Automatica 47 (2011) 829–837

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Pseudospectral methods for solving infinite-horizon optimal control problems[☆]

Divya Garg, William W. Hager, Anil V. Rao [*]

University of Florida, Gainesville, FL 32611, USA

ARTICLE INFO

Article history:
Available online 1 March 2011

Keywords:
Optimal control
Pseudospectral methods
Nonlinear programming

ABSTRACT

An important aspect of numerically approximating the solution of an infinite-horizon optimal control problem is the manner in which the horizon is treated. Generally, an infinite-horizon optimal control problem is approximated with a finite-horizon problem. In such cases, regardless of the finite duration of the approximation, the final time lies an infinite duration from the actual horizon at $t = +\infty$. In this paper we describe two new direct pseudospectral methods using Legendre–Gauss (LG) and Legendre–Gauss–Radau (LGR) collocation for solving infinite-horizon optimal control problems numerically. A smooth, strictly monotonic transformation is used to map the infinite time domain $t \in [0, \infty)$ onto a half-open interval $\tau \in [-1, 1)$. The resulting problem on the finite interval is transcribed to a nonlinear programming problem using collocation. The proposed methods yield approximations to the state and the costate on the entire horizon, including approximations at $t = +\infty$. These pseudospectral methods can be written equivalently in either a differential or an implicit integral form. In numerical experiments, the discrete solution exhibits exponential convergence as a function of the number of collocation points. It is shown that the map $\phi : [-1, +1) \to [0, +\infty)$ can be tuned to improve the quality of the discrete approximation.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Over the last decade, pseudospectral methods have become increasingly popular in the numerical solution of optimal control problems. Benson (2004), Benson, Huntington, Thorvaldsen, and Rao (2006), Elnagar, Kazemi, and Razzaghi (1995), Elnagar and Razzaghi (1997), Fahroo and Ross (2008), Garg et al. (2009), Garg et al. (2010), Huntington (2007), Rao et al. (2010) and Williams (2004) Pseudospectral methods are a class of *direct collocation* methods where the optimal control problem is transcribed to a nonlinear programming problem (NLP) by parameterizing the state and the control using global polynomials and collocating the differential–algebraic equations using nodes obtained from a Gaussian quadrature. The three most commonly used sets of collocation points are *Legendre–Gauss* (LG), *Legendre–Gauss–Radau* (LGR), and *Legendre–Gauss–Lobatto* (LGL) points. These three sets of points are obtained from the roots of a Legendre polynomial and/or linear combinations of a Legendre polynomial and its derivatives. All three sets of points are defined on the domain $[-1, 1]$, but differ significantly in that the LG points include *neither* of the endpoints, the LGR points include *one* of the endpoints, and the LGL points include *both* of the endpoints. In recent years, the two most well-documented pseudospectral methods are the *Legendre–Gauss–Lobatto pseudospectral method* (Elnagar et al., 1995) and the *Legendre–Gauss pseudospectral method* (Benson et al., 2006; Garg et al., 2010; Rao et al., 2010).

In this paper we describe two new pseudospectral methods for the numerical solution of nonlinear infinite-horizon optimal control problems based on either LG or LGR collocation. For either scheme, a smooth, strictly monotonic change of variables is used to map the domain of the infinite time interval $t \in [0, \infty)$ to a finite half open time interval $\tau \in [-1, +1)$. The resulting finite horizon problem is discretized using either LG or LGR collocation. Our collocation schemes avoid the singularity at $\tau = +1$ introduced by the change of variables. Furthermore, in the LG scheme, an explicit formula is derived to include the state at the horizon (that is, at $t = +\infty$) in the NLP, while in the LGR scheme the state at $t = +\infty$ is included in the NLP as a variable in state approximation. Thus, either scheme developed in this paper yields an estimate for the state on the *entire* horizon. In addition, we also present the transformed adjoint systems that relate the Lagrange multipliers of the NLP to the costate of the continuous control problem.

We note that an LGR pseudospectral method for approximating the solution of nonlinear infinite-horizon optimal control problems has been previously developed in Fahroo and Ross (2008).

☆ The authors gratefully acknowledge partial support for this research from the US Army Research Office under Grant 55173-CI, the US Air Force Research Laboratory under Contract FA8651-08-D-0108, the National Science Foundation under Grant 0620286, and the US Office of Naval Research under Grant N00014-11-1-0068. The material in this paper was partially presented at the 2010 AIAA Guidance, Navigation, and Control Conference, August 2–5, 2010, Toronto, Ontario, Canada. This paper was recommended for publication in revised form under the direction of Editor Berç Rüstem.

* Corresponding author. Tel.: +1 352 392 5523; fax: +1 352 392 7303.
E-mail addresses: divyagarg2002@ufl.edu (D. Garg), hager@ufl.edu (W.W. Hager), anilvrao@ufl.edu (A.V. Rao).

0005-1098/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2011.01.085

In Fahroo and Ross (2008), the infinite-horizon problem is converted to a finite-horizon problem using the change of variables $t = (1+\tau)/(1-\tau)$. The finite interval $\tau \in [-1, +1]$ corresponds to the infinite time interval $t \in [0, \infty)$ in the original infinite-horizon problem. The finite-horizon problem is then discretized by collocation at the LGR points. Because the final LGR point is strictly smaller than $\tau = +1$, the singularity in the change of variables $t = (1+\tau)/(1-\tau)$ is avoided. Fahroo and Ross (2008) then provides two numerical examples that illustrate the approach.

The contributions of our paper are as follows. First, we derive two new pseudospectral methods using LG and LGR points for solving the nonlinear infinite-horizon optimal control problem. Our LGR scheme is different from the LGR scheme in Fahroo and Ross (2008). In either of our methods, the state and the costate at the horizon (that is, $t = +\infty$), are variables in the discrete approximation. As a result, the state and the costate solutions are obtained on the entire horizon. Second, our methods have the property that the differentiation matrices are rectangular and full rank, leading to the property that either of our discrete approximations can be written equivalently as an integral method. Third, we consider a general change of variables $t = \phi(\tau)$ of an infinite-horizon problem to a finite-horizon problem. We find that better approximations to the continuous-time problem can be attained by using a function $\phi(\tau)$ that grows more slowly than $t = (1+\tau)/(1-\tau)$ near the point $\tau = +1$. Hence, by tuning the choice of the transformation, the accuracy of the discretization can be improved.

This paper is organized as follows. In Section 2 we provide a description of our notation. In Section 3 we state the infinite-horizon optimal control problem. Section 4 describes our Gauss and Radau pseudospectral methods for solving infinite-horizon optimal control problems. In addition, we show that the first-order optimality conditions associated with our Gauss and Radau methods are equivalent to pseudospectral schemes for the continuous costate equation. In Section 5 we demonstrate the methods on two examples. Finally, in Section 6 we provide conclusions.

2. Notation

Throughout the paper, we employ the following notation. First, we treat all vector functions of time as *row* vectors; that is, $\mathbf{x}(\tau) = [x_1(\tau), \ldots, x_n(\tau)] \in \mathbb{R}^n$, where n is the continuous-time dimension of $\mathbf{x}(\tau)$. $\mathbf{B}^{\mathsf{T}}$ denotes the transpose of a matrix $\mathbf{B}$. Given $\mathbf{a}$ and $\mathbf{b} \in \mathbb{R}^n$, $\langle \mathbf{a}, \mathbf{b} \rangle$ is their dot product. If $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^m$, then $\nabla \mathbf{f}$ is the m by n Jacobian matrix whose ith row is ∇f_i. In particular, the gradient of a scalar-valued function is a row vector. If $\phi : \mathbb{R}^{m\times n} \to \mathbb{R}$ and $\mathbf{X}$ is an m by n matrix, then $\nabla \phi$ denotes the m by n matrix whose (i, j) element is $(\nabla \phi(\mathbf{X}))_{ij} = \partial \phi(\mathbf{X})/\partial X_{ij}$. If $\mathbf{A}$ is a matrix, then $\mathbf{A}_{i:j}$ is the submatrix formed by rows i through j, while $\mathbf{A}_i$ is the ith row of $\mathbf{A}$. The Kronecker delta function is defined by $\delta_{ii} = 1$ and $\delta_{ij} = 0$ if $i \neq j$.

3. Infinite-horizon optimal control problem

Consider the infinite-horizon optimal control problem:

$$\begin{aligned} &\min J = \int_0^\infty g(\mathbf{x}(t), \mathbf{u}(t))\mathrm{d}t \\ &\text{subject to} \\ &\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)), \quad \mathbf{x}(0) = \mathbf{x}_0, \end{aligned} \tag{1}$$

where $g : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$, and $\dot{\mathbf{x}}$ denotes the time derivative of $\mathbf{x}$. We make the change of variables $t = \phi(\tau)$ where ϕ is a differentiable, strictly monotonic function of τ that maps the interval $[-1, 1]$ onto $[0, \infty)$. Three examples of such a function are

$$\phi_a(\tau) = \frac{1+\tau}{1-\tau}, \tag{2}$$

$$\phi_b(\tau) = \log\left(\frac{2}{1-\tau}\right), \tag{3}$$

$$\phi_c(\tau) = \log\left(\frac{4}{(1-\tau)^2}\right). \tag{4}$$

The change of variables $\phi_a(\tau)$ was originally proposed in Fahroo and Ross (2008), while the transformations $\phi_b(\tau)$ and $\phi_c(\tau)$ are introduced in this paper. These latter changes of variables produce slower growth in $t = \phi(\tau)$ as τ approaches $+1$, than that of $\phi_a(\tau)$. As we will see in the numerical experiments, better discretization can be achieved by tuning the change of variables to the problem.

Define $T(\tau) = \mathrm{d}\phi/\mathrm{d}\tau \equiv \phi'(\tau)$. After changing variables from t to τ, the infinite-horizon optimal control problem becomes

$$\begin{aligned} &\min J = \int_{-1}^{+1} T(\tau) g(\mathbf{x}(\tau), \mathbf{u}(\tau))\mathrm{d}\tau \\ &\text{subject to} \\ &\dot{\mathbf{x}}(\tau) = T(\tau)\mathbf{f}(\mathbf{x}(\tau), \mathbf{u}(\tau)), \quad \mathbf{x}(-1) = \mathbf{x}_0. \end{aligned} \tag{5}$$

Here $\mathbf{x}(\tau)$ and $\mathbf{u}(\tau)$ denote the state and the control as a function of the new variable τ. Formally, the first-order optimality conditions for the finite horizon control problem Eq. (5), also called the Pontryagin minimum principle, are

$$\dot{\boldsymbol{\lambda}}(\tau) = -T(\tau)\nabla_x H(\mathbf{x}(\tau), \mathbf{u}(\tau), \boldsymbol{\lambda}(\tau)), \quad \boldsymbol{\lambda}(1) = \mathbf{0}, \tag{6}$$

$$\mathbf{0} = \nabla_u H(\mathbf{x}(\tau), \mathbf{u}(\tau), \boldsymbol{\lambda}(\tau)), \tag{7}$$

where $H(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = g(\mathbf{x}, \mathbf{u}) + \langle \boldsymbol{\lambda}, \mathbf{f}(\mathbf{x}, \mathbf{u}) \rangle$ is the Hamiltonian for Eq. (1).

4. Pseudospectral methods for infinite-horizon optimal control problems

In this section we formulate discrete approximations to the nonlinear infinite-horizon optimal control problem described in Section 3. These discrete schemes are based on global collocation using either Gauss or Radau collocation points. As will be seen, these two schemes differ in their treatment of the horizon. For the Gauss quadrature scheme, the state at the horizon is included by quadrature, while for the Radau scheme, the state at the horizon is included in the state approximation.

4.1. Infinite-horizon Gauss pseudospectral method

Consider the LG collocation points $(\tau_1, \ldots, \tau_N)$ on the interval $(-1, 1)$ and two additional *noncollocated* points $\tau_0 = -1$ (the initial time) and $\tau_{N+1} = 1$ (the terminal time, corresponding to $t = +\infty$). The state is approximated by a polynomial of degree at most N as

$$\begin{aligned} &\mathbf{x}(\tau) \approx \sum_{j=0}^{N} \mathbf{X}_j L_j(\tau), \\ &L_j(\tau) = \prod_{\substack{k=0 \\ k\neq j}}^{N} \frac{\tau - \tau_k}{\tau_j - \tau_k}, \quad j = 0, \ldots, N, \end{aligned} \tag{8}$$

where $\mathbf{X}_j \in \mathbb{R}^n$ and L_j is a basis of Nth-degree Lagrange polynomials. Notice that the basis includes the function L_0 corresponding to the initial time $\tau_0 = -1$, but not a function corresponding to $\tau_{N+1} = +1$. Differentiating the series of Eq. (8) and evaluating at the collocation point τ_i gives

$$\dot{\mathbf{x}}(\tau_i) \approx \sum_{j=0}^{N} \mathbf{X}_j \dot{L}_j(\tau_i) = \mathbf{D}_i \mathbf{X}, \tag{9}$$

where $\mathbf{D}_i$ is the ith row of $\mathbf{D}$,

$$D_{ij} = \dot{L}_j(\tau_i), \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} \mathbf{X}_0 \\ \vdots \\ \mathbf{X}_N \end{bmatrix}.$$

The rectangular $N \times (N+1)$ matrix $\mathbf{D}$ formed by the coefficients D_{ij}, $(i = 1, \ldots, N;\ j = 0, \ldots, N)$ is the *Gauss Pseudospectral differentiation matrix* since it transforms the state approximation at $\tau_0, \ldots, \tau_N$ to the derivative of the state approximation at the collocation points $\tau_1, \ldots, \tau_N$.

Let $\mathbf{U}$ be an $N \times m$ matrix whose ith row $\mathbf{U}_i$ is an approximation to the control $\mathbf{u}(\tau_i)$, $1 \le i \le N$. Our discrete approximation to the system dynamics $\dot{\mathbf{x}}(\tau) = T(\tau)\mathbf{f}(\mathbf{x}(\tau), \mathbf{u}(\tau))$ is obtained by evaluating the system dynamics at each collocation point and replacing $\dot{\mathbf{x}}(\tau_i)$ by its discrete approximation $\mathbf{D}_i\mathbf{X}$. Hence, the discrete approximation to the system dynamics is

$$\mathbf{D}_i\mathbf{X} = T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i), \quad 1 \le i \le N. \tag{10}$$

It is important to observe that the left-hand side of Eq. (10) contains approximations for the state at the initial point plus the LG points while the right-hand side contains approximations for the state (and control) at only the LG points. The objective function in Eq. (5) is approximated by a Legendre–Gauss quadrature as

$$J = \int_{-1}^{+1} T(\tau) g(\mathbf{x}(\tau), \mathbf{u}(\tau)) d\tau \approx \sum_{i=1}^{N} w_i T(\tau_i) g(\mathbf{X}_i, \mathbf{U}_i),$$

where w_i is the quadrature weight associated with τ_i. The state at the horizon is estimated by quadrature as

$$\mathbf{x}(+1) = \mathbf{x}(-1) + \int_{-1}^{+1} T(\tau)\mathbf{f}(\mathbf{x}(\tau), \mathbf{u}(\tau)) d\tau$$
$$\approx \mathbf{X}_0 + \sum_{i=1}^{N} w_i T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \equiv \mathbf{X}_{N+1}, \tag{11}$$

where $\mathbf{X}_{N+1}$ is treated as an additional variable. Rearranging Eq. (11), the following equality constraint is then added in the discrete approximation:

$$\mathbf{X}_0 - \mathbf{X}_{N+1} + \sum_{i=1}^{N} w_i T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) = \mathbf{0}. \tag{12}$$

It is noted that adding the constraint of Eq. (12) and the variable $\mathbf{X}_{N+1}$ does not change the number of degrees of freedom because Eq. (12) is the same size as $\mathbf{X}_{N+1}$. The continuous-time nonlinear infinite-horizon optimal control problem of Eq. (1) is then approximated by the following NLP:

$$\min \sum_{i=1}^{N} w_i T(\tau_i) g(\mathbf{X}_i, \mathbf{U}_i)$$

subject to

$$T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) - \mathbf{D}_i\mathbf{X} = \mathbf{0}, \quad 1 \le i \le N, \tag{13}$$
$$\mathbf{X}_0 - \mathbf{X}_{N+1} + \sum_{i=1}^{N} w_i T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) = \mathbf{0},$$
$$\mathbf{x}_0 - \mathbf{X}_0 = \mathbf{0}.$$

Although the change of variables $t = \phi(\tau)$ must have a singularity at $\tau = +1$, we never evaluate $T(\tau) = \phi'(\tau)$ at the singularity in Eq. (13), rather we evaluate T at the quadrature points which are all strictly less than 1.

The first-order optimality conditions for Eq. (13), also called the KKT conditions, are obtained by differentiating the Lagrangian $\mathcal{L}$ with respect to the components of $\mathbf{X}$ and $\mathbf{U}$. The Lagrangian associated with Eq. (13) is

$$\mathcal{L}(\boldsymbol{\Lambda}, \boldsymbol{\Lambda}_0, \boldsymbol{\Lambda}_{N+1}, \mathbf{X}, \mathbf{U})$$
$$= \sum_{i=1}^{N} \left(w_i T(\tau_i) g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) - \mathbf{D}_i\mathbf{X} \rangle \right)$$
$$+ \langle \boldsymbol{\Lambda}_0, \mathbf{x}_0 - \mathbf{X}_0 \rangle + \left\langle \boldsymbol{\Lambda}_{N+1}, \mathbf{X}_0 - \mathbf{X}_{N+1} + \sum_{i=1}^{N} w_i T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \right\rangle, \tag{14}$$

where $\boldsymbol{\Lambda}$ is an $N \times n$ matrix of Lagrange multipliers associated with the collocation points, and $\boldsymbol{\Lambda}_0$ and $\boldsymbol{\Lambda}_{N+1}$ are each $1 \times n$ row vectors of Lagrange multipliers associated with the initial condition and the quadrature equation, respectively. Differentiating the Lagrangian with respect to $\mathbf{X}_0$, $\mathbf{X}_{N+1}$, $\mathbf{X}_i$ and $\mathbf{U}_i$, $1 \le i \le N$, gives us the optimality conditions

$$\boldsymbol{\Lambda}_0 = -\mathbf{D}_0^{\mathsf{T}}\boldsymbol{\Lambda}, \tag{15}$$
$$\boldsymbol{\Lambda}_{N+1} = \mathbf{0}, \tag{16}$$
$$\mathbf{D}_i^{\mathsf{T}}\boldsymbol{\Lambda} = T(\tau_i)\nabla_x \left(w_i g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \rangle \right), \tag{17}$$
$$\mathbf{0} = T(\tau_i)\nabla_u \left(w_i g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \rangle \right), \tag{18}$$

where $\mathbf{D}_i^{\mathsf{T}}$ is the ith row of $\mathbf{D}^{\mathsf{T}}$.

Using an approach similar to that of Hager (2000), the first-order optimality conditions given in Eqs. (15)–(18) can be reformulated so that they become a discretization of the first-order optimality conditions for the continuous control problem given in Eq. (5). First, define the following expressions:

$$\boldsymbol{\lambda}_0 = \boldsymbol{\Lambda}_0, \qquad \boldsymbol{\lambda}_{N+1} = \boldsymbol{\Lambda}_{N+1},$$
$$\boldsymbol{\lambda} = \mathbf{W}^{-1}\boldsymbol{\Lambda}, \qquad \mathbf{D}^{\dagger} = -\mathbf{W}^{-1}\mathbf{D}_{1:N}^{\mathsf{T}}\mathbf{W}, \tag{19}$$

were $\mathbf{W}$ is the diagonal matrix whose ith diagonal element is w_i. Making these substitutions in Eqs. (15)–(18), we can rewrite the optimality conditions as

$$\boldsymbol{\lambda}_0 = -\mathbf{D}_0^{\mathsf{T}}\mathbf{W}\boldsymbol{\lambda}, \tag{20}$$
$$\boldsymbol{\lambda}_{N+1} = \mathbf{0}, \tag{21}$$
$$\mathbf{D}_i^{\dagger}\boldsymbol{\lambda} = -T(\tau_i)\nabla_x H(\mathbf{X}_i, \mathbf{U}_i, \boldsymbol{\lambda}_i), \quad 1 \le i \le N, \tag{22}$$
$$\mathbf{0} = \nabla_u H(\mathbf{X}_i, \mathbf{U}_i, \boldsymbol{\lambda}_i). \tag{23}$$

Hence, this transformation makes the discrete optimality conditions look very similar to the continuous Pontryagin minimum principle of Eq. (6)–(7). There are two basic differences: The continuous derivative $\dot{\boldsymbol{\lambda}}(\tau)$ is replaced by the discrete analog $\mathbf{D}_i^{\dagger}\boldsymbol{\lambda}$ and the definition of initial costate of Eq. (20) does not appear in the continuous optimality conditions.

In Garg et al. (2010), we show that $\mathbf{D}^{\dagger}$ is a differentiation matrix for polynomials of degree N; more precisely, if p is a polynomial of degree at most N with values $p_i = p(\tau_i)$, $1 \le i \le N$, then

$$(\mathbf{D}^{\dagger}\mathbf{p})_i = \dot{p}(\tau_i), \quad 1 \le i \le N.$$

Hence, the system of equations in (22), represents a pseudospectral scheme for the costate equation based on polynomials of degree N. In Garg et al. (2010) we have also shown that

$$\boldsymbol{\lambda}_0 = -\mathbf{D}_0^{\mathsf{T}}\boldsymbol{\Lambda} = \sum_{j=1}^{N} w_j T_j \nabla_x H(\mathbf{X}_j, \mathbf{U}_j, \boldsymbol{\lambda}_j). \tag{24}$$

On the other hand, because Eq. (6) holds, the continuous costate satisfies

$$\boldsymbol{\lambda}(-1) = \boldsymbol{\lambda}(+1) - \int_{-1}^{+1} \dot{\boldsymbol{\lambda}}(\tau) d\tau$$
$$= \int_{-1}^{+1} T(\tau)\nabla_x H(\mathbf{x}(\tau), \mathbf{u}(\tau), \boldsymbol{\lambda}(\tau)) d\tau. \tag{25}$$

The right side of Eq. (24) represents a quadrature approximation to the right side of Eq. (25). Hence, Eq. (20) is in fact an approximation to the continuous-time initial costate. We refer to the LG collocation method developed in this section as the infinite-horizon version of the *Gauss pseudospectral method* (Benson, 2004; Benson et al., 2006; Garg et al., 2010; Huntington, 2007; Rao et al., 2010)

4.2. Infinite-horizon Radau pseudospectral method

Consider the LGR collocation points $-1 = \tau_1 < \cdots < \tau_N < +1$, and the additional *noncollocated* point $\tau_{N+1} = 1$. The state is then approximated by a polynomial of degree at most N as

$$\mathbf{x}(\tau) \approx \sum_{j=1}^{N+1} \mathbf{X}_j L_j(\tau),$$
$$L_j(\tau) = \prod_{\substack{k=1\\k\neq j}}^{N+1} \frac{\tau - \tau_k}{\tau_j - \tau_k}, \quad j = 1, \ldots, N+1, \tag{26}$$

where L_j is a basis of Nth-degree Lagrange polynomials. For the Radau scheme, the Lagrange interpolation points are $\tau_1 = -1$ through $\tau_{N+1} = +1$, while for the Gauss scheme, the interpolation points were $\tau_0 = -1$ through $\tau_N < +1$. Again, differentiating the series Eq. (26) and evaluating at the collocation point τ_i gives

$$\dot{\mathbf{x}}(\tau_i) \approx \sum_{j=1}^{N+1} \mathbf{X}_j \dot{L}_j(\tau_i) = \sum_{j=1}^{N+1} D_{ij}\mathbf{X}_j = \mathbf{D}_i\mathbf{X}, \quad i = 1, \ldots, N, \tag{27}$$

where

$$D_{ij} = \dot{L}_j(\tau_i) \quad \text{and} \quad \mathbf{X} = \begin{bmatrix} \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_{N+1} \end{bmatrix}.$$

Unlike the LG scheme of the previous section, the state $\mathbf{X}_{N+1}$ at the horizon appears in the state discretization Eq. (27). The discrete approximation to the control problem, using LGR scheme of this section, is almost the same as the LG scheme Eq. (13) except for the absence of the quadrature equation and the index on the initial condition $\mathbf{X}_1$ below corresponding to $\tau_1 = -1$:

$$\begin{aligned} &\min \sum_{i=1}^{N} w_i T(\tau_i) g(\mathbf{X}_i, \mathbf{U}_i) \\ &\text{subject to} \\ &T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) - \mathbf{D}_i\mathbf{X} = \mathbf{0}, \quad 1 \le i \le N \\ &\mathbf{x}_0 - \mathbf{X}_1 = \mathbf{0}. \end{aligned} \tag{28}$$

As with the Gauss pseudospectral method, it is important to observe that the left-hand side of the collocation equations contains approximations for the state at the LGR points plus the terminal point $\tau = +1$, which corresponds to $t = +\infty$. The right-hand side of the collocation equations contains approximations of the state at only the LGR points. Moreover, because the state $\mathbf{X}$ in Eq. (28) contains an additional component $\mathbf{X}_{N+1}$ corresponding to $\tau_{N+1} = +1$, the state at the horizon is a variable in our Radau state discretization. Again, we point out that the singularity in $T(\tau)$ at $\tau = +1$ is avoided in Eq. (28) since we evaluate T at the collocation points τ_i, $1 \le i \le N$, where $\tau_N < 1$.

In order to better relate the Radau discretization to the continuous control problem, we utilize the Lagrangian

$$\mathcal{L}(\boldsymbol{\Lambda}, \mathbf{X}, \mathbf{U}) = \langle \boldsymbol{\mu}, \mathbf{x}_0 - \mathbf{X}_1 \rangle + \sum_{i=1}^{N} \Big(w_i T(\tau_i) g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, T(\tau_i)\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) - \mathbf{D}_i\mathbf{X} \rangle \Big).$$

The optimality conditions, obtained by differentiating the Lagrangian with respect to the states $\mathbf{X}_1, \ldots, \mathbf{X}_{N+1}$ and the controls $\mathbf{U}_1, \ldots, \mathbf{U}_N$, are

$$\mathbf{D}_i^{\mathsf{T}}\boldsymbol{\Lambda} = T(\tau_i)\nabla_x \left(w_i g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \rangle \right) - \delta_{1i}\boldsymbol{\mu}, \tag{29}$$
$$\mathbf{0} = \mathbf{D}_{N+1}^{\mathsf{T}}\boldsymbol{\Lambda}, \tag{30}$$
$$\mathbf{0} = T(\tau_i)\nabla_u \left(w_i g(\mathbf{X}_i, \mathbf{U}_i) + \langle \boldsymbol{\Lambda}_i, \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i) \rangle \right), \tag{31}$$

where $1 \le i \le N$. The $\boldsymbol{\mu}$ term only enters into the first equation corresponding to differentiation with respect to $\mathbf{X}_1$. The boundary condition $\mathbf{D}_{N+1}^{\mathsf{T}}\boldsymbol{\Lambda} = \mathbf{0}$ arises from differentiating the Lagrangian with respect to $\mathbf{X}_{N+1}$.

In a manner similar to that for the infinite-horizon Gauss pseudospectral method, the first-order optimality conditions of Eqs. (29)–(31) arising from the infinite-horizon Radau pseudospectral method can be reformulated so that they resemble the first-order optimality conditions for the continuous control problem Eq. (5). First, we introduce the following expressions:

$$\begin{aligned} \boldsymbol{\lambda} &= \mathbf{W}^{-1}\boldsymbol{\Lambda}, \\ \boldsymbol{\lambda}_{N+1} &= \mathbf{D}_{N+1}^{\mathsf{T}}\boldsymbol{\Lambda}, \\ \mathbf{D}^{\dagger} &= -\mathbf{W}^{-1}\mathbf{D}_{1:N}^{\mathsf{T}}\mathbf{W} - \frac{1}{w_1}\mathbf{e}_1\mathbf{e}_1^{\mathsf{T}}, \end{aligned} \tag{32}$$

where $\mathbf{e}_1$ is the first column of the identity matrix. Substituting Eq. (32) into the conditions of Eqs. (29)–(31) gives

$$\mathbf{D}_i^{\dagger}\boldsymbol{\lambda} = -T(\tau_i)\nabla_x H(\mathbf{X}_i, \mathbf{U}_i, \boldsymbol{\lambda}_i) + \frac{\delta_{1i}}{w_1}(\boldsymbol{\mu} - \boldsymbol{\lambda}_1), \tag{33}$$
$$\boldsymbol{\lambda}_{N+1} = \mathbf{0}, \tag{34}$$
$$\mathbf{0} = \nabla_u H(\mathbf{X}_i, \mathbf{U}_i, \boldsymbol{\lambda}_i), \quad 1 \le i \le N. \tag{35}$$

The $\boldsymbol{\lambda}_1$ term in Eq. (33) emerges from the $\mathbf{e}_1$ term in the definition of $\mathbf{D}^{\dagger}$. Moreover, from the definition of $\boldsymbol{\lambda}_{N+1}$, the following identity can be derived (see Garg et al., 2009, for the details):

$$\boldsymbol{\mu} = \boldsymbol{\lambda}_{N+1} + \sum_{j=1}^{N} w_j T_j \nabla_x H(\mathbf{X}_i, \mathbf{U}_i, \boldsymbol{\lambda}_i). \tag{36}$$

Eq. (36) represents a quadrature approximation to the fundamental theorem of calculus Eq. (25). The right-hand side is the quadrature approximation to the costate $\boldsymbol{\lambda}_1$ at the initial time. Consequently, Eq. (36) is a subtle way of enforcing the equality $\boldsymbol{\mu} = \boldsymbol{\lambda}_1$ in Eq. (33). In Garg et al. (2009), we also show that $\mathbf{D}^{\dagger}$ is a differentiation matrix for polynomials of degree $N-1$; more precisely, if p is a polynomial of degree at most $N-1$ with values $p_i = p(\tau_i)$, $1 \le i \le N$, then

$$(\mathbf{D}^{\dagger}\mathbf{p})_i = \dot{p}(\tau_i), \quad 1 \le i \le N.$$

Hence, the system of equations in Eq. (33), represents a pseudospectral scheme for the costate equation based on polynomials of degree $N-1$. We refer to the LGR collocation method developed in this section as the infinite-horizon version of the *Radau pseudospectral method*.

4.3. Integrated forms

In formulating the pseudospectral schemes, the left side of the state equations in Eqs. (13) and (28) contained the derivatives of Lagrange polynomials. As shown in Garg et al. (2009) and Garg et al. (2010), we can invert the nonsingular part of the differentiation matrix $\mathbf{D}$ to write the discrete dynamics in the form

$$\mathbf{X}_i = \mathbf{x}_0 + \sum_{j=1}^{N} A_{ij} T_j \mathbf{f}(\mathbf{X}_j, \mathbf{U}_j), \tag{37}$$

where $1 \leq i \leq N$ for LG collocation and $2 \leq i \leq N+1$ for LGR collocation. Here the matrix elements A_{ij} can be expressed as the integrals of Lagrange interpolating polynomials associated with the collocation points. More precisely, for LG collocation, we have

$$A_{ij} = \int_{-1}^{\tau_i} L_j^{\dagger}(\tau) \mathrm{d}\tau, \qquad L_j^{\dagger} = \prod_{\substack{k=1 \\ k \neq j}}^{N} \frac{\tau - \tau_k}{\tau_j - \tau_k}, \quad i, j = 1, \ldots, N. \tag{38}$$

For LGR collocation, the right side of Eq. (38) defines $A_{i+1,j}$. Computationally, the differential formulation of Eq. (10) of the system dynamics is more convenient since any nonlinear terms in $\mathbf{f}$ retain their sparsity in the discretization, while for the integrated version of Eq. (37), the nonlinear terms are nonsparse due to multiplication by the dense matrix $\mathbf{A}$.

5. Examples

We now consider two examples of the pseudospectral methods developed in Section 4. In the first example we perform both an error analysis and a computation time comparison. In the second, more complex example, we focus on the computation time required to obtain a solution using each method.

5.1. Infinite-horizon control of a one-dimensional nonlinear system

Consider the following nonlinear infinite-horizon optimal control problem:

$$\min J = \frac{1}{2} \int_0^{\infty} \left(\log^2 y(t) + u(t)^2\right) \mathrm{d}t$$

subject to

$$\dot{y}(t) = y(t) \log y(t) + y(t) u(t), \quad y(0) = 2. \tag{39}$$

The exact solution to this problem is

$$y^*(t) = \exp(x^*(t)), \qquad u^*(t) = -\alpha x^*(t),$$
$$\lambda^*(t) = \alpha \exp(-x^*(t)) x^*(t), \qquad x^*(t) = 2\exp(-t\sqrt{2}), \tag{40}$$

where $\alpha = 1 + \sqrt{2}$. The example of Eq. (39) was solved for $N = (5, 10, 15, 20, 25, 30)$ using both the infinite-horizon pseudospectral methods described in Section 4 and the approach of Fahroo and Ross (2008) with the three strictly monotonic transformations of the domain $\tau \in [-1, +1)$ given in Eqs. (2)–(4). All solutions were obtained using an Apple MacBook Pro running Mac OS-X version 10.6.4 with 4 GB of DDR3 1 GHz memory, MATLAB Version 7.10.0.499 (R2010a), and the NLP solver SNOPT (Gill, Murray, & Saunders, 2002, 2005) with optimality and feasibility tolerances of 1×10^{-10} and 2×10^{-10}, respectively. Furthermore, the following initial guess of the solution was used:

$$y(\tau) = y_0, \qquad u(\tau) = \tau, \quad \tau \in [-1, 1]. \tag{41}$$

The maximum base ten logarithm of the state, the control, and the costate errors are defined as

$$E_y = \max_k \log_{10} \left|\mathbf{Y}_k - y^*(\tau_k)\right|$$
$$E_u = \max_k \log_{10} \left|\mathbf{U}_k - u^*(\tau_k)\right| \tag{42}$$
$$E_\lambda = \max_k \log_{10} \left|\boldsymbol{\lambda}_k - \lambda_y^*(\tau_k)\right|.$$

In Eq. (42) the index k spans the approximation points in the case of either the state and the costate and spans only the collocation points in the case of the control. We remind the reader that the state and the costate obtained on the *entire horizon* with the index $N+1$ corresponding to the state and the costate at $\tau = +1$, or equivalently, at $t = +\infty$.

The errors obtained using the Gauss and Radau methods of this paper are shown, respectively, in Figs. 1 and 2 alongside the error obtained using the method of Fahroo and Ross (2008) with the transformation given in Eqs. (2) and (4). It is seen for all three transformations and for both methods of this paper, the state, the control, and the costate errors decrease in essentially a linear manner until $N = 30$, demonstrating an approximately exponential convergence rate. Furthermore, it is observed that either the Gauss or Radau method of this paper yields approximately the same error for a particular value of N and choice of transformation. Moreover, it is seen that the errors are largest and smallest, respectively, using the transformations of Eqs. (3) and (4). In fact, the transformation of Eq. (4) is at least one order of magnitude more accurate than either of the other two transformations. Finally, it is seen that the errors from the two methods of this paper using the transformation of Eq. (4) are significantly smaller than those obtained using the method of Fahroo and Ross (2008) (where the transformation of Eq. (2) are used). When the transformation of Eq. (4) is used, however, the state errors from the method of Fahroo and Ross (2008) are nearly the same as those obtained using the Gauss and Radau methods, while the control and the costate errors are approximately one order of magnitude larger using the method of Fahroo and Ross (2008).

We now analyze the computational performance of each method for this example. Fig. 3 shows the CPU time as a function of N for the different methods, where it is seen that the change of variables given in either $t = \phi_b(\tau)$ or $t = \phi_c(\tau)$ Eqs. (3) and (4), respectively, are more computationally efficient than the change of variables $t = \phi_a(\tau)$ given in Eq. (2). More specifically, it is seen that for $N > 5$ using the change of variables $t = \phi_a(\tau)$ nearly doubles the required CPU time when compared with either of the other two transformations. Next, Fig. 3 also suggests that the CPU time is not related to the choice of the collocation method since the time to solve the discrete problem using either the Gauss or Radau scheme of this paper or the Radau scheme of Fahroo and Ross (2008) is about the same. Instead, the CPU time depends on the choice of the change of variables $t = \phi(\tau)$.

The different behavior of the functions given in Eqs. (2)–(4) is understood if we apply the change of variables to the continuous solution. The optimal state in the transformed coordