral Research Associate, Department of Mechanical and Astronautical Engineering. E-mail: qgong@nps.edu.

where $J : \mathcal{X} \times \mathcal{U} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, while the second specification requires that the initial and final states meet some specified endpoint conditions,

$$(x(t_0), x(t_f), t_0, t_f) \in \mathbb{E} \subset \mathbb{R}^{N_x} \times \mathbb{R}^{N_x} \times \mathbb{R} \times \mathbb{R} \quad (4)$$

where t_0 and t_f are initial and final times under consideration. The time interval, $t_f - t_0$, may be finite or infinite. Infinite horizon problems are directly addressed in Refs. [6] and [7]. Thus, we are addressing a general problem of optimal feedback control for constrained nonlinear systems. Note that the cost functional is not necessarily limited to a quadratic criterion or the nonlinear dynamics limited to control-affine nonlinear systems. The limitations we impose on the system constraints will be clarified later but it essentially reduces to articulating all the constraints in terms of functional inequalities where the governing functions are smooth. This does not necessarily imply that the problem data (i.e. the stipulated sets in the problem definition) are smooth or that the functions describing these sets are smooth. A number of practical nonsmooth problems may be transformed (without simplification) to smooth problems by introducing additional control variables and constraints.⁸

Our main motivation for considering such difficult problems is aerospace applications. Aerospace problems remain one of the most challenging problems in control theory as optimality, in addition to the usual problems of constraints and nonlinearities, dictates the engineering feasibility of a space mission. For example, launch vehicle ascent guidance^{9,10} and orbit control⁸ problems demand fuel minimality as part of their engineering feasibility requirements. Even in problems where such an obvious criterion is not stated, optimality is inherently required for feasibility, safety and other considerations. For example, if a standard quaternion-based feedback control law is implemented for slewing the NPSAT1 spacecraft discussed in Sec. VII, it leads to an erroneous conclusion that a horizon-to-horizon scan is not feasible over an area of coverage because the attitude maneuver time (about 50 minutes) exceeds the coverage time (about 10 minutes). On the other hand, a minimum-time maneuver is equal to about half the coverage time (5 minutes) indicating that a horizon-to-horizon scan is indeed physically realizable.^{11,12} In other words, traditional feedback controls that are not based on optimality considerations may limit the performance of the system well under its true capabilities as evidently possible from the physics of the problem. Traditional feedback control laws may also diminish safety margins.¹³ For example, one aspect of reentry safety is the size of the footprint: the larger the footprint, the safer the entry guidance as it implies the availability of additional landing sites for exigency operations. Thus, footprint maximization is part of the entry guidance requirements.¹⁴ Consequently, entry guidance algorithms that are not based on optimal control compromise safety.¹⁵ It is apparent that both guidance and control problems in aerospace applications have optimality specified either directly or indirectly. Typically, vehicle guidance problems are finite-horizon problems while stabilization problems are infinite-horizon control problems. In extending this concept for general nonlinear control systems (i.e. not necessarily aerospace) we follow Clarke¹⁶ and note that there are essentially two principal objectives in all of control theory:

- 1) *Positional*: the state trajectory, $t \mapsto x$, must remain in or approach a given set of \mathbb{R}^{N_x} ;
- 2) *Transfer*: the control program, $t \mapsto u$, must transfer the state of the system from a given set in \mathbb{R}^{N_x} to a target set in \mathbb{R}^{N_x} .

Traditionally, optimal control theory has been largely used to address the transfer objective. The positional (i.e. stability) objective can also be addressed by way of optimal control theory by linking the Bellman value function to a Lyapunov function as facilitated by nonsmooth calculus.^{2,16} For the type of feedback control promulgated in this paper, we have provided such a link in a complementary paper,⁶ but under a stronger assumption of smoothness of the Lyapunov function. In this paper, we outline the issues, principles and certain implementation details. This allows us to discuss, at the very fundamental level, the key differences, between traditional feedback control theory and our proposal outlined in Sec. IV. We support Sec. IV by developing a key lemma in Sec. V that links the meaning of real-time computation to the Lipschitz constant of the dynamical system. In Sec. VI, we derive a few simple rules of thumb to facilitate an implementation of our ideas. A demonstration of our ideas is discussed in Sec. VII by way of ground-test results for the time-optimal feedback control of NPSAT1 (see Fig. 1) an experimental spacecraft conceived, designed and built at the Naval Postgraduate School, and scheduled to be launched in Fall 2007.




Figure 1. Artist's rendition of NPSAT1 in orbit.

II. Theoretical Issues in Optimal Feedback Control

A framework for solving the optimal feedback control problem posed in Sec. I is the well-known dynamic programming method. In this approach, we need to find a function, $\varphi : \mathbb{R} \times \mathbb{R}^{N_x} \rightarrow \mathbb{R}$, that solves the Hamilton-Jacobi-Bellman (HJB) partial differential equation,

$$\mathcal{H}(\partial_x \varphi(t, x), x, t) + \partial_t \varphi(t, x) = 0 \quad \forall t \in [t_0, t_f] \quad (5)$$

where \mathcal{H} is the lower Hamiltonian² given by,

$$\mathcal{H}(\lambda, x, t) := \min_{u \in \mathbb{U}} H(\lambda, x, u, t) \quad (6)$$

H is the control Hamiltonian,

$$H(\lambda, x, u, t) := F(x, u, t) + \lambda^T f(x, u, t) \quad (7)$$

and $\lambda \in \mathbb{R}^{N_x}$ is the covector, $\partial_x \varphi(t, x)$. A point worth emphasizing at this juncture and not widely discussed in textbooks is the implication of (6) on (5). That is, in order to merely construct the HJB equation, we need to obtain a closed-form solution to the optimization problem,

$$\text{(HMC)} \quad \begin{cases} \text{Minimize} & H(\lambda, x, u, t) \\ \text{Subject to} & u \in \mathbb{U}(t, x) \end{cases} \quad (8)$$

When \mathbb{U} is prescribed in terms of functional inequalities, Problem HMC is a nonlinear programming problem. Thus, in order to merely construct the HJB equation, the dynamic programming method requires a closed-form solution to a nonlinear programming problem! This is one reason why complex engineering problems seeking the HJB route are posed in a simplified manner that facilitates a closed-form solution to Problem HMC. This is part of a prevailing philosophy in control theory of seeking exact solutions to simple problems. In striking contrast, our goals are to seek approximate, but highly accurate solutions to complex more realistic problems. In any case, even under such simplifications, the HJB equation does not have a (C^1 -smooth) solution for φ even for simple problems. One of the earliest and best-known criticisms of the HJB approach is due to Pontryagin et al.¹⁷ This issue was resolved in 1983 by Crandall and Lions¹⁸ who developed the notion of viscosity solutions to the HJB equation. Recent updates to these ideas are described in Ref. [19]. Considering that obtaining even smooth solutions to partial differential equations is itself a very challenging problem, it is no surprise that obtaining viscosity solutions to the HJB equation are even more problematic. Finally, even if φ is smooth, there is, in general, no continuous function, $u = k(x)$, that satisfies the HJB equation,

$$F(x, k(x), t) + \partial_x \varphi^T f(x, k(x), t) + \partial_t \varphi(t, x) = 0 \quad (9)$$

Thus, even if all the hurdles described above are overcome, we cannot hope to find a continuous, optimal feedback law. That continuous feedback controls, do not exist even for stabilizing control systems was proved by Brockett as far back as 1983 by his famous counterexample of the nonholonomic integrator.²⁰ Thus, in developing a general feedback control theory for the problem defined in Sec. I, we need to allow for the possibility of a discontinuous feedback law. When $k(x)$ is discontinuous, $f(x, k(x), t)$ is discontinuous even if we had assumed $f(x, u, t)$ to be continuous in all its variables. Obviously, feedback is expected to alter the system dynamics, but a discontinuous feedback law changes the dynamics to such an extent that even the very definition of a solution, $x(\cdot)$, for a discontinuous differential equation is called into question. This is because when $f(x, k(x), t)$ is continuous, we say that $x(\cdot)$ is a solution to the differential equation,

$$\dot{x} = f(x, k(x), t) \quad (10)$$

if $\dot{x}(t)$ coincides with $f(x(t), k(x(t)), t)$ for all t . In the discontinuous case, Carathéodory's extension of this concept uses the "almost everywhere" moniker together with $x(\cdot)$ being absolutely continuous (see Appendix). This solution concept has served very well for open-loop systems,

$$\dot{x} = f(x, u(t), t) \quad (11)$$

and forms the basis of open-loop optimal control theory;^{17,21} however, for closed-loop systems, it is well-known that the Carathéodory solution concept turns out to be quite unsatisfactory from the point of view of existence and uniqueness.²² This is because (11) has discontinuities at most in the independent variable, t , whereas (10) allows discontinuities in the dependent variable, x , as well. Hence, the question of providing a suitable definition of a solution for discontinuous differential equations is of fundamental importance. Among the large variety of definitions proposed in the literature, the Filippov solution concept²² is widely used in control applications. The Filippov solution is based on a convexification of the differential inclusion associated with a differential equation. In 1994, Ryan²³ and Coron and Rosier²⁴ extended Brockett's conclusion by proving the validity of his result for discontinuous feedback if $x(\cdot)$ is a Filippov solution. As a result of such outstanding issues in discontinuous differential equations, vigorous research in this field continues; see, for example, Ref. [25] and the references contained therein. In recent years, a neo-classic notion of closed-loop system sampling has been advanced as a means to address such long-standing theoretical difficulties. This solution concept is not only different from the Filippov or Carathéodory solution concepts, it is also one of the most obvious and practical ways of engineering a "digital" solution to a controlled differential equation. In other words, what has happened in recent years is the close juxtaposition of practical methods for control with advance mathematics, particularly nonsmooth calculus. This is in sharp distinction to the development of control theory of prior years that were largely focused on seeking closed-form or "analytical" expressions for feedback representations, $k(x)$, wherein k was designed by means of elementary operations (e.g. multiply, add) over elementary programmable functions (e.g. polynomials, transcendental representations). In the concept proposed in this paper, $k(x)$ is computed by way of "designer" functions that have no discernable analytical representations. Thus, theory and computation are now fundamentally intertwined in first principles itself. A preliminary stability theory from this perspective is provided in our companion paper;⁶ here, we outline the basic principles.

III. Sample and Hold Feedback Control

We first consider arbitrary time-invariant feedback control laws, $\tilde{k} : \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_u}$ applied to autonomous dynamical systems, $\dot{x} = f(x, u)$. In the next section, we consider the more general *time-varying* feedback law, $k : \mathbb{R} \times \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_u}$, as this is central to our new approach.

