Outline of Ten Main Lectures by Dr. Hager William

Computational Methods in Optimal Control

1. Background

An optimal control problem is a variational problem that involves partial or ordinary differential equations describing the dynamics of a system and an objective such as minimization of energy or cost or time or drag. There are two types of variables in an optimal control problem, the state variables are the solution of the dynamic equations, and the control variables are parameters within the dynamics. The goal is to optimize the objective while satisfying the constraints. Besides the constraints given by the system dynamics, additional constraints may arise, such as bounds on components of the control or state variables.

Optimal control theory has played an increasingly important role in numerous science and engineering applications, ranging from the ground-breaking 1969 Apollo moon landing, through the ubiquitous Kalman filter, to the effective treatment of cancer and the optimal hot rolling of steel. The initial applications of optimal control were in the aerospace industry, while in recent decades, optimal control has been applied to a wide range of models in fields including medicine, biology, chemical manufacturing, economics, mechanics, and materials science.

Due to the growing complexity of optimal control applications, solutions are often obtained by numerical algorithms. The lectures of this NSF-CBMS conference on computational methods in optimal control will focus on discretization and solution techniques. As part of the program, attendees will be provided access to the recently developed GPOPS-II software for numerically solving optimal control problems with MATLAB using *hp*-adaptive orthogonal collocation techniques. Attendees will be encouraged to bring their own problems, or will be given problems to solve during the afternoon laboratory and discussion sessions.

There is a fundamental difference between solving a differential equation and solving a control problem with the same differential equation constraint. We illustrate this with the following simple example:

$$\begin{aligned} &\text{minimize} \quad \frac{1}{2}\int_0^1 u(t)^2 + 2x(t)^2 \, dt \\ &\text{subject to} \quad \dot{x}(t) = .5x(t) + u(t), \quad x(0) = 1, \end{aligned} \tag{1}$$

with the optimal solution

$$x^*(t) = \frac{2e^{3t} + e^3}{e^{3t/2}(2 + e^3)}, \quad u^*(t) = \frac{2(e^{3t} - e^3)}{e^{3t/2}(2 + e^3)}.$$

In an undergraduate course on ordinary differential equations, we might introduce the students to the Euler's method, and then point out that the improved Euler's method yields a much more

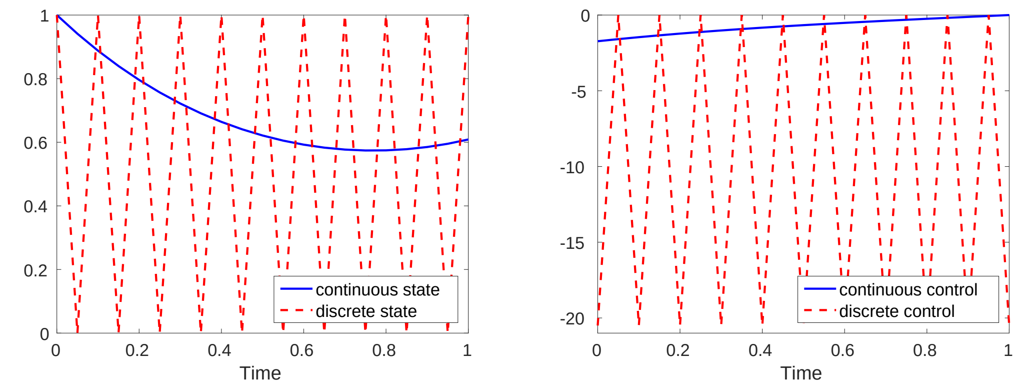

Figure 1: Discrete and Continuous Solutions

accurate solution. The improved Euler discretization of (1) is as follows:

$$
\begin{aligned}
&\text{minimize} \quad \frac{h}{2}\sum_{k=0}^{N-1} u_{k+1/2}^2 + 2x_{k+1/2}^2 \\
&\text{subject to} \quad x_{k+1/2} = x_k + \tfrac{h}{2}(.5x_k + u_k), \\
&\qquad\qquad\quad\; x_{k+1} = x_k + h(.5x_{k+1/2} + u_{k+1/2}), \quad x_0 = 1.
\end{aligned}
$$

