

Engineering Notes

ENGINEERING NOTES are short manuscripts describing new developments or important results of a preliminary nature. These Notes should not exceed 2500 words (where a figure or table counts as 200 words). Following informal review by the Editors, they may be published within a few months of the date of receipt. Style requirements are the same as for regular contributions (see inside back cover).

Convergence of the Costates Does Not Imply Convergence of the Control

Fariba Fahroo* and I. Michael Ross†

Naval Postgraduate School, Monterey, California 93943

DOI: 10.2514/1.37331

I. Introduction

SOLVING an optimal control problem using a digital computer implies discrete approximations. Since the 1960s, there have been well-documented [1–3] naïve applications of Pontryagin’s principle in the discrete domain. Although its incorrect applications continue to this day, the origin of the naïveté is quite understandable because one has a reasonable expectation of the validity of Pontryagin’s principle within a discrete domain. That an application of the Hamiltonian minimization condition is not necessarily valid in a discrete domain [1,4] opens up a vast array of questions in theory and computation [2,5]. These questions continue to dominate the meaning and validity of discrete approximations and computational solutions to optimal control problems [6–10]. Among these questions is the convergence of discrete approximations in optimal control.

About the year 2000, there were a number of key discoveries on the convergence of discrete approximations [9,11–14]. Among other things, Hager [9] showed, by way of a counterexample, that a convergent Runge–Kutta (RK) method may not converge. This seemingly contradictory result is actually quite simple to explain [10]: long-established convergence results on ordinary differential equations do not necessarily apply for optimal control problems. Thus, an RK method that is convergent for an ordinary differential equation may not converge when applied to an optimal control problem. Not only does this explain the possibility of erroneous results obtained through computation, it also explains why computational optimal control has heretofore been such a difficult problem. The good news is that if a proper RK method is used (those developed by Hager [9]), convergence can be assured under a proper set of conditions.

Whereas RK methods have a long history of development for ordinary differential equations, pseudospectral (PS) methods have had a relatively short history of development for optimal control. In parallel to Hager’s [9] discovery on RK methods, recent developments [8,15–17] show that the convergence theory for PS approximations in optimal control is sharply different from that used in solving partial differential equations. Furthermore, the

convergence theory for PS approximations is also different from the one used in RK approximations to optimal control.

A critical examination of convergence of approximations using the new theories developed in recent years has not only begun to reveal the proper computational techniques for solving optimal control problems, it has also exposed the fallacy of long-held tacit assumptions. For instance, Ross [18] showed, by way of a simple counterexample, that an *indirect method* generates the wrong answer, whereas a *direct method* generates the correct solution. *This counterexample exposed the fallacy of the long-held belief that indirect methods are more accurate than direct methods.*

In this Note, we show, by way of another counterexample, that the convergence of the costates does not imply convergence of the control. This result appears to have more impact on the convergence of PS approximations in optimal control than the convergence of RK approximations because of the significant differences between the two theories; consequently, we restrict our attention to the impact of this result on PS methods, noting, nonetheless, the generality of this assertion.

II. Counterexample

Consider the following optimal control problem proposed by Gong et al. [15]:

$$\mathbf{x} = (x, v) \in \mathbb{R}^2, \quad u \in \mathbb{R}$$

$$(G) \left\{ \begin{array}{l} \text{Minimize} \quad J[\mathbf{x}(\cdot), u(\cdot)] = \int_0^1 v(t)u(t) dt \\ \text{Subject to} \quad \dot{x}(t) = v(t) \\ \quad \quad \quad \dot{v}(t) = -v(t) + u(t) \\ \quad \quad \quad v(t) \geq 0 \\ \quad \quad \quad 0 \leq u(t) \leq 2 \\ \quad \quad \quad (x(0), v(0)) = (0, 1) \\ \quad \quad \quad (x(1), v(1)) = (1, 1) \end{array} \right.$$