In following Clarke,¹⁶ we let $\pi = \{t_i\}_{i \geq 0}$ be a partition of the time interval, $[0, \infty)$, with $t_i \rightarrow \infty$ as $i \rightarrow \infty$. The quantity,

$$diam(\pi) := \sup_{i \geq 0} (t_{i+1} - t_i)$$

is called the diameter of π . Given an initial condition, x_0 , a π -trajectory, $x(\cdot)$, corresponding to $\tilde{k}(x)$ is defined as follows:¹⁶ From t_0 to t_1 , we generate a classical solution to the differential equation,

$$\dot{x}(t) = f(x(t), \tilde{k}(x_0)), \quad x(t_0) = x_0, \quad t_0 \leq t \leq t_1 \quad (12)$$

A classical solution exists for (12) because f is Lipschitz-continuous with respect to its first argument. At $t = t_1$, we set $x_1 = x(t_1)$ and restart the system with $u = \tilde{k}(x_1)$,

$$\dot{x}(t) = f(x(t), \tilde{k}(x_1)), \quad x(t_1) = x_1, \quad t_1 \leq t \leq t_2$$

Continuing in this sample-and-hold manner, we end up with a closed-loop trajectory corresponding to a piecewise constant open-loop control. *Note that this π -solution concept is not to be confused with an Euler polygonal arc.* Under the concept of a π -trajectory, it is possible to rigorously prove a rich number of theorems that have practical consequences. These theorems directly connect the “epsilons and deltas” of abstract mathematics to practical digital computer implementations of optimal control.^{16,26} An implementation of the sample-and-hold feedback concept is illustrated in Fig. 2. In order to maintain the focus of this paper

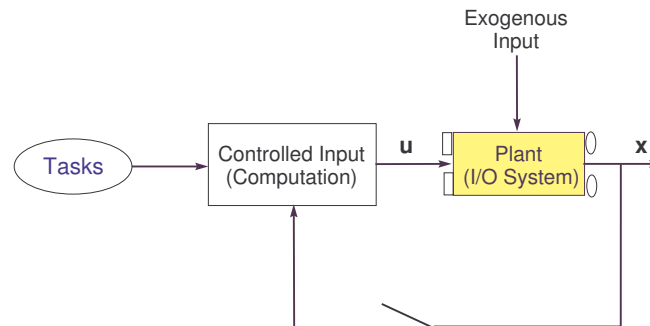


Figure 2. Illustrating closed-loop control by a sampling technique.

on the main elements of feedback control, we assume state information to be available on demand. The problem of nonlinear state estimation is addressed elsewhere.^{27,28} The switch in Fig. 2 remains open over open time intervals, $(t_{i+1} - t_i)$, and is closed only over zero measure (in time). It is clear that a π trajectory is indeed equivalent to a zero-order-hold (ZOH) digital implementation of a continuous-time input signal; however, note that *in the π -solution concept, the π -trajectories are viewed as fundamental objects that are not explicitly reliant upon a digital computer.* To drive home this point, we note that we would consider π -solutions even if we had analog computers only (i.e. even if digital computers were nonexistent). In this case, a π -solution would be implemented by the switch mechanism illustrated in Fig. 2. Thus, in a mathematical analysis of solutions to differential equations, we now consider sequences of π -trajectories parameterized by $diam(\pi)$. In principle, the limiting situation $diam(\pi) \rightarrow 0$ can then be connected to traditional feedback control theory; however, when traditional control theory is implemented via a digital computer, it is clear that the practice of the two concepts merge before $diam(\pi) \rightarrow 0$. Thus, in the new perspective promoted here, we discard the limiting condition, $diam(\pi) \rightarrow 0$, in favor of an analysis to considering what happens when $diam(\pi) < \delta_\pi$ where $\delta_\pi > 0$ is a given number. *Under this philosophy, we no longer demand exactness.* Thus, for example, rather than demand $\|x(t)\| \rightarrow 0$ as $t \rightarrow \infty$ as in standard asymptotic stability theory, we only require that there exist some $T > 0$ such that $\|x(t)\| < \epsilon_x \forall t \geq T$ where $\epsilon_x > 0$ is a specified tolerance. This is the notion of practical stability^{29,30} introduced in the 1960s. In adopting this concept, we require that $\|x(t)\| < \epsilon_x \forall \epsilon_x \geq \epsilon_{ses} > 0$ where $\epsilon_{ses} > 0$ is a given number based on sensor and estimation errors. Even under simulation, $\epsilon_x \neq 0$ as a result of computational errors. As further elaborated in later sections of this paper, these new concepts illustrate how theory and computation have become intertwined at the fundamental level.

As a matter of comparison, we note that a ZOH implementation of a standard constant gain feedback control law, $u = Kx$, where K is some gain matrix, generates a π trajectory. It is therefore clear that a *practical (digital) implementation of any closed loop controller involves an open-loop strategy.* The most widely used open-loop strategy is a *hold* implementation. This concept is illustrated in Fig. 3. Also shown in Fig. 3 is a practical modification of the π trajectory when the time to compute $\tilde{k}(x)$ is non-negligible as can happen with legacy systems employing old computer technology. An extension of the ZOH concept is the first-order hold. In a first-order (or higher-order) hold, $u(t)$ is obtained by extrapolation to the future from the past. Such implementations require memory storage of an appropriate number of previous measurements to perform the extrapolation. *Thus, a second point to note here is that a natural extension of this concept generates feedback implementation policies that are backward-looking approximation to the continuous-time*

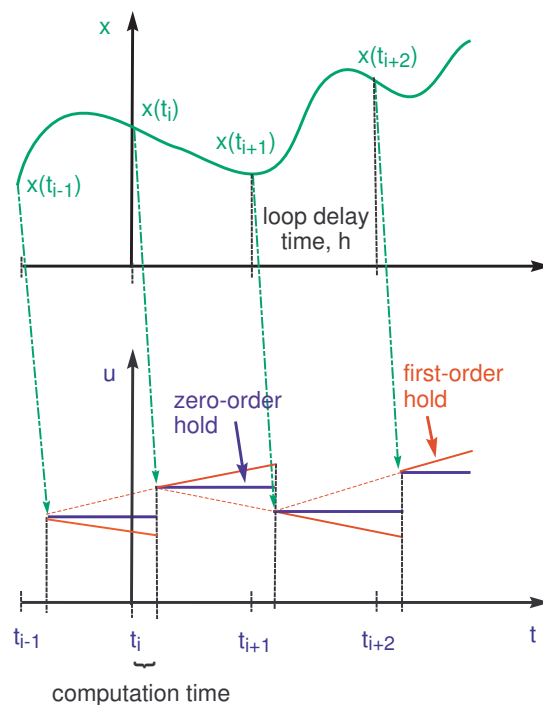


Figure 3. Standard sampling implementations of control signals illustrating the “backward-looking” feedback policy; also shown are computational times that are assumed to be less than the time-delay in receiving sensor updates.

signal, $u(t) = \tilde{k}(x(t))$. That is, we have a staircase approximation for a zero-order hold, a (backward) secant approximation in a first-order hold etc.

In summary, a sample-and-hold feedback controller generates a π -trajectory as a fundamental theoretical and practical solution to a controlled differential equation whether or not the feedback controller is smooth. Note that these ideas do not address the determination of $\tilde{k}(x)$; thus, finding a feedback map that guarantees stability of the closed-loop system continues to dominate ongoing research in nonlinear control theory.

IV. A Forward-Looking Approach to Feedback

Now suppose that the feedback law was time-varying: $u = k(t, x)$. In principle, we could implement it in exactly the same style as suggested in Sec. III and all the prior conclusions would hold. That is, we could continue to implement a sample and hold procedure,

$$u(t) = k(t_i, x(t_i)) \quad t \in [t_i, t_{i+1}] \quad (13)$$

or its standard extensions to a first- or higher-order hold. We would indeed arrive at this conclusion by the standard procedure¹⁷ in optimal control theory in dealing with non-autonomous systems by introducing a new state variable, $\hat{x} \in \mathbb{R}$, such that

$$\hat{\dot{x}} = 1, \quad \hat{x}(t_0) = t_0 \quad (14)$$

By augmenting (14) with the dynamics as in,

$$\begin{aligned} \dot{x} &= f(x, u, \hat{x}) \\ \hat{\dot{x}} &= 1 \end{aligned}$$

we can define a new state, $\tilde{x} = (x, \hat{x})$ and transform $k(t, x)$ quite simply to $\tilde{k}(\tilde{x}) = k(\hat{x}, x)$. In Ref. [31], Sussmann notes that this procedure adopted by Pontryagin et al¹⁷ to address non-autonomous systems is quite limiting from a theoretical standpoint as (14) assumes that $f(x, u, t)$ is of class C^1 with respect to t , a much stronger requirement than f being merely measurable in t , a sufficient condition for the existence of a Carathéodory solution. Here, we propose that this idea is also quite unsatisfactory for practical feedback

control. This is because a time-varying feedback law, $k(t, x)$, inherently contains an element of prediction as a result of its explicit dependence on the clock time, t . Hence, it is clear that a better policy for control in between samples is a forward-looking implementation,

$$u(t) = k(t, x(t_i)) \quad t \in [t_i, t_{i+1}] \quad (15)$$

This strategy does not pose the problem of the meaning of a solution for a discontinuous differential equation,

$$\dot{x} = f(x, k(t, x(t_i)), t)$$

as the discontinuities between the samples are now only with respect to the independent variable, t , and not the dependent variable, x . Thus, any possible discontinuity we have to contend with are in the open-loop segments over the sample time. Hence, we can define a solution over the sample segment in the standard Carathéodory sense and then glue all these Carathéodory-solution pieces in the same manner as the π -trajectory. Hence, we distinguish this concept as a Carathéodory- π trajectory. The motivation for choosing Carathéodory segments, and not for instance the Filippov solution, is motivated by optimal control considerations. That is, we will generate the open-loop segments by computing open-loop optimal controls. One motivation for optimality is, as indicated in Sec. I, aerospace problems. Regardless, optimal control also has the advantage of guaranteeing stability through the Bellman value function serving as a natural Lyapunov function.⁶ Thus, the existence of an open-loop optimal control is now directly tied to constructing closed-loop solutions. As already noted in Sec. II, a rich number of theorems are available for optimal open-loop trajectories defined in the Carathéodory sense. Consequently, we are naturally led to a Carathéodory modification of the π trajectory. The practical consequence of this Carathéodory- π trajectory is that over any time segment, $[t_i, t_{i+1}]$, the control is continuously updated by the clock information, t . This implementation implies that the controlled input box in Fig. 2 (i.e. the computation box) includes a memory storage device for dumping the signal, $u(t) = k(t, x(t_i)) \quad t \in [t_i, t_{i+1}]$ that is to be fed to the plant in between samples. The potentially extra cost of this new implementation compared to the ZOH policy is a clock and a memory storage device to store the signal, $u(t), t \in [t_i, t_{i+1}]$. Note the key difference between this policy and a higher-order hold. In a higher-order hold implementation (which also requires memory), the control

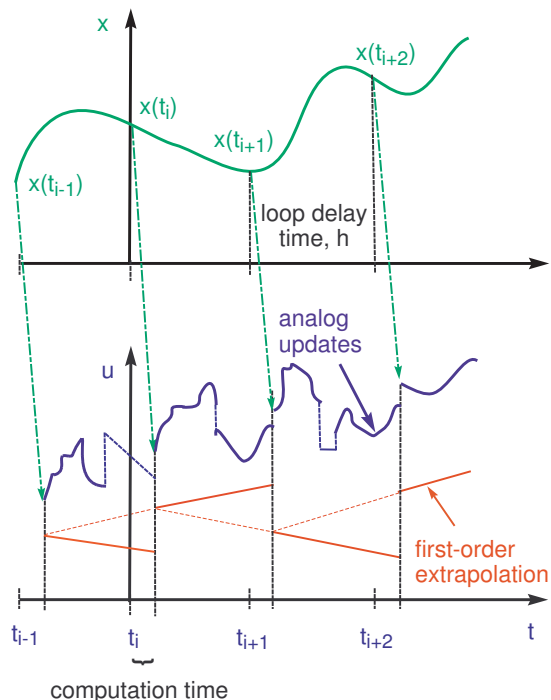


Figure 4. A proposed implementation of control signals in a “clock-based” forward-looking feedback policy for the generation of Carathéodory- π solutions, and its contrast to the “backward-looking” first-order hold.

values in between samples is based on the *past controls* with current values obtained based on *extrapolation*.

In the implementation of (15) the control signal is based on *current controls*, the currency being indicated by the clock time. These points are illustrated in Fig. 4. Note that it is not possible to use a clock-based feedback implementation if the feedback control law is independent of the clock time as in $u = \tilde{k}(x)$. Thus, a time-varying feedback control law, $u = k(t, x)$, is integral to our ideas.

At first glance, the clock-based feedback implementation may appear to deliver a minimal gain in performance at significant expense. The reason this is not necessarily true is that with regards to (13), δ_π must be substantially smaller than the diameter of π required of (15) for any given ϵ_x . This is a direct effect of utilizing the clock information in the feedback law, $k(t, x)$. Deferring the details of this discussion to Sec. V, (15) implies that we have traded the requirement of a small δ_π to the purchase of a clock and a memory unit. The higher allowable δ_π now implies that we can afford to have a very sophisticated feedback strategy (e.g. optimal) than traditional analytical expressions for k . A computation of this sophisticated feedback law, $u = k(t, x)$, is now permitted to have a sampling time far greater than that afforded by (13). Thus, we abandon the notion of seeking closed-form expressions for k and resort to its fundamental form, $k : \mathbb{R} \times \mathbb{R}^{N_x} \rightarrow \mathbb{R}^{N_u}$ where the map is given by some computational algorithm. In essence, a “large” value of δ_π is now used to support the computation time for the control algorithm. Hence, we add an explicit clock to Fig. 2 as shown in Fig. 5 to emphasize the shift in philosophy that clock times are used to gener-

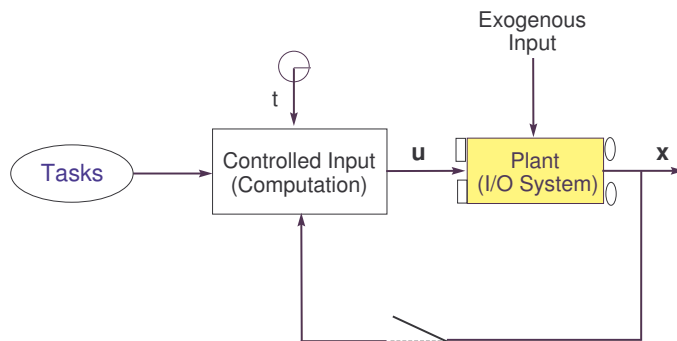


Figure 5. Implementing a clock-based feedback controller via computation of open-loop controls.

ate open-loop “analog” signals, $u(t) = k(t, x(t_i))$ thereby subverting the traditional concept of real-time as something equivalent to a very small δ_π . Intuitively, it is clear that the requirement on computation time depends upon the “time constant” of the plant dynamics among other things. The Lipschitz constant of f acts as this fundamental time constant as developed in the next section.

V. Feedback-Based Requirements on Computation Time

Consider a nonlinear *model* of a control system,

$$\dot{x} = f(x, u, t; p_0) \quad (16)$$

where $p_0 \in \mathbb{R}^{N_p}$ is a constant of system parameters (such as inertia, mass etc.). The purpose of feedback is to largely manage various uncertainties such as imperfections in plant modeling, unmodeled/unknown exogenous input (see Fig. 5), estimation errors, etc. To this end, we let $t \mapsto \zeta(t)$, be some function in L_{loc}^∞ so that

$$\dot{x} = f(x, u, t; p) + \zeta(t) \quad (17)$$

represents the dynamics of the real system (plant). In (17), p denotes the actual plant parameters. Thus $p = p_0$ and $\|\zeta\|_{L^\infty} = 0$ corresponds to perfect modeling. We implement a feedback policy as follows; see Fig. 6. At $t = t_i$, we compute $[t_i, t_f] \mapsto u = k(t, x_R(t_i))$ where x_R is the state of the real system (plant). In the infinite horizon case, we compute $[t_i, \infty) \mapsto u$ by way of a domain transformation technique.^{6,7} In any event, under the action of an open-loop control, $[t_i, t_{i+1}] \mapsto k(t, x_R(t_i))$, the state of the model at t_{i+1} , is given by,

$$x_M(t_{i+1}) = x_R(t_i) + \int_{t_i}^{t_{i+1}} f(x_M(t), k(t, x_R(t_i)), t; p_0) dt \quad (18)$$

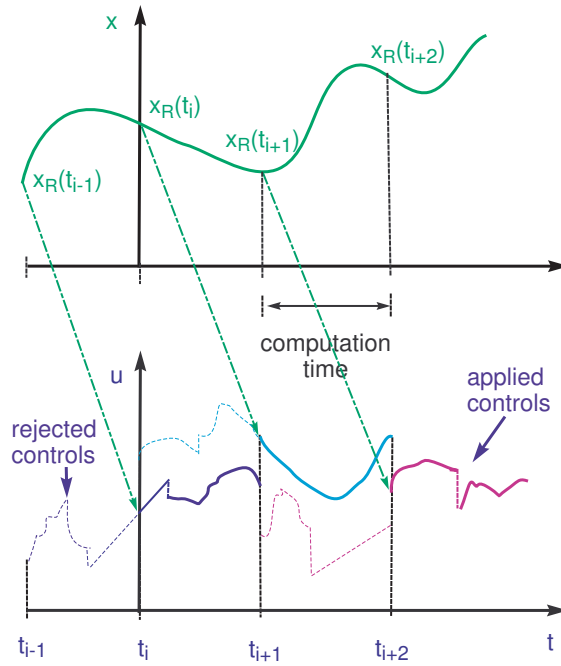


Figure 6. Practical implementation of control signals in clock-based feedback control; the control, $u(t) = k(t, x_R(t_i)), t \in [t_i, t_{i+1}]$ is computed but not applied.

Let t_{i+1} be the instant where the control $t \mapsto k(t, x_R(t_i))$ is available for application to the plant (see Fig. 6); that is,

$$\tau_i := t_{i+1} - t_i \quad (19)$$

is the representative computational time. Thus, the state of the plant at t_{i+1} is determined by the action of the control, $[t_i, t_{i+1}] \mapsto k(t, x_R(t_{i-1}))$, and is given by,

$$x_R(t_{i+1}) = x_R(t_i) + \int_{t_i}^{t_{i+1}} f(x_R(t), k(t, x_R(t_{i-1})), t; p) dt + \int_{t_i}^{t_{i+1}} \zeta(t) dt \quad (20)$$

That is, $[t_i, t_{i+1}] \mapsto x_R$ differs from the ideal/model trajectory, $[t_i, t_{i+1}] \mapsto x_M$, due to the effects of the computational delay time as well as deviations caused by $t \mapsto \zeta(t)$ and p . Subtracting (18) and (20) we have,

$$x_R(t_{i+1}) - x_M(t_{i+1}) = \int_{t_i}^{t_{i+1}} f(x_R(t), k(t, x_R(t_{i-1})), t; p) dt - \int_{t_i}^{t_{i+1}} f(x_M(t), k(t, x_R(t_i)), t; p_0) dt + \int_{t_i}^{t_{i+1}} \zeta(t) dt \quad (21)$$

Although it is possible to estimate an upper bound for the right hand side of (21) under existing assumptions, a stronger result is possible under the following additional assumption.

Assumption 1

- i) For each x and almost all t , the function, $(u, p) \mapsto f(x, u, t; p)$, is Lipschitz continuous.
- ii) For each x , the function, $t \mapsto k(t, x)$, is in L^∞_{loc} .

Under Assumption 1, we have from (21),

$$\|x_R(t_{i+1}) - x_M(t_{i+1})\| \leq Lipf_x \int_{t_i}^{t_{i+1}} \|x_R(t) - x_M(t)\| dt + Lipf_u \int_{t_i}^{t_{i+1}} \|k(t, x_R(t_{i-1})) - k(t, x_R(t_i))\| dt + Lipf_p \|p - p_0\| \tau_i + \|\zeta\|_{L^\infty} \tau_i \quad (22)$$

where, $Lipf$ denotes the Lipschitz constants of f with respect to the subscripted variables.

Lemma 1 Let $\delta > 0$ be the smallest number such that,

- a) $\|k(\cdot, x_R(t_{i-1})) - k(\cdot, x_R(t_i))\|_{L^\infty} \leq \alpha_1 \delta \frac{Lipf_x}{Lipf_u}$
- b) $\|\zeta\|_{L^\infty} \leq \alpha_2 \delta Lipf_x$,
- c) $\|p - p_0\| \leq \alpha_3 \delta \frac{Lipf_x}{Lipf_p}$

for some $\alpha_i \in (0, 1)$ and $\sum_i \alpha_i = 1$. Then, for any given $\epsilon > 0$, if

$$\tau_i \leq \frac{W(r)}{Lipf_x} \quad (23)$$

it implies that $\|x_R(t_{i+1}) - x_M(t_{i+1})\| \leq \epsilon$, where $W(r)$ is the Lambert W function (see Appendix), and r is the ratio, $r := \epsilon/\delta$.