Here $h = 1/N$ is the stepsize and (x_k, u_k) is the discrete approximation at $t = kh$. The objective function is nonnegative, and the discrete problem has an explicit solution which makes the objective function zero:

$$u_k = -(2N + 0.5), \quad x_k = 1, \quad u_{k+1/2} = x_{k+1/2} = 0.$$

The discrete solution, shown in Figure 1, clearly does not converge to the solution of the continuous problem as N tends to infinity. On the other hand, changing the first equation in the discretization to

$$x_{k+1/2} = x_k + \frac{h}{2}(.5x_k + u_{k+1/2})$$

results in a discretization which yields a second-order approximation to the continuous solution [5].

The stark contrast between discretizations of initial-value problems and discretizations of control problems where one not only solves an equation, but also solves an optimization problem, has led to a fundamentally new theory for the analysis of discrete approximations and for the convergence of numerical algorithms. The analysis centers around perturbed inclusions. The necessary optimality conditions for the numerical approximation to the continuous control problem can be expressed in the following abstract form:

$$\text{Find } \mathbf{w} \in \mathcal{X} \text{ such that } \mathcal{T}^h(\mathbf{w}) \in \mathcal{F}^h(\mathbf{w}). \tag{2}$$

Here $\mathcal{X}$ is a Banach space, h might denote a discretization parameter, $\mathcal{T}^h : \mathcal{X} \to \mathcal{Y}$, $\mathcal{Y}$ is a linear normed space, and $\mathcal{F}^h : \mathcal{X} \to 2^{\mathcal{Y}}$. A solution $\mathbf{w}^h$ of (2), represents an approximation to a continuous solution of the original problem. Since the continuous solution is usually not an element of $\mathcal{X}$, we often project the continuous solution into $\mathcal{X}$ (for example, evaluate the continuous solution at the mesh points). We let $\mathbf{w}^*$ denote this projection of the continuous solution.

To estimate the distance between $\mathbf{w}^h$ and $\mathbf{w}^*$, we first measure the distance from $\mathcal{T}^h(\mathbf{w}^*)$ to $\mathcal{F}^h(\mathbf{w}^*)$, or equivalently, a parameter $\boldsymbol{\delta}$ is introduced for which

$$\mathcal{T}^h(\mathbf{w}^*) + \boldsymbol{\delta} \in \mathcal{F}^h(\mathbf{w}^*).$$

The goal is to relate the distance $\|\mathbf{w}^h - \mathbf{w}^*\|$ to the norm $\|\boldsymbol{\delta}\|$ of the residual $\boldsymbol{\delta}$. This is done by the study of a linearization of (2):

$$\text{Find } \mathbf{w} \in \mathcal{X} \text{ such that } \mathcal{L}(\mathbf{w}) + \boldsymbol{\pi} \in \mathcal{F}^h(\mathbf{w}). \tag{3}$$

Here $\mathcal{L}$ is a linear operator approximating $\mathcal{T}^h$ (for example, the derivative at $\mathbf{w}^*$ if it exists), and $\boldsymbol{\pi} \in \mathcal{Y}$ represents a parameter. An analysis of the Lipschitz properties of solutions of (3) with respect to the parameter $\boldsymbol{\pi}$ eventually leads to an estimate for the distance between $\mathbf{w}^h$ and $\mathbf{w}^*$ relative to $\|\boldsymbol{\delta}\|$ which yields convergence rates for algorithms as well as bounds for how the error in a discrete approximation depends on a discretization parameter. See [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12] and the references therein.

Different classes of algorithms or different types of discretizations lead to different types of analysis. The proposed conference and monograph will cover this theory and its application to a variety of numerical methods for solving optimal control problems. This will provide a foundation both for continuing research on numerical methods for optimal control, and for variational problems in general, and will yield practical guidance for how to achieve convergent algorithms for solving optimal control problems. The conference and monograph are particularly opportune due to the recent development of orthogonal collocation techniques, which have the potential for solving control problems with much greater speed and accuracy than was previously possible, and which have led to many open problems, some that are easy to state, and with potentially broad interests in the area of computational data science and engineering.

3. Brief Description of the Lectures

This 5-days workshop will include ten main lectures during the morning sessions, which are accompanied with carefully designed computer lab sessions in the afternoon focusing on computer implementations of the algorithms. A brief description of the lectures follows.

Lecture 1. General classes of discretizations and algorithms for solving optimal control problems will be introduced including Runge-Kutta schemes, orthogonal collocation methods, penalty and multiplier methods, and sequential quadratic programming. The behavior of discretization schemes for initial-value problems will be contrasted with their behavior for control problems.

Lecture 2. We start to show how each of the schemes introduced in lecture 1 can be reformulated as an inclusion of the form (2). This requires a discussion of the first-order optimality conditions for both the control problem, and a mathematical programming problem. These conditions in the context of optimal control are often referred to as the Pontryagin minimum principle, while they are the Karush-Kuhn-Tucker (KKT) conditions in the context of mathematical programming.

Lecture 3. We continue to show how to connect the schemes of lecture 1 to an inclusion (2). The KKT conditions often do not resemble the Pontryagin minimum principle. But a suitable transformation of the KKT multipliers leads to a transformed adjoint system [8] that resembles the costate equation in the Pontryagin minimum principle. Different schemes require different transformations.

Lecture 4. The analysis of finite difference approximations to initial value problems is often based on the Lax equivalence theorem which roughly states that a consistent approximation to a well-posed linear initial value problem is convergent if and only if it is stable. A corresponding result will be developed in the context of the inclusion (2) with stability in the Lax equivalence theorem replaced by a Lipschitz property for the linearization (3). In a sense, the convergence result for (2) generalizes the Lax equivalence theorem since $\mathcal{T}$ may be nonlinear and an equation has been replaced by an inclusion, which allows the treatment of variational inequalities.

Lecture 5. A key step in the convergence analysis for a algorithm is the analysis of Lipschitz stability for a solution of the linearized problem (3) as function of the parameter $\boldsymbol{\pi}$. We show how to reformulate the analysis of (3) into the analysis of a related quadratic programming problem (QP). Different algorithms for the control problem lead to different variations of the QP; but after the reduction to a QP, there are a variety of tools for studying dependence of a solution on a parameter.

Lecture 6. Using the theoretical foundation developed in the previous lectures, the general classes of algorithms and discretizations introduced in Lecture 1 will be analyzed. We start with penalty/multiplier methods and sequential quadratic programming where the analysis can be performed in an infinite dimensional setting, and then proceed to Runge-Kutta methods, where the analysis is finite dimensional. An important step in the analysis of Runge-Kutta methods is to bound the residual. The size of the residual depends both on the order of the original Runge-Kutta scheme, and the order of a new Runge-Kutta scheme associated with the transformed adjoint system. Consequently, the class of Runge-Kutta methods that can be used to solve optimal control problems is a subset of the class that can be used to solve an initial value problem.

Lecture 7. This lecture begins the analysis of orthogonal collocation methods in which the state is approximated by a polynomial and the dynamics are collocated at the roots of an orthogonal polynomial. It is shown how the KKT multipliers for the discrete problem can be used to construct a polynomial that approximates the continuous adjoint variable; as a result the KKT conditions can be reformulated in a polynomial spaces. Analysis is developed for schemes based on collocation at either Radau or Gauss quadrature points. For collocation at the Lobatto quadrature points, new issues arise and the theory remains relatively open.

Lecture 8. The analysis of the residual in orthogonal collocation methods is reduced to the analysis of the derivative error associated with a Lagrange interpolant that passes through the state at the collocation points and at the initial time. For the Gauss collocation points, this analysis can be performed in a Sobolev space setting, leading to a tight bound for the derivative error. For the Radau collocation points, additional technicalities arise and the analysis is done in an infinity norm, which involves the consideration of Lebesgue constants. Nonetheless, for both the Gauss and Radau collocation points, an exponential convergence rate holds for the discrete approximation when the solution to the optimal control problem is smooth. The analysis of Lipschitz stability for the linearized problem (3) leads to a number of intriguing properties for orthogonal collocation schemes, which are key to their convergence, and which have been verified numerically, but not proved (see http://users.clas.ufl.edu/hager/papers/prize.pdf for an example).

Lecture 9. The exponential convergence rate for orthogonal collocation schemes requires a smooth solution. However, when there are constraints in the problem formulation, such as bounds on the state or control components, a solution may be nonsmooth. *hp*-orthogonal collocation methods can be used to achieve fast convergence in the nonsmooth case. A mesh is introduced, similar to what might be done with a Runge-Kutta scheme, however, on each mesh interval a different orthogonal polynomial is used to approximate the state variable. The *hp*-schemes attempt to use a polynomial for convergence in regions where the solution is smooth, while mesh intervals are moved or added where smoothness is lost.

Lecture 10. Techniques for solving the discretized control problem are examined. General purpose optimization packages can be used, or in certain cases, more specialized techniques can be used that exploit the structure of the optimal control problem. In particular, it is shown that the transformed adjoint system yields an efficient mechanism for computing the derivative of the objective with respect to the control variable, which can be exploited in gradient descent methods. Shooting methods are another approach which might work well when a good initial guess for the problem solution is known.

References

[1] A. Dontchev, W. W. Hager, A. Poore, and B. Yang, *Optimality, stability, and convergence in nonlinear control*, Applied Math. and Optim., 31 (1995), pp. 297–326.

[2] A. L. Dontchev and W. W. Hager, *Lipschitzian stability in nonlinear control and optimization*, SIAM J. Control Optim., 31 (1993), pp. 569–603.

[3] ——, *Lipschitzian stability for state constrained nonlinear optimal control*, SIAM J. Control Optim., 36 (1998), pp. 696–718.

[4] ——, *The Euler approximation in state constrained optimal control*, Math. Comp., 70 (2001), pp. 173–203.

[5] A. L. Dontchev, W. W. Hager, and V. M. Veliov, *Second-order Runge-Kutta approximations in constrained optimal control*, SIAM J. Numer. Anal., 38 (2000), pp. 202–226.

[6] ——, *Uniform convergence and mesh independence of Newton's method for discretized variational problems*, SIAM J. Control Optim., 39 (2000), pp. 961–980.

[7] W. W. Hager, *Multiplier methods for nonlinear optimal control*, SIAM J. Numer. Anal., 27 (1990), pp. 1061–1080.

[8] ——, *Runge-Kutta methods in optimal control and the transformed adjoint system*, Numer. Math., 87 (2000), pp. 247–282.

[9] ——, *Numerical analysis in optimal control*, in International Series of Numerical Mathematics, K.-H. Hoffmann, I. Lasiecka, G. Leugering, J. Sprekels, and F. Tröltzsch, eds., vol. 139, Basel/Switzerland, 2001, Birkhauser Verlag, pp. 83–93.

[10] W. W. Hager, H. Hou, S. Mohapatra, and A. V. Rao, *Convergence rate for an hp-collocation method applied to unconstrained optimal control*, (2016, arXiv.org/abs/1605.02121).

[11] W. W. Hager, H. Hou, and A. V. Rao, *Convergence rate for a Gauss collocation method applied to unconstrained optimal control*, J. Optim. Theory Appl., 169 (2016), pp. 801–824.

[12] W. W. Hager, J. Liu, S. Mohapatra, A. V. Rao, and X.-S. Wang, *Convergence rate for a gauss collocation method applied to constrained optimal control*, SIAM J. Control Optim., 56 (2018), pp. 1386–1411.