This problem describes a particle of unit mass ($m = 1$) moving in a linear resistive medium ($\alpha = -1$), where x is the position, v is the velocity, and u is the applied force (see Fig. 1). The optimal control problem is to minimize the total amount of work done. By a direct application of the minimum principle, it can be easily shown that the exact values of the costates are given by

$$\lambda_x(t) = -2 \quad (1)$$

$$\lambda_v(t) = -1 \quad (2)$$

and the exact optimal primal solution is given by

$$x^*(t) = t \quad (3)$$

$$v^*(t) = 1 \quad (4)$$

$$u^*(t) = 1 \quad (5)$$

Thus, the exact optimal cost is given by

$$J[\mathbf{x}^*(\cdot), u^*(\cdot)] = 1 \quad (6)$$

Received 27 February 2008; revision received 2 May 2008; accepted for publication 6 May 2008. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/08 \$10.00 in correspondence with the CCC.

*Professor, Department of Applied Mathematics; ffahroo@nps.edu. Associate Fellow AIAA.

†Professor, Department of Mechanical and Astronautical Engineering; imross@nps.edu. Associate Fellow AIAA.




Fig. 1 Schematic for problem G.

An application of two Legendre PS methods [19,20] based on Gauss and Gauss–Radau points generates the states and costates shown in Figs. 2 and 3, respectively. Clearly, both the states and costates have converged. Despite this strong result, the optimal control does not converge, as shown in Fig. 4.

III. Explaining the Result

Let $\mathbf{y} = (\mathbf{x}, \boldsymbol{\lambda})$ be the state–costate pair. Then, for a general optimal control problem, under appropriate conditions, the state–costate dynamics can be written as

$$\dot{\mathbf{y}} = \mathbf{g}(\mathbf{y}, \mathbf{u}) \quad (7)$$

For the purposes of brevity, we use the notation L^1 and L^∞ to denote Lebesgue spaces in a finite domain Ω , rather than the more elaborate notations $L^1(\Omega)$ and $L^\infty(\Omega)$. Under appropriate technical assumptions [known as Carathéodory conditions (see [21,22])], we can guarantee that if $\mathbf{u}(\cdot) \in L^\infty$, then there exists an absolutely continuous solution $\mathbf{y}(\cdot)$. This is part of the foundations on which Pontryagin’s theory rests. What is not true is the converse. That is, if $\mathbf{y}(\cdot)$ is absolutely continuous, there is no guarantee that $\mathbf{u}(\cdot)$ will be in L^∞ ; however, we can guarantee that $\mathbf{u}(\cdot)$ will be in L^1 , under appropriate technical assumptions. Because L^1 is a different and larger space than L^∞ , it is possible to have two control functions, $\mathbf{u}_a(\cdot)$ and $\mathbf{u}_b(\cdot)$, such that for any small $\varepsilon > 0$ and any large $M > 0$, the following holds:

$$\|\mathbf{u}_a(\cdot) - \mathbf{u}_b(\cdot)\|_{L^1} < \varepsilon, \quad \text{but } \|\mathbf{u}_a(\cdot) - \mathbf{u}_b(\cdot)\|_{L^\infty} > M \quad (8)$$

That is, two functions can be very close together in the L^1 norm but very far apart in the L^∞ metric. Thus, when the states and costates converge absolutely, we can only guarantee that the controls converge in L^1 but not in the L^∞ norm. This type of weak convergence appears to have occurred in Figs. 2–4, as we will now verify.

By inspection of Figs. 2 and 3, it is clear that the states and costates are converging in the strong L^∞ norm. To determine the appropriate way in which the controls are converging, we first note that

$$\|u^*(\cdot)\|_{L^1} = 1 \quad \text{exactly} \quad (9)$$

Computing $\|u^N(\cdot)\|_{L^1}$ using the Gauss and Gauss–Radau quadrature formulas for the respective Legendre PS methods yields

$$\|u_{\text{Gauss}}^{30}(\cdot)\|_{L^1} = 0.9946, \quad \|u_{\text{Gauss–Radau}}^{30}(\cdot)\|_{L^1} = 0.9962 \quad (10)$$

In other words, although the controls at the Gauss and Gauss–Radau points have not converged in the L^∞ norm, they do seem to be converging weakly in the L^1 topology. To further test this hypothesis, we compute $\|u^N(\cdot)\|_{L^1}$ for various values of N ; the result of this exercise is shown in Fig. 5. It is clear that $\|u^N(\cdot)\|_{L^1}$ for both PS methods are converging toward 1 for increasing N .

Although such a mathematical explanation may be sufficient to explain a numerical result, what is desirable from an engineering view point is a possible means to fix the problem. That is, from an engineering view point, what is needed is a computational tool that can provide convergence of the control in the L^∞ norm. This can be done by solving the problem using the Legendre PS method based on Gauss–Lobatto points.

IV. Fixing the Problem

In recognizing that a PS method for solving optimal control problems requires a new approach for convergence analysis, we




Fig. 2 PS state trajectory for problem G for two flavors of a PS method.




Fig. 3 PS costate trajectory for problem G for two flavors of a PS method.




Fig. 4 Exact and PS controls to problem G for two flavors of a PS method.




Fig. 5 Verification of the weak convergence of $u^N(\cdot)$ for problem G.

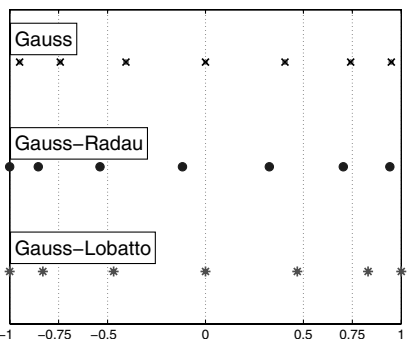


Fig. 6 Legendre quadrature nodes for $N = 7$.

developed [19] a unified theory for PS methods. Among other things, this theory indicates [20] that *the proper PS method for solving generic finite-horizon optimal control problems must be based on Gauss-Lobatto points*. A complete explanation of this point is beyond the scope of this Note, but we briefly provide some heuristics here for the purpose of completeness while noting that a detailed mathematical explanation is provided in [20].

Under the absence of any preselection of a set of interpolation points, the best quadrature nodes are the Gauss points. If any one point is arbitrarily preselected (say, -1), Gauss points are no longer the best points; the new set of best points are Gauss-Radau points. This fact is illustrated in Fig. 6, in which best is defined in the sense of the L^2 error norm. In these particular cases, such quadrature nodes are all known as Legendre quadrature nodes because all such points are based on Legendre polynomials; consequently, *all such PS methods are different flavors of the Legendre PS method*. In any event, if two points are preselected for interpolation (say, -1 and

$+1$), then neither the Gauss nor the Gauss-Radau points are the best points; the new set of best points are the Gauss-Lobatto points. Thus, the natural selection of quadrature nodes for solving boundary-value problems are Gauss-Lobatto points. Because an optimal control problem is fundamentally a boundary-value problem, that is, conditions are specified at the boundary points -1 and $+1$, it is clear that the Gauss-Lobatto points are the most natural choice of quadrature points. This is why if one chooses another PS method such as the Chebyshev PS method, for instance, then it is also based on Gauss-Lobatto points or, more precisely, the Chebyshev-Gauss-Lobatto points, which are different from the Legendre-Gauss-Lobatto points because they are based on minimizing the L^∞ error norm. What happens then if a PS method is naively based on, say, Gauss-Radau points or just Gauss points? According to our theory [19,20], convergence is not guaranteed. Just as in Hager's theory for RK methods for optimal control, in our theory for PS methods, an application of the wrong PS method does not guarantee a wrong solution; rather, it is quite possible to obtain the correct solution by a wrong PS method on one or more problems. But because a proof of convergence of a method based on select problems is not proof at all, what is indeed rigorous is just one counterexample to prove a negative. This is precisely the role problem G has served for the Legendre PS method based on Gauss and Gauss-Radau points.

A solution to problem G using the Legendre PS method based on Gauss-Lobatto points is shown in Fig. 7. It is clear that this Legendre PS method has indeed found the right answer in the right norm.

It is important to note that a complete theory for discrete approximations to optimal control problems is far from complete. Although a number of important advances have been made in recent years, development of convergence theories is an active research topic [8,10,15,17]. In developing a theory, it is quite important to know the parameters that bound the validity of the assumptions. What this counterexample has shown is that regardless of the method used, convergence of the costates does not imply convergence of the control, but convergence of the control does indeed imply convergence of the state-costate pair under mild assumptions.

V. Impact on PS Approximations

Given that PS methods have now progressed to flight applications [23–25], it is especially important that convergence issues for pseudospectral methods be critically examined. As alluded to in Sec. I, convergence of PS approximations for optimal control simply cannot be assumed on the grounds that the method has been shown to converge in other areas of mathematics. This is why the development of convergence results for PS approximations is an active field [8,15,17]. Furthermore, these new convergence results have been successfully used in the development of spectral [26] and guess-free algorithms [27] for solving optimal control problems. Such algorithms may fail or perform poorly when incorrect PS methods are used. For instance, it is clear from Fig. 5 that even though $\|u^N(\cdot)\|_{L^1}$ converges, it does not converge at a spectral rate; that is, Gauss-Radau and Gauss points do not even offer weak spectral convergence. In sharp contrast, the Legendre PS method based on

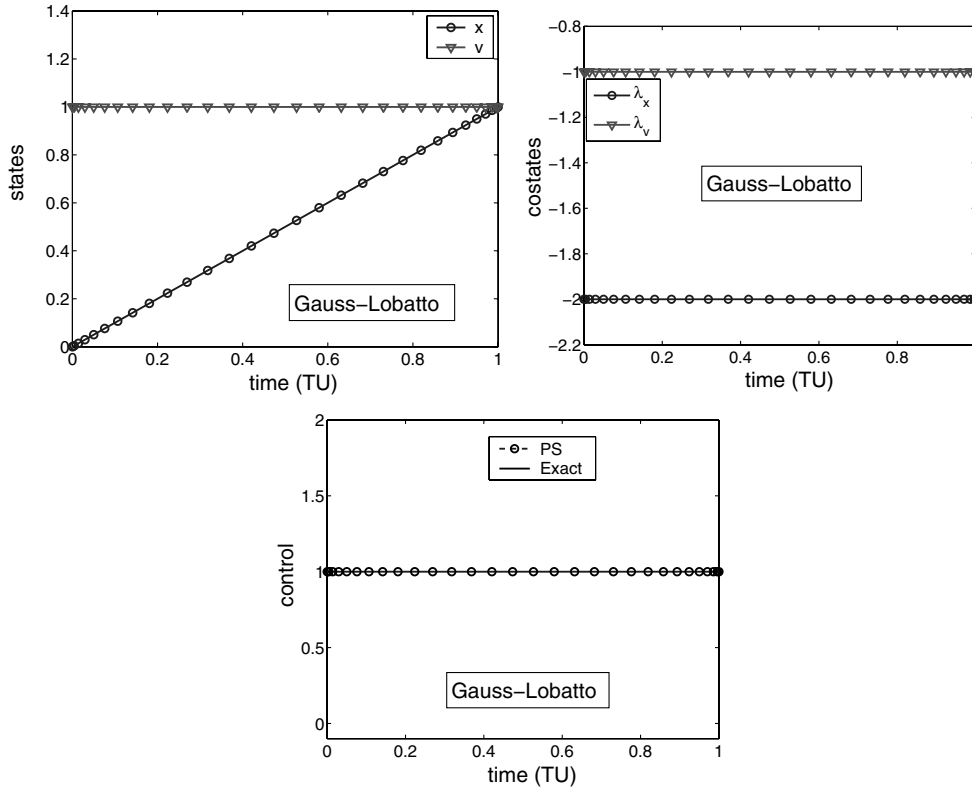


Fig. 7 Exact and PS solution for problem G using the correct PS method.

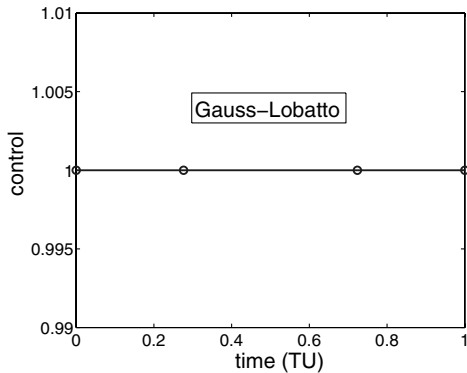


Fig. 8 Illustrating the exactness of the correct Legendre PS solution for problem G.

Gauss-Lobatto points not only converges pointwise (cf. Fig. 7), but its spectral convergence property implies that $N = 30$ is an overestimate of the number of nodes needed for convergence. This point is illustrated in Fig. 8, which shows that $N = 4$ Gauss-Lobatto points are sufficient to generate the exact solution. Figures 4 and 8 also underline an important point when comparing the performance of multiple numerical methods: the proper baseline to be used is not equality in the number of nodes, but equality in error tolerances. Under such a proper benchmark, the Legendre PS method based on

Gauss-Lobatto points has achieved the exact solution for $N = 4$, whereas those based on Gauss-Radau and Gauss points do not achieve this performance even when the number of nodes is increased by an order of magnitude.

A corollary to the convergence of PS solutions to exact optimal solutions is the convergence of the PS cost value J^N . Because L^1 is a larger space than L^∞ , a numerical examination of the PS cost by different PS methods on any number of problems may lead one to wrong conclusions. To illustrate this point, we compute J^N for problem G using all three PS methods. The result of this exercise is illustrated in Table 1. For $N = 4$ or 8, J^N , the cost using Gauss and Gauss-Radau quadrature rules for the corresponding PS methods, is larger than the cost using the Gauss-Lobatto approach. Based on this information, although one might arrive at the correct conclusion that the Gauss-Lobatto approach is superior to the other two approaches, it would be quite incorrect to state that the Gauss-Radau approach is better than the Gauss method, because both techniques have failed. As the number of nodes are increased, it is possible to draw even more erroneous conclusions. For instance, for $N = 25$, the Gauss and Gauss-Radau costs are lower than the exact minimum! In the absence of knowledge of the exact cost (namely, 1.0000), one may erroneously conclude that the Gauss-Radau points provide *better* solutions than the Gauss-Lobatto points and that the Gauss points provide the best solutions. Not only are such conclusions erroneous relative to the three PS methods, they are also absurd in implying that one can achieve solutions that yield cost values below the minimum!

VI. Conclusions

No number of examples can prove a claim, whereas a single counterexample is sufficient to disprove one. Furthermore, a counterexample for one method may not serve as a counterexample for another. This is why designing counterexamples is not an easy task. The counterexample presented in this Note is not some clever mathematical construct but an actual motion-planning problem designed by Gong et al. [15]. Counterexamples in optimal control theory have had a rich history in proving the incorrectness of certain hypotheses. In this Note, we have shown that Gong et al.'s motion-planning problem serves as an excellent counterexample to illustrate

Table 1 Illustrating the possibility of drawing wrong conclusions through cost value comparisons

N	Cost values			Exact
	Gauss	Gauss-Radau	Gauss-Lobatto	
4	1.2586	1.1220	1.0000	1.0000
8	1.0330	1.0140	1.0000	1.0000
15	0.9987	0.9987	1.0000	1.0000
25	0.9965	0.9979	1.0000	1.0000

the possibility of drawing any number of incorrect conclusions about computational optimal control methods.

We can now assert one fundamental fact about discrete approximations in optimal control theory: the controls may not converge pointwise (i.e., in the strong L^∞ norm) even if the states and costates converge to the right values. The impact of this result on pseudospectral methods is subtle but clear: *the proper pseudospectral method for solving a finite-horizon optimal control problem must be based on Gauss–Lobatto points*. Rigorous convergence theorems have been developed and proved along these lines. The counterexample also reveals that convergence cannot be proved for pseudospectral methods based on Gauss–Radau and Gauss points, except possibly under additional restrictive hypotheses. As in the case of Runge–Kutta methods, great care must be exercised in selecting the proper pseudospectral method for solving optimal control problems. In following this precept, we note that the preceding conclusions cannot be extended for infinite-horizon problems due to the same old reason that what works in an infinite domain may not work in a finite domain and vice versa. Consequently, although a Legendre pseudospectral method based on Gauss–Radau points may be incorrect for solving generic finite-horizon problems, it is quite an appropriate method for solving infinite-horizon optimal control problems. An important cautionary note that comes about as a corollary to the convergence result presented in this Note is the possibility of drawing absurd conclusions based on comparing cost values as a relative figure of merit.

Acknowledgments

We gratefully acknowledge partial funding for this research provided to one of the authors (Ross) by the U.S. Air Force Office of Scientific Research (AFOSR) under AFOSR grant F1ATA0-60-6-2G002. We also take this opportunity to express our appreciation to Qi Gong, with whom we had many stimulating discussions about the convergence of pseudospectral methods.

References

- [1] Halkin, H., "A Maximum Principle of the Pontryagin Type for Systems Described by Nonlinear Difference Equations," *SIAM Journal on Control*, Vol. 4, No. 1, 1966, pp. 90–111.
doi:10.1137/0304009
- [2] Cullum, J., "Discrete Approximations to Continuous Optimal Control Problems," *SIAM Journal on Control*, Vol. 7, No. 1, Feb. 1969, pp. 32–49.
doi:10.1137/0307003
- [3] Chang, S. S. L., "On Convexity and the Maximum Principle for Discrete Systems," *IEEE Transactions on Automatic Control*, Vol. AC-11, Feb. 1966, pp. 121–123.
- [4] Mordukhovich, B. S., and Shvartsman, I., "The Approximate Maximum Principle in Constrained Optimal Control," *SIAM Journal on Control and Optimization*, Vol. 43, No. 3, 2004, pp. 1037–1062.
doi:10.1137/S0363012903433012
- [5] Polak, E., *Optimization: Algorithms and Consistent Approximations*, Applied Mathematics Series, Vol. 124, Springer, Heidelberg, Germany, 1997, Chap. 4.
- [6] Mordukhovich, B. S., *Variational Analysis and Generalized Differentiation, 1: Basic Theory*, Grundlehren der Mathematischen Wissenschaften, Vol. 330, Springer, Berlin, 2005.
- [7] Mordukhovich, B. S., *Variational Analysis and Generalized Differentiation, 2: Applications*, Grundlehren der Mathematischen Wissenschaften, Vol. 331, Springer, Berlin, 2005.
- [8] Gong, Q., Ross, I. M., Kang, W., and Fahroo, F., "Connections Between the Covector Mapping Theorem and Convergence of Pseudospectral Methods for Optimal Control," *Computational Optimization and Applications* (to be published).
- [9] Hager, W. W., "Runge–Kutta Methods in Optimal Control and the Transformed Adjoint System," *Numerische Mathematik*, Vol. 87, No. 2, 2000, pp. 247–282.
doi:10.1007/s002110000178
- [10] Ross, I. M., "A Roadmap for Optimal Control: The Right Way to Commute," *Annals of the New York Academy of Sciences*, Vol. 1065, Dec. 2005, pp. 210–231.
doi:10.1196/annals.1370.015
- [11] Fahroo, F., and Ross, I. M., "A Second Look at Approximating Differential Inclusions," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 1, Jan.–Feb. 2001, pp. 131–133.
- [12] Fahroo, F., and Ross, I. M., "Costate Estimation by a Legendre Pseudospectral Method," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 2, Mar.–Apr. 2001, pp. 270–277.
- [13] Fahroo, F., and Ross, I. M., "Trajectory Optimization by Indirect Spectral Collocation Methods," AIAA Paper 2000-4028, 2000.
- [14] Ross, I. M., and Fahroo, F., "A Pseudospectral Transformation of the Covectors of Optimal Control Systems," *Proceedings of the 1st IFAC/IEEE Symposium on System Structure and Control*, Elsevier, Oxford, Dec. 2001.
- [15] Gong, 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.
doi:10.1109/TAC.2006.878570
- [16] Kang, W., Gong, Q., Ross, I. M., and Fahroo, F., "On the Convergence of Nonlinear Optimal Control Using Pseudospectral Methods for Feedback Linearizable Systems," *International Journal of Robust and Nonlinear Control*, Vol. 17, No. 14, Sept. 2007, pp. 1251–1277.
doi:10.1002/mc.1166
- [17] Kang, W., Ross, I. M., and Gong, Q., "Pseudospectral Optimal Control and Its Convergence Theorems," *Analysis and Design of Nonlinear Systems*, Springer, Berlin, Nov. 2007, pp. 109–124.
- [18] Ross, I. M., "A Historical Introduction to the Covector Mapping Principle," *Advances in the Astronautical Sciences*, Univelt, San Diego, CA, Vol. 123, 2006, pp. 1257–1278.
- [19] Fahroo, F., and Ross, I. M., "On Discrete-Time Optimality Conditions for Pseudospectral Methods," AIAA Paper 2006-6304, 2006.
- [20] Fahroo, F., and Ross, I. M., "Advances in Pseudospectral Methods for Optimal Control," AIAA Paper 2008-7309, 2008.
- [21] Biles, D. C., and Binding, P. A., "On Carathéodory's Conditions for the Initial Value Problem," *Proceedings of the American Mathematical Society*, Vol. 125, No. 5, May 1997, pp. 1371–1376.
doi:10.1090/S0002-9939-97-03942-7
- [22] Ross, I. M., Sekhvat, P., Fleming, A., and Gong, Q., "Optimal Feedback Control: Foundations, Examples and Experimental Results for a New Approach," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 2, Mar.–Apr. 2008, pp. 307–321.
- [23] Bedrossian, N., Bhatt, S., Lammers, M., Nguyen, L., and Zhang, Y., "First Ever Flight Demonstration of Zero Propellant Maneuver Attitude Control Concept," AIAA Paper 2007-6734, 2007.
- [24] Bedrossian, N., Bhatt, S., Lammers, M., and Nguyen, L., "Zero Propellant Maneuver: Flight Results for 180° ISS Rotation," 20th International Symposium on Space Flight Dynamics, Annapolis, MD, NASA CP-2007-214158, 24–28 Sept. 2007.
- [25] Gong, Q., Kang, W., Bedrossian, N., Fahroo, F., Sekhvat, P., and Bollino, K., "Pseudospectral Optimal Control for Military and Industrial Applications," *Proceedings of the 46th IEEE Conference on Decision and Control*, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, Dec. 2007, pp. 4128–4142.
- [26] Gong, Q., Fahroo, F., and Ross, I. M., "A Spectral Algorithm for Pseudospectral Methods in Optimal Control," *Journal of Guidance, Control, and Dynamics*, Vol. 31, No. 3, May–June 2008, pp. 460–471.
doi:10.2514/1.32908
- [27] Ross, I. M., and Gong, Q., "Guess-Free Trajectory Optimization," AIAA Paper 2008-6273, 2008.