Proof: From (22) and conditions a), b), and c) of the lemma, we have,

$$\|x_R(t_{i+1}) - x_M(t_{i+1})\| \leq Lipf_x \int_{t_i}^{t_{i+1}} \|x_R(t) - x_M(t)\| dt + \delta Lipf_x \tau_i \quad (24)$$

From Gronwall's Lemma (see Appendix), (24) reduces to,

$$\|x_R(t_{i+1}) - x_M(t_{i+1})\| \leq \delta Lipf_x \tau_i \exp(Lipf_x \tau_i) \quad (25)$$

From the definition of the Lambert W function (see Appendix), it can be easily verified that for $y, z \in \mathbb{R}_+$, if $z \leq W(y)$ then $ze^z \leq y$; hence, from Eqs. (23) and (25), we have $\|x_R(t_{i+1}) - x_M(t_{i+1})\| \leq r\delta = (\epsilon/\delta)\delta = \epsilon$.

Remark 1 Under perfection, that is, if $\|\zeta\|_{L^\infty} = 0$ and $\|p - p_0\| = 0$, the only contribution to δ is from condition a) of the lemma. If $t \mapsto k(t, x(t))$ is an optimal control, then it follows from Bellman's principle of optimality and local uniqueness that $k(t, x(t_{i-1})) = k(t, x(t_i))$. Hence, under perfection and optimality, $\delta = 0$ for arbitrary computational delay. Hence, $r \rightarrow \infty$ for any $\epsilon > 0$. Since $W(r) \rightarrow \infty$ as $r \rightarrow \infty$, (23) conforms to the fact that under ideal conditions, open-loop optimal controls are sufficient if these controls are generated to meet performance and constraint specifications.

Remark 2 Under the conditions of Remark 1, we have $\delta \neq 0$ for sample-and-hold feedback except under the special situation, $k(t, x(t_{i-1})) = k(t, x(t_i)) = \text{constant} \forall t_i$. This is one reason why the Carathéodory- π feedback solution concept allows a longer computational delay than ZOH sampling.

Remark 3 The clock-based feedback controller has an inherent safety component in the following sense: If the feedback signals ("wires") were cut off in Fig. 5, open-loop constrained optimal control signals remain as plant inputs in contrast to a constant signal of a ZOH implementation.

A. Applications of Lemma 1

As noted in Sec. I, one of the objectives of control theory is positional. That is, for the positional objective of stabilizing a system about the origin, we require,

$$\|x_R(t)\| \leq \epsilon_{Req} \quad \forall t \geq T \quad (26)$$

where $\epsilon_{Req} > 0$ is a given number based on design requirements and consistent with the accuracy of the state estimator (and hence sensor errors). It is well-known¹⁶ that stability can be achieved by optimal control principles. Hence, if optimal controls are generated at the rate specified by (23), stability can be achieved under appropriate technical conditions. A proof of such a theorem, along with several example problems, are discussed in Ref. [6]. Thus, (23) provides a sufficient condition for the definition of "real-time" in real-time optimal control.

For finite-horizon feedback control (as in vehicle guidance), we require that $x_R(t_i)$ $t_i \in \pi$ be such that for each $x_R(t_i)$, there exist open loop controls, $[t_i, t_f] \mapsto k(t, x_R(t_i))$ such that at time, t_f ,

$$(x_R(t_f), t_f) \in \mathbb{E}_f \subset \mathbb{E}$$

Thus, if open-loop optimal controls are generated starting at $t = t_0$, it is possible to guide a vehicle to its target by generating a Carathéodory- π feedback solution. An example of such a strategy for entry vehicle guidance is developed in Ref. [5].

B. Connections to Continuous-Valued Feedback Control

It is worth noting that we never find or compute, $(t, x) \mapsto k(t, x) \forall x \in \mathbb{X}$; rather, we only compute its “semi-discrete, semi-analog” representation in the sense that we find $k(t, x)$ at only discrete points in state space but allow $k(t, x)$ to be computed at all times, t . Thus, even if there were to be a unique feedback controller, we only consider its semi-discrete representations.

From a theoretical point of view, we may consider limit points of convergent subsequences of these semi-discrete solutions as forming generalized solutions to nonsmooth differential equations, $\dot{x} = f(x, k(t, x), t)$, $x(t_0) = x^0$. This is an open research problem. This point, once again, illustrates the convergence of theory with practice!

VI. Practical Implementation of the Principles

The principles developed in the previous section deal with the “original” control problem head on. That is, instead of seeking to minimize quadratic cost functions that may have no practical bearing on the actual control system, or simplify the dynamical model to derive closed-form solutions, we propose to deal directly with the original problem itself. As pointed in Sec. I, our motivations for this route largely stem from aerospace applications. Philosophically, what we seek are approximate solutions to the exact (i.e. original) problem rather than exact solutions to the approximate problem. This leads us to the following precepts:

A. Precision in Practical Control

The accuracy afforded by the sensors/estimation provides a lower limit on the precision attainable by any control system, regardless of the sophistication of the control law. That is, no amount of sophistication in control theory can achieve precision finer than the resolution of the sensors. Thus, if the precision of the sensors/estimation is $\epsilon_{ses} > 0$ in x (e.g. 3σ value), then it is impossible to (verifiably) achieve any stability requirement that has $\epsilon_{Req} < \epsilon_{ses}$. Alternatively, for a given precision requirement, we must choose sensors with precision such that, $\epsilon_{ses} \leq \epsilon_{Req}$. Typically, we choose $\epsilon_{ses} \approx 0.1\epsilon_{Req}$; that is, we choose sensors to be about an order of magnitude finer in sensing than the requirement imposed on stability. Thus, this implies that if we seek to transfer a system from an initial point, $x_0 = x^0$, to a final point, $x_f = x^f$, we replace this notion, without prejudice, by seeking controls that transfer the state from an initial epsilon set, such as an epsilon ball,

$$x_0 \in \mathbb{B}(x^0, \epsilon_{ses}) \quad (27)$$

to a final epsilon set,

$$x_f \in \mathbb{B}(x^f, \epsilon_{Req}) \quad (28)$$

where $\mathbb{B}(x^*, r)$ denotes a closed ball of radius r centered around x^* ,

$$\mathbb{B}(x^*, r) := \{x : \|x - x^*\| \leq r\} \quad (29)$$

This point once again illustrates the merging of approximation theory with control theory at the fundamental level. In Sec. VII, we revisit such epsilon sets explicitly with the epsilon sets defined in terms of the infinity norm.

B. Impact of Computational Inexactness on Precision

In demonstrating the preceding principles for general nonlinear systems by way of numerical computations, we note that it is quite impossible to obtain exact solutions. That is, numerical calculations automatically introduce $\zeta(t) \neq 0$; hence a perfect numerical simulation of a perfect plant is quite impossible. Alternatively, we regard numerical errors as simulating state estimation errors with $\epsilon_{ses} = \epsilon_{num} > 0$, associated with a perfect plant (i.e. model) and imperfect measurements. We use this notion explicitly in Sec. VII. Thus, traditional theoretical issues such as $\epsilon \rightarrow 0$ must be replaced by other epsilon-delta concepts; for instance, the existence of a δ such that for all $\epsilon \geq \epsilon^* > 0$, something happens (see Lemma 1). Furthermore, since the central point of feedback is to manage uncertainties, the very notion of an exact solution is an anathema since if exact modeling was possible, open-loop control would suffice and feedback would be unnecessary. Consequently, the fact that our methods rely on approximations is not to be construed as a shortcoming; rather, we contend that approximations are essential and inherent for the design of feedback control itself.

What we have done, and demonstrate experimentally in Sec. VII, is to bring computational principles to the very foundations of control theory rather than as an afterthought. Thus, quantities such as ϵ_{num} , ϵ_{ses} , ϵ_{Req} etc., must be properly coordinated with the physics of the problem in a manner that a mathematically realizable solution exists if a physically realizable solution is possible. This is essentially a problem of designing the control system.

C. Computation of Real-Time Open-Loop Controls

The principles described in Sections III-V rely on a sufficiently fast generation of open-loop controls. Thus, if open-loop controls can be generated as demanded by (23), closed loop is achieved quite simply by generating Carathéodory- π solutions. In recent years, it has become quite apparent that pseudospectral (PS) methods^{32–40} are capable of generating optimal open-loop controls within fractions of a second,^{36,37,39} even when implemented in legacy hardware running MATLAB[®]. This implies that real-time optimal controls can be generated for systems with large Lipschitz constants; that is, systems with fast dynamics. This simply follows by re-writing (23) as,

$$Lipf_x \leq \frac{W(r)}{\tau_c}$$

where τ_c is the largest computational time. Nonlinear systems for which such PS-based feedback controls have been demonstrated range from the classical problem of swinging up a pendulum,⁶ to satellite attitude control,¹¹ to reentry vehicle guidance.⁵ Solutions to over a hundred problems are documented in Ref. [32] and the references contained therein. In Sec. VII, we provide experimental results that demonstrate this concept for the attitude control of NPSAT1.

PS methods for computing optimal controls are readily available by means of the software package, DIDO.⁴¹ DIDO is a minimalist’s approach to solving optimal control problems. Only the problem formulation as described in Sec. I is required. DIDO incorporates the Covector Mapping Principle^{40,42–45} (CMP) which essentially allows dualization to commute with discretization so that the necessary conditions for optimality can be automatically verified.

D. Generation of Time-Smooth Controls

In general, the controls that generate the Carathéodory- π solutions are nonsmooth with respect to time. This nonsmoothness arises due to two key facts:

1. The controls are allowed to be discontinuous in between sampling times, and
2. The controls are allowed to be discontinuous at the sampling times.

The discontinuity at the sampling time is no different from those that occur in the standard ZOH controller; however, the discontinuity in between the sampling time is allowed in order to achieve high performance (i.e. optimality) as motivated by problems in aerospace and introduced in Sec. I. In certain applications, discontinuous controls may not be desirable and smooth control may be stated as a requirement. As noted earlier, the traditional approach to designing such control systems is by way of choosing quadratic cost functions. This work-around fails in our proposed method as it can only prevent discontinuities in between sampling times and not the discontinuities at the sampling time. One option is to impose continuity of the controls at the sampling times. Imposing such continuity conditions has two drawbacks:

- i) The controls are, at best, continuous and may continue to be nonsmooth with respect to time, and
- ii) A theoretical analysis of the system would be unnecessarily difficult due to the differences between the posed problem and the implemented solution.

A far simpler approach to generating time-smooth controls is simply by dynamic extension. That is, we now consider the state space to be a subset of $\mathbb{R}^{N_x} \times \mathbb{R}^{N_u}$ where the new states are given by,

$$\bar{x} = [x^T, u^T]$$

The dynamics of the original states are, as before, given by $\dot{x} = f(x, u, t)$. The dynamics for the “state,” u , is now given by,

$$\dot{u} = \bar{u} \tag{30}$$

where \bar{u} is the new control variable that is chosen from some appropriate compact set,

$$\bar{u} \in \bar{\mathcal{U}} \subset \mathbb{R}^{N_u}$$

Since we automatically require that $\bar{u}(\cdot) \in L^\infty$, (30) guarantees that $u(\cdot) \in W^{1,\infty}$. In other words, (30) ensures in a single stroke that the control trajectory for the original controls, $t \mapsto u$, will be smooth at both the sampling time as well as in between samples. The inertia introduced for the original controls, u , by way of $\bar{\mathcal{U}}$ is at the discretion of the designer. Clearly, higher-order smoothness may be attained in a similar fashion. Since the new problem has a new vector field, $[f(x, u), \bar{u}]$, a theoretical analysis can be carried out in concert with practical implementation; that is, there would be no gap between theory and practice under this implementation. *Note however, that if the original problem had only pure control constraints, the new problem would have both state and control constraints.* Such implementations of control generation have been performed successfully for a number of problems; see Refs. [5, 46, 47].

VII. Example: Non-Eulerian Spacecraft Slew Control

In this section we demonstrate a successful experimental application of PS methods for slewing NPSAT1 (see Fig. 1). Numerical simulation results are discussed in Ref. [11]; here, we advance the technique to ground testing. An airbearing setup of the spacecraft for ground testing is shown in Fig. 7. NPSAT1 uses




Figure 7. Air bearing setup for the ground test of NPSAT1.

magnetic actuators and a pitch momentum wheel for attitude control. One experiment onboard the NPSAT1 spacecraft is to demonstrate in flight the application of the new principles developed in the previous sections. One component of this experiment is the performance of the minimum-time feedback controller when the pitch momentum wheel is caged or failed. Thus, the cost function is simply given by the transfer time while the attitude motion is governed by magnetic torquers.

A. Dynamical Model for the Airbearing Setup

Choosing the standard quaternion and body rates as the state variables, we have $x = (q, \omega) \in \mathbb{R}^7$, where

- $q = (q_1, q_2, q_3, q_4)$: quaternion of the body frame with respect to the laboratory frame,
- $\omega = (\omega_x, \omega_y, \omega_z)$: angular velocity of the body frame with respect to the laboratory frame expressed in the body frame.

We suppress many of the differential geometric concepts involved in the motion of a spacecraft as all of these notions have a differential algebraic counterpart. Thus, for example, when we write $q \in \mathbb{R}^4$, we do not imply that q takes values in all of \mathbb{R}^4 but that q is simply a vector of four variables.

It is quite straightforward²⁸ to show that the dynamical equations of motion for the NPSAT1 airbearing set up are given by,

$$\dot{q}_1(t) = \frac{1}{2} [\omega_x(t)q_4(t) - \omega_y(t)q_3(t) + \omega_z(t)q_2(t)] \quad (31)$$

$$\dot{q}_2(t) = \frac{1}{2} [\omega_x(t)q_3(t) + \omega_y(t)q_4(t) - \omega_z(t)q_1(t)] \quad (32)$$

$$\dot{q}_3(t) = \frac{1}{2} [-\omega_x(t)q_2(t) + \omega_y(t)q_1(t) + \omega_z(t)q_4(t)] \quad (33)$$

$$\dot{q}_4(t) = \frac{1}{2} [-\omega_x(t)q_1(t) - \omega_y(t)q_2(t) - \omega_z(t)q_3(t)] \quad (34)$$

$$\dot{\omega}_x(t) = \frac{I_2 - I_3}{I_1} \omega_y(t)\omega_z(t) - \frac{mgl}{I_1} C_{23}(q(t)) + \frac{1}{I_1} [B_z(q(t), t)u_2(t) - B_y(q(t), t)u_3(t)] \quad (35)$$

$$\dot{\omega}_y(t) = \frac{I_3 - I_1}{I_2} \omega_x(t)\omega_z(t) + \frac{mgl}{I_2} C_{13}(q(t)) + \frac{1}{I_2} [B_x(q(t), t)u_3(t) - B_z(q(t), t)u_1(t)] \quad (36)$$

$$\dot{\omega}_z(t) = \frac{I_1 - I_2}{I_3} \omega_x(t)\omega_y(t) + \frac{1}{I_3} [B_y(q(t), t)u_1(t) - B_x(q(t), t)u_2(t)] \quad (37)$$

where $(I_1, I_2, I_3) = (2.60, 2.87, 1.45) \text{ kg}\cdot\text{m}^2$ are the principal moments of inertia of the NPSAT1 airbearing assembly; $m = 59.0 \text{ kg}$ is the mass of the table platform; g is the gravitational constant; $l = 0.56 \text{ mm}$ is the distance from the center of mass of the table to the center of rotation; $C_{ij}(q)$ is the quaternion-parameterized ij -th element of the direction cosine matrix,

$$C(q) = \begin{bmatrix} q_1^2 - q_2^2 - q_3^2 + q_4^2, & 2(q_1q_2 + q_3q_4), & 2(q_1q_3 - q_2q_4) \\ 2(q_1q_2 - q_3q_4), & q_2^2 - q_1^2 - q_3^2 + q_4^2, & 2(q_2q_3 + q_1q_4) \\ 2(q_1q_3 + q_2q_4), & 2(q_2q_3 - q_1q_4), & q_3^2 - q_1^2 - q_2^2 + q_4^2 \end{bmatrix}$$

$(B_x(q, t), B_y(q, t), B_z(q, t))$ are the components of the Earth's magnetic field in the body frame,

$$\begin{pmatrix} B_x(q, t) \\ B_y(q, t) \\ B_z(q, t) \end{pmatrix} = C(q) \begin{pmatrix} B_1(t) \\ B_2(t) \\ B_3(t) \end{pmatrix} \quad (38)$$

and $(B_1(t), B_2(t), B_3(t))$ are the components of the Earth's magnetic field in the laboratory frame. A detailed description of the measurements of this field at the Naval Postgraduate School's NPSAT1 Laboratory are described in Ref. [48]. Based on these measurements the magnetic field was averaged to the constant values $(B_1(t), B_2(t), B_3(t)) = (0.24970, 0.09993, 0.39280) \text{ Gauss}$. That is, given the various uncertainties imbedded in the NPSAT1 ground test assembly, it was determined that no substantial advantages would be obtained by taking laboratory fluctuations in the magnetic field. For on-orbit applications, however, the B -field is taken to be time-varying and determined from the dipole model,

$$\begin{aligned} B_1(t) &= \frac{M_e}{r_0^3} [\cos(\omega_0 t) [\cos(\epsilon) \sin(i) - \sin(\epsilon) \cos(i) \cos(\omega_e t)] - \sin(\omega_0 t) \sin(\epsilon) \sin(\omega_e t)] \\ B_2(t) &= -\frac{M_e}{r_0^3} [\cos(\epsilon) \cos(i) + \sin(\epsilon) \sin(i) \cos(\omega_e t)] \\ B_3(t) &= \frac{2M_e}{r_0^3} [\sin(\omega_0 t) [\cos(\epsilon) \sin(i) - \sin(\epsilon) \cos(i) \cos(\omega_e t)] + 2 \cos(\omega_0 t) \sin(\epsilon) \sin(\omega_e t)] \end{aligned}$$

where $M_e = 7.943 \times 10^{15} \text{Wb.m}$ is the magnetic dipole moment of the Earth, $\epsilon = 11.7^\circ$ is the magnetic dipole tilt, $i = 35.4^\circ$ is the orbit inclination of NPSAT1, $\omega_e = 7.29 \times 10^{-5} \text{rad/s}$ is the spin rate of the Earth. Note also that there are other differences between the airbearing and flight models of NPSAT1; for instance, the mgl -terms in (35)-(37) are replaced by gravity gradient torques.^{11,28}

The controls, $(u_1, u_2, u_3) \in \mathbb{R}^3$ are the actuator dipole moments on NPSAT1. Clearly, the dynamics of NPSAT1 are quite complex with substantial nonlinearities.

B. State and Control Spaces

That the quaternion variables must lie on S^3 generates a state space given by,

$$\mathbb{X} = \{(q, \omega) \in \mathbb{R}^7 : \|q\|_2 = 1, q_4 \geq 0\}$$

The controls on NPSAT1 are bounded by the maximum available dipole moment; this generates a control space given by,

$$\mathbb{U} = \{u \in \mathbb{R}^3 : \|u\|_\infty \leq 33 \text{ A.m}^2\}$$

Thus, the NPSAT1 control system contains both state and control constraints.

In accordance with (1), we have provided a mathematical model for the dynamical system, obviously a nonlinear system. Furthermore, since the dipole moment is directly controlled by passing current to the solenoid, it is essentially inertialess; hence, $\mathcal{U} = L^\infty([t_0, t_f], \mathbb{R}^3)$. Thus, to each $u(\cdot) \in L^\infty$, f satisfies C^1 -Carathéodory conditions.

C. Slew Control Problem

As briefly noted in Sec. I, the transfer objective of NPSAT1 is to slew the spacecraft in minimum time. Thus, the cost function for slew is simply given by transfer time,

$$J_T[x(\cdot), u(\cdot), t_T] := t_T - t_0$$

The optimal control problem is to maneuver NPSAT1 in minimum time from a rest to rest state. A benchmark set of infinite-precision boundary conditions are given by,

$$\begin{aligned} [q_1(t_0), q_2(t_0), q_3(t_0), q_4(t_0)] &= [0, 0, \sin(\phi_0/2), \cos(\phi_0/2)] \equiv [q_1^0, q_2^0, q_3^0, q_4^0] \\ [\omega_x(t_0), \omega_y(t_0), \omega_z(t_0)] &= [0, 0, 0] \equiv [\omega_x^0, \omega_y^0, \omega_z^0] \\ [q_1(t_T), q_2(t_T), q_3(t_T), q_4(t_T)] &= [0, 0, \sin(\phi_f/2), \cos(\phi_f/2)] \equiv [q_1^f, q_2^f, q_3^f, q_4^f] \\ [\omega_x(t_T), \omega_y(t_T), \omega_z(t_T)] &= [0, 0, 0] \equiv [\omega_x^f, \omega_y^f, \omega_z^f] \end{aligned}$$

with $\phi_0 = 0$ and $\phi_f = 135^\circ$.

As noted in Sec. VII, the boundary conditions can only be attained within the context of sensor/estimation precision. Based on NPSAT1's sensor and estimation accuracies, the practical boundary conditions on NPSAT1 for a 135 degree slew are given by,

$$q_i(t_0) \in \mathbb{B}(q_i^0, 0.1\varepsilon_{q_i}) \quad q_i(t_f) \in \mathbb{B}(q_i^f, \varepsilon_{q_i}), \quad i = 1, 2, 3, 4 \quad (39)$$

$$\omega_j(t_0) \in \mathbb{B}(\omega_j^0, 0.1\varepsilon_\omega) \quad \omega_j(t_f) \in \mathbb{B}(\omega_j^f, \varepsilon_\omega), \quad j = x, y, z \quad (40)$$

where

$$(\varepsilon_{q_1}, \varepsilon_{q_2}, \varepsilon_{q_3}, \varepsilon_{q_4}) = (0.087, 0.087, 0.034, 0.081)$$

and $\varepsilon_\omega = 0.0025/s$. These numbers are based on the capability of sensors onboard NPSAT1 in detecting deviations from the target set of point conditions to about a tenth of the mission requirements. This is why the initial conditions are known to within the sensor capability while the final conditions are targeted to the specified requirements. Complete details on the sensors and filtering algorithms are discussed in Ref. [28].




Figure 8. PS-computed open-loop optimal controls, $t \mapsto u_i$, $i = 1, 2, 3$, for the NPSAT1 air bearing model.

D. Experimental Results

Shown in Fig. 8 are the time-optimal open loop controls generated for the NPSAT1 ground test model. The controls are optimal in the sense that they satisfy all the necessary conditions for optimality as demonstrated in Refs. [11] and [12]; hence, strictly speaking they are only extremals. For the purposes of brevity, we do not describe these optimality tests as they are extensively discussed elsewhere.^{11,12} The open-loop controls applied to the experimental set up is slightly shifted from the computed controls as implied in Fig. 8 because of various time-delay issues such as the wireless communication (see Fig. 7) to the onboard single board computer (SBC) on NPSAT1. The quaternion-responses of the spacecraft, both model-expected and experimentally achieved, are shown in Fig. 9. As expected, the open-loop controls drive the model to the

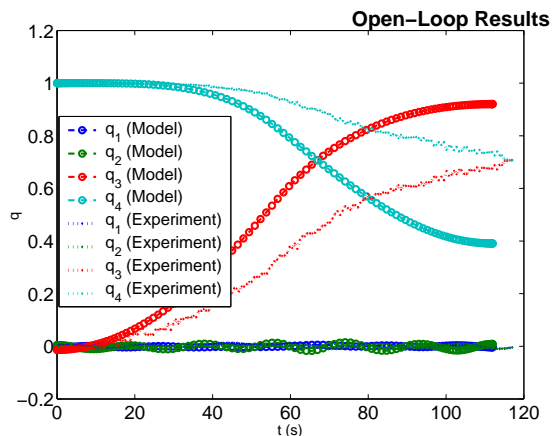


Figure 9. Open-loop quaternion trajectory, $t \mapsto q_i$, $i = 1, 2, 3, 4$, for the NPSAT1 air bearing model and experimental set up. The experimental response is measured from the laser diode shown in Fig. 7.

target conditions but the actual experiment fails to meet the requirement as a result of uncertainties and imperfections, the hallmark of any engineering set up.

The closed-loop control trajectories generated along the lines detailed in Sec. IV are shown in Fig. 10 along with the open-loop controls. Having interacted with the plant (i.e. experimental set up), the closed-loop control trajectories are indeed substantially different from the open-loop controls. The closed-loop response, $t \mapsto (q_3, q_4)$ is shown in Fig. 11. It is clear that, this time, the closed-loop response achieves all the required objectives as correlated by both the laser diode measurement as well as the unscented Kalman filter (UKF) estimates, the details of which are described in Ref. [28]. Note that the closed-loop response does not track a reference trajectory; rather, the closed-loop response simply meets all the performance specifications, and hence generates an autonomous trajectory. This is one of the many ways in which our control system design is quite different from those advocated in standard control theory; that is, we simply deal with the original




Figure 10. PS-computed closed-loop control trajectories, $t \mapsto u_i$, $i = 1, 2, 3$, for the NPSAT1 air bearing experimental set up; compare Fig. 6.

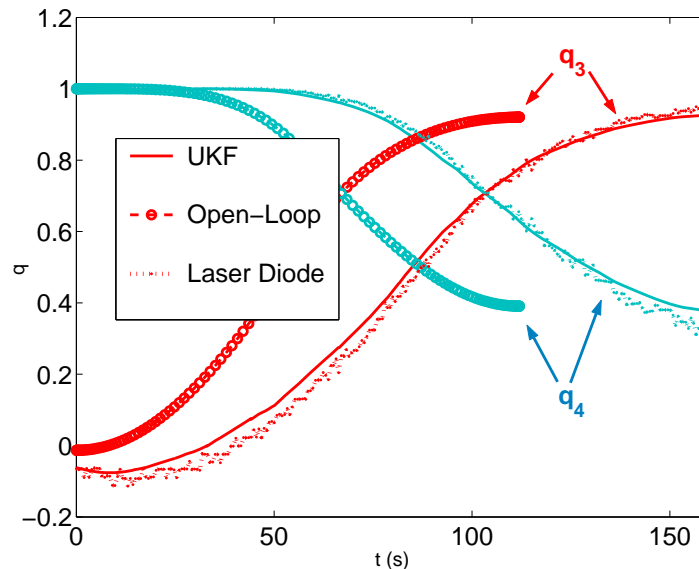


Figure 11. Closed-loop quaternion trajectory, $t \mapsto (q_3, q_4)$, for the NPSAT1 air bearing experimental set up. In addition to the laser diode measurement, the experimental response is also obtained from an implementation of an unscented Kalman filter (UKF); see Ref. [28] for measurement details.

specifications directly. Should the original specifications call for a tracking controller, this would be achieved through a proper problem definition.

The closed-loop response, $t \mapsto (q_1, q_2)$, shown in Fig. 12 requires some additional explanation. First, note that the UKF estimates do indicate that the target conditions are indeed obtained. The reason the laser diode measurements do not correlate with the UKF estimates is because the laser diode is accurate only up to 2.5° which maps to a 0.018 error ball in q_1 and 0.04 in q_2 ; hence, the laser diode outputs shown in Fig. 12 are really not trustworthy. On the other hand, the UKF measurements are indeed reliable as explained in detail in Ref. [28]. These two plots illustrate some of the points argued in Secs. III and VI on the merging of “epsilon balls” in theory, computation and experiment.

As a matter of completeness, the angular velocity response of the closed-loop system is shown in Fig. 13. It is clear that the final rest conditions are indeed obtained well within the requirements of the ω epsilon ball. Finally, as a point of further validation of the experimental results, the online sensing and estimation of the local magnetic field is shown in Fig. 14. In addition to validating the excellent performance of the UKF, this plot illustrates how magnetic sensing and actuation could be performed at the same time.

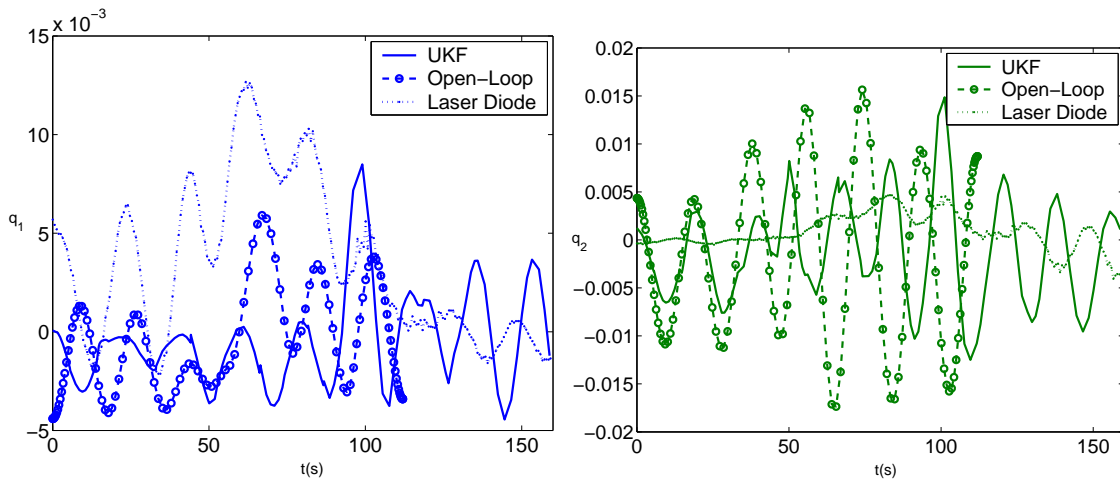


Figure 12. Closed-loop quaternion trajectory, $t \mapsto (q_1, q_2)$, for the NPSAT1 air bearing experimental set up. Here, the UKF outputs are more accurate estimates of the states than the noisy measurements of the laser diode; see Ref. [28] for measurement details.




Figure 13. Closed-loop angular velocity trajectory components, $t \mapsto w_i, i = 1, 2, 3$, for the NPSAT1 air bearing experimental set up. The open-loop trajectory is the response of the mathematical model. All control system specifications are met without resorting to tracking model trajectories.

E. Sample and Hold Feedback: A Brief Discussion on Experimental Results

PS methods are a means to generate clock-based feedback controls that result in Carathéodory- π solutions. Since Carathéodory- π solutions are generalizations of the π solutions, PS methods may also be used to implement a sample and hold feedback controller. The purpose of performing this exercise to illustrate the differences in performances between clock-based and sample-and-hold feedback when everything else is the same. To this end, we first note that for a sufficiently fast computational time, there ought to be no substantive differences between the π and the Carathéodory- π solutions, all other things being equal. This analogy is similar to the expectation that Euler solutions should be the same as Runge-Kutta solutions, provided everything else is equal. Even under such equivalence, it is customary to prefer Runge-Kutta integration simply because it is more efficient, and has more desirable stability and convergence properties.




Figure 14. Sensed (dashed line) and UKF-estimated (solid line) local magnetic field during closed-loop control of NPSAT1.

With this principle in mind, it is worth comparing the practical performances of the clock-based and sample-and-hold feedback controllers on the NPSAT1 testbed when both feedback controllers are generated using the same PS method. The experimental result of the system response to a sample and hold feedback controller is shown in Fig. 15. It is quite clear that this controller, despite being PS-generated, has failed. Although




Figure 15. Failed performance of NPSAT1 to a sample and hold feedback control.

not shown here for the purpose of brevity, it can be numerically demonstrated that a sampled data feedback controller would perform successfully if the computation time were to be substantially faster. Thus, the requirements on real-time computation are indeed high when the feedback is in a sample and hold form, but ironically, the computational “burden” on our proposed clock-based feedback implementation is indeed low and theoretically justified by Lemma 1 and the arguments of Sec. V. In connecting our theory with experimental results, the closed-loop trajectory shown in Fig. 15 is a π solution while results reported in

Sec. VII.E are Carathéodory- π solutions. As an illustrative point of comparison, we conclude this section with Fig. 16 wherein a sample-and-hold implementation for $t \mapsto u_1$ is plotted alongside a clock-based feedback implementation. Comparing the plots, the stark differences are readily apparent.

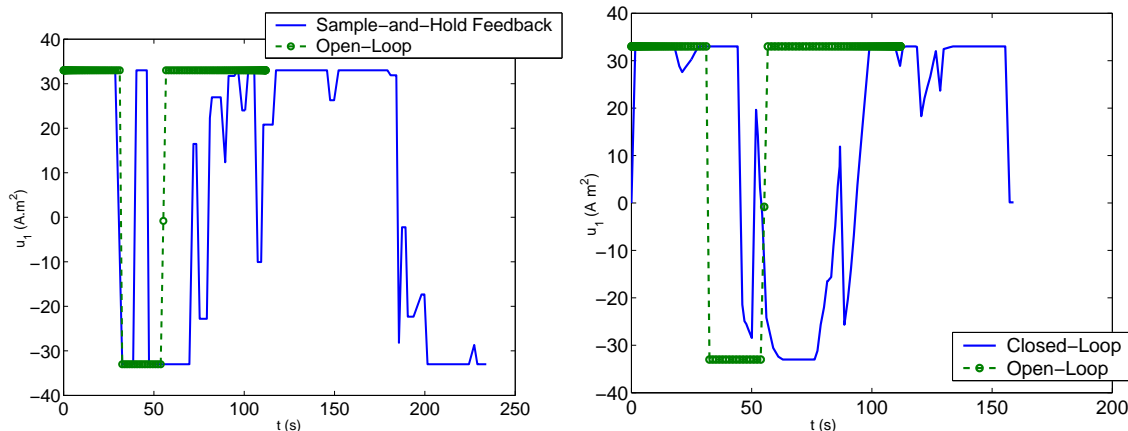


Figure 16. Failed sample and hold (left) and successful clock-based (right) feedback controls generated using PS methods.

VIII. A Plethora of Other Examples

It is quite clear from Sec. VII that sophisticated PS feedback control laws are not only possible, but that they are practically implementable. Although the experimental results reported in Sec. VII are the first of its kind, there is substantial evidence to indicate that these ideas are readily portable to virtually any nonlinear system that satisfies the mild conditions indicated in Sec. I. For example, in Ref. [39], optimal solutions to a robotics problem was obtained at approximately 30Hz frequency. Feedback solutions (in the sense of Carathéodory- π) to a number of problems that include the classically difficult problem of swing up of an inverted pendulum are presented in Ref. [6]. Reentry guidance for an X-33 class launch vehicle based on the same principles described in this paper is discussed in Ref. [5]. Solutions that satisfy the real-time computational conditions set forth in Sec. V have been reported for such diverse problems as spacecraft formation design and control,⁴⁹ low-thrust trajectory optimization,⁵⁰ singularity-free steering of control moment gyros,⁵¹ and loitering of unmanned air vehicles⁵² to name just a few. Almost all of the computations reported in these papers and elsewhere were performed with un-optimized versions of the DIDO software package⁴¹ running in a MATLAB[©] environment under MS Windows[©] operating systems. Despite such run-time overhead, PS controls can be computed in real-time for many systems. All of these recent results and the experimental demonstrations reported in this paper show that real-time generation of optimal controls for complex nonlinear systems are demonstrable under adverse conditions. Under such a framework, the new main problems in control theory are sharply focused at its very foundations, as theory and computation continue to merge at the level of first principles itself. An excellent overview of such fundamental theoretical problems are described by Clarke;¹⁶ see also Ref. [2].

IX. Conclusions

When open-loop controls can be generated in real time, it implies feedback control. Such feedback controls are nontraditional in the sense that they do not have a recognizable analytical representation, $x \mapsto k(x)$. Consequently, traditional analysis based on explicit Lyapunov functions cannot be carried out. A new type of analysis is necessary wherein closed-loop performance can be guaranteed to meet design specifications. This analysis lies at the intersection of optimal control theory, approximation theory, and modern computational principles. Recent advances in sample and hold feedback solutions reveal that the concept avoids many theoretical pitfalls associated in defining solutions to nonsmooth differential equations arising from nonsmooth feedback maps; however, their practical implementation requires a very fast computation of open loop controls. By introducing a Carathéodory- π solution concept, we keep or extend the theoretical benefits

of a sample-and-hold controller while alleviating the computational burden. Thus, the proposed pseudospectral (PS) feedback controllers that generate Carathéodory- π solutions are implementable for systems with fast dynamics. Although mathematically sophisticated, PS controllers are easy to implement. Given that they can demonstrably generate solutions in fractions of a second to a few seconds, it is easy to conclude that a unified approach to effective feedback control is indeed possible for a rich variety of dynamical systems. In addition, to their effective implementation by PS methods, the clock-based feedback controllers have a built-in safety factor. Ground test results for the NPSAT1 spacecraft indicate a successful implementation of these concepts. Flight tests are scheduled after the launch of NPSAT1.

X. Acknowledgments

We gratefully acknowledge funding for this research provided in part by the Secretary of the Air Force and the Air Force Office of Scientific Research under AFOSR Grant F1ATA0-60-6-2G002. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

We are particularly grateful to Jim Horning, Ron Phelps, and Dan Sakoda for their expert engineering assistance in the experimental set up. This project could not have been completed without their help.

XI. Appendix

A. Carathéodory Solution^{21, 53}

An initial value (Cauchy) problem,

$$\dot{x} = g(x, t), \quad x(t_0) = x^0$$

where $g(x, t)$ is a time-varying vector field, $g : \mathbb{R}^{N_x} \times \mathbb{R} \rightarrow \mathbb{R}^{N_x}$, is said to have a Carathéodory solution if there exists an absolutely continuous function, $x^*(\cdot) : R \supset [t_0, t_f] \rightarrow \mathbb{R}^{N_x}$, such that

$$\dot{x}^*(t) = g(x^*(t), t) \text{ for a.a. } t \in [t_0, t_f] \quad \text{and} \quad x^*(t_0) = x^0$$

B. L^1 -Carathéodory Conditions⁵³

Let $g(x, t)$ be a time-varying vector field, $g : \mathbb{R}^{N_x} \times \mathbb{R} \rightarrow \mathbb{R}^{N_x}$, defined over a tube,

$$\mathbb{T}(\xi, \epsilon) := \{(x, t) : t_0 \leq t \leq t_f, \|x - \xi(t)\| \leq \epsilon\}$$

around an absolutely continuous function, $\xi : [t_0, t_f] \rightarrow \mathbb{R}^{N_x}$. If $g(x, t)$ is

1. *measurable* in t for each x ,
2. *continuous* in x for almost all t , and
3. $\|g(x, t)\| \leq \varrho(t)$ for each x and almost all t and some $\varrho \in L^1([t_0, t_f], \mathbb{R}^{N_x})$,

then g is said to satisfy L^1 -Carathéodory conditions.

C. Carathéodory Existence Theorem^{21, 53}

If g is L^1 -Carathéodory, then there exists a Carathéodory solution for the initial value problem, $\dot{x} = g(x, t)$, $x(t_0) = x^0$ if $(x_0, t_0) \in \mathbb{T}(\xi, \epsilon)$.

D. C^1 -Carathéodory Conditions³¹

Let g be L^1 -Carathéodory. If

1. $g(x, t)$ is C^1 in x for almost all t ,
2. $\partial_x g(x, t)$ is measurable in t for each x , and
3. $\|g(x, t)\| + \|\partial_x g(x, t)\| \leq \varrho(t)$ for all $(x, t) \in \mathbb{T}(\xi, \epsilon)$

then g is said to satisfy C^1 -Carathéodory conditions.

E. The Lambert W function⁵⁴

The multi-valued function, $\mathbb{R} \ni x \rightarrow W(x)$, given implicitly by

$$x = W(x)e^{W(x)} \quad (41)$$

is called the Lambert W function. A detailed description of this function along with its historical origins and many applications are described in Ref. [54]. For $x \geq 0$, $W(x)$ is single-valued.

F. Gronwall's Lemma⁵⁵

Let $[t_0, t_f] \mapsto y(t) \in \mathbb{R}$ be an integrable function that satisfies Gronwall's inequality,

$$y(t) \leq a(t) + \int_{t_0}^t b(s)y(s) ds$$

where a and b are continuous, nonnegative, bounded functions with $t \mapsto a(t)$ nondecreasing over the interval, $[t_0, t_f]$; then,

$$y(t) \leq a(t)e^{B(t)}$$

where

$$B(t) := \int_{t_0}^t b(s) ds$$

References

- ¹Vinter, R. B., *Optimal Control*, Birkhäuser, Boston, MA, 2000.
- ²Clarke, F. H., Ledyae, Y. S., Stern, R. J., and Wolenski, P. R., *Nonsmooth Analysis and Control Theory*, Springer-Verlag, New York, NY, 1998.
- ³Clarke, F. H., *Optimization and Nonsmooth Analysis*, SIAM Publications, Philadelphia, PA, 1990.
- ⁴Shaffer, P. J., Ross, I. M., Oppenheimer, M. W. and Doman, D. B., "Optimal Trajectory Reconfiguration and Retargeting for a Reusable Launch Vehicle," *AIAA Guidance, Navigation and Control Conference*, San Francisco, CA, 2005.
- ⁵Bollino, K., Ross, I.M. and Doman, D., "Optimal Nonlinear Feedback Guidance for Reentry Vehicles," *AIAA Guidance, Navigation and Control Conference*, Keystone, CO, 2006. AIAA-2006-6074.
- ⁶Ross, I. M., Gong, Q., Fahroo, F. and Kang, W., "Practical Stabilization Through Real-Time Optimal Control," *2006 American Control Conference*, Minneapolis, MN, June 14-16 2006.
- ⁷Fahroo F., and Ross, I. M., "Pseudospectral Methods for Infinite-Horizon Nonlinear Optimal Control Problems," *Proc. AIAA Guid., Nav. and Control Conf.*, San Francisco, CA, August 15-18, 2005.
- ⁸Ross, I. M., "Space Trajectory Optimization and L^1 -Optimal Control Problems," *Modern Astrodynamics*, edited by P. Gurfil, Elsevier, St. Louis, MO, 2006, Chapter 6, pp. 155-188.
- ⁹Lu, P., Sun, H. and Tsai, B., "Closed-Loop Endoatmospheric Ascent Guidance," *Journal of Guidance, Control and Dynamics*, Vol.26, No. 2, 2003, pp.283-294.
- ¹⁰Rea, J., "Launch Vehicle Trajectory Optimization Using a Legendre Pseudospectral Method," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, Austin, TX, August 2003. Paper No. AIAA 2003-5640.
- ¹¹Sekhavat, P., Fleming A., and Ross, I. M., "Time-Optimal Nonlinear Feedback Control for the NPSAT1 Spacecraft," *Proceedings of the 2005 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, AIM 2005, 24-28 July 2005 Monterey, CA.
- ¹²Fleming, A. "Real-time optimal slew maneuver design and control," Astronautical Engineer's Thesis, US Naval Postgraduate School, December 2004.
- ¹³Fahroo, F., Doman, D. and Ngo, A., "Modeling Issues in Footprint Generation for Reusable Launch Vehicles," *Proceedings of the 2003 IEEE Aerospace Conference*, Big Sky, MT, March 2003.
- ¹⁴Fahroo, F., Doman, D. and Ngo, A., "Footprint Generation for Reusable Launch Vehicles Using a Direct Pseudospectral Method," *Proceedings of the American Control Conference*, Denver, CO June 2003.
- ¹⁵Fahroo, F. and Doman, D., "A Direct Method for Approach and Landing Trajectory Reshaping with Failure Effect Estimation," *2004 AIAA Conference on Guidance, Navigation and Control*, August 2004, Providence, RI.
- ¹⁶Clarke, F., "Lyapunov functions and feedback in nonlinear control," *Optimal Control, Stabilization and Nonsmooth Analysis*, M.S. de Queiroz, M. Malisoff, P. Wolenski (Eds.), Lecture Notes in Control and Information Sciences Vol. 301, pp. 267-282, Springer-Verlag, 2004.
- ¹⁷Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., and Mischenko, E. F., *The Mathematical Theory of Optimal Processes*, Wiley-Interscience, New York, N.Y., 1962.
- ¹⁸Crandall, M. G. and Lions, P.-L., "Viscosity Solutions of Hamilton-Jacobi Equations," *Transactions of the American Mathematical Society*, No. 277, 1983, pp. 1-42.
- ¹⁹Bardi, M. and Capuzzo-Dolcetta, I., *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Birkhäuser, Boston, MA, 1997.
- ²⁰Brockett, R., "Asymptotic Stability and Feedback Stabilization," *Differential Geometric Control Theory*, R. Brockett, R. Millman, and H. Sussmann, eds., Birkhäuser, Boston, 1983, pp. 181191.
- ²¹Young, L. C., *Lectures on the Calculus of Variations and Optimal Control Theory*, Saunders, Philadelphia, PA, 1969.
- ²²Filippov, A. F., *Differential Equations with Discontinuous Right-Hand Side*, D. Reidel, Boston, 1989.

- ²³Ryan, E. P., "On Brockett's Condition for Smooth Stabilizability and its Necessity in a Context of Nonsmooth Feedback," *SIAM Journal of Control and Optimization*, Vol. 32, 1994, pp. 1597-1604.
- ²⁴Coron, J.-M. and Rosier, L., "A Relation Between Continuous Time-Varying and Discontinuous Feedback Stabilization," *Journal of Math. System, Estimation, and Control*, No. 4, 1994, pp. 67-84.
- ²⁵Piccoli, B., and Sussmann, H. J., "Regular Synthesis and Sufficiency Conditions for Optimality," *SIAM Journal of Control Optimization*, Vol. 39, No. 2, 2000, pp. 359-410.
- ²⁶Ross, I. M., "A Roadmap for Optimal Control: The Right Way to Commute," *Annals of the New York Academy of Sciences*, Vol. 1065, New York, N.Y., January 2006.
- ²⁷Gong, Q., Ross, I. M. and Kang, W., "A Pseudospectral Observer for Nonlinear Systems," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, August 2005, San Francisco, CA.
- ²⁸Sekhvat, P., Gong, Q. and Ross, I. M., "Unscented Kalman Filtering: NPSAT1 Ground Test Results," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 21-24 August 2006, Keystone, CO.
- ²⁹Weiss, L. and Infante, E. F., "Finite Time Stability Under Perturbing Forces and on Product Spaces," *IEEE Transaction on Automatic Control*, Vol. AC-12, No. 1, February 1967, pp. 54-59.
- ³⁰Dorato, P., "Comment on 'Finite Time Stability Under Perturbing Forces and on Product Spaces,'" *IEEE Transaction on Automatic Control*, Vol. AC-12, June 1967, pp. 340.
- ³¹Sussmann, H. J., "Geometry and Optimal Control," *Mathematical Control Theory*, J. Baillieul and J. C. Willems, Eds., Springer-Verlag, 1998, pp. 140-198.
- ³²Fahroo, F. and Ross, I. M., "On Discrete-Time Optimality Conditions for Pseudospectral Methods," *Proceedings of the AIAA/AAS Astrodynamics Conference*, Keystone, CO, August 2006. AIAA-2006-6304.
- ³³Elnagar, J., Kazemi, M. A. and Razzaghi, M., "The Pseudospectral Legendre Method for Discretizing Optimal Control Problems," *IEEE Transactions on Automatic Control*, Vol. 40, No. 10, 1995, pp. 1793-1796.
- ³⁴Ross, I. M., and Fahroo, F., "Legendre Pseudospectral Approximations of Optimal Control Problems," *Lecture Notes in Control and Information Sciences*, Vol. 295, Springer-Verlag, New York, 2003.
- ³⁵Fahroo, F. and Ross, I. M., "Costate Estimation by a Legendre Pseudospectral Method," *Journal of Guidance, Control and Dynamics*, Vol. 24, No. 2, 2001, pp. 270-277.
- ³⁶Ross, I. M. and Fahroo, F., "Issues in the Real-Time Computation of Optimal Control," *Mathematical and Computer Modelling*, An International Journal, Vol. 43, Issues 9-10, May 2006, pp.1172-1188. (Special Issue: Optimization and Control for Military Applications).
- ³⁷Gong, Q., Kang, W. and Ross, I. M., "A Pseudospectral Method for the Optimal Control of Constrained Feedback Linearizable Systems," *IEEE Transactions on Automatic Control*, Vol. 51, No. 7, July 2006, pp. 1115-1129.
- ³⁸Trefethen, L. N., *Spectral Methods in MATLAB*, SIAM, Philadelphia, PA, 2000.
- ³⁹Ross, I. M. and Fahroo, F., "A Unified Framework for Real-Time Optimal Control," *Proceedings of the IEEE Conference on Decision and Control*, Maui, December, 2003.
- ⁴⁰Gong, Q., Ross, I. M., Kang, W. and Fahroo, F., "On the Pseudospectral Covector Mapping Theorem for Nonlinear Optimal Control," to appear in *Proceedings of the the 45th IEEE Conference on Decision and Control*, San Diego, CA, December 13-15, 2006.
- ⁴¹Ross, I. M., *User's Manual for DIDO: A MATLAB Application Package for Solving Optimal Control Problems*, Tomlab Optimization, Sweden, February 2004.
- ⁴²Hager, W. W., "Numerical Analysis in Optimal Control," *International Series of Numerical Mathematics*, Hoffmann, K.-H. Lasiecka, I. Leugering, G., Sprekels, J., and Troeltzsch, F., Eds., Birkhäuser, Basel, Switzerland, 2001, Vol. 139, pp. 83-93.
- ⁴³Mordukhovich, B. S., *Variational Analysis and Generalized Differentiation, I: Basic Theory*, vol. 330 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences] Series, Springer, Berlin, 2005.
- ⁴⁴Mordukhovich, B. S., *Variational Analysis and Generalized Differentiation, II: Applications*, vol. 331 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences] Series, Springer, Berlin, 2005.
- ⁴⁵Ross, I. M., "A Historical Introduction to the Covector Mapping Principle," *AAS/AIAA Astrodynamics Specialist Conference*, Tahoe, NV, August 8-11, 2005, Paper AAS 05-332.
- ⁴⁶Josselyn S. and Ross, I. M., "A Rapid Verification Method for the Trajectory Optimization of Reentry Vehicles," *Journal of Guidance, Control and Dynamics*, Vol. 26, No. 3, 2003.
- ⁴⁷Stevens, R. and Ross, I. M., "Preliminary Design of Earth-Mars Cyclers Using Solar Sails," *Journal of Spacecraft and Rockets*, Vol. 41, No. 4, 2004.
- ⁴⁸Herbert, E. W., "NPSAT1 Magnetic Attitude Control System Algorithm Verification, Validation and Air-Bearing Tests," M.S. Thesis in Electrical Engineering, Naval Postgraduate School, September 2004.
- ⁴⁹Infeld, S. I., Josselyn, S. B., Murray, W. and Ross, I. M., "Design and Control of Libration Point Spacecraft Formations," *Journal of Guidance, Control and Dynamics*, to appear.
- ⁵⁰Ross, I. M., Gong, Q. and Sekhavat, P., "A Simple Technique for Low-Thrust High-Accuracy Trajectory Optimization," *16th AAS/AIAA Space Flight Mechanics Conference*, Tampa, FL, January 22-26, 2006. AAS 06-150.
- ⁵¹Fleming, A. and Ross, I. M., "Singularity-Free Optimal Steering of Control Moment Gyros," *Proceedings of the 2005 AAS/AIAA Astrodynamics Specialist Conference*, August 7-11, 2005 Lake Tahoe, CA. AAS 05-418.
- ⁵²Harada, M., Bollino, K., and Ross, I. M., "Minimum-Fuel Circling for an Unmanned Aerial Vehicle," *2005 JSASS-KSAS Joint International Symposium on Aerospace Engineering*, Nagoya, Japan, October 12-15, 2005.
- ⁵³Biles, D. C. and Binding, P. A., "On Carathéodory's Conditions for the Initial Value Problem," *Proceedings of the American Mathematical Society*, Volume 125, Number 5, May 1997, Pages 1371-1376.
- ⁵⁴Corless, R. M., Gonnet, G. H., Hare, D.E.G., Jeffrey, D. J., and Knuth, D. E., "On the Lambert W Function," *Advances in Computational Mathematics*, Vol. 5, 1996, pp. 329-359.
- ⁵⁵Zabczyk, J., *Mathematical Control Theory: An Introduction*, Birkhäuser, Boston, 1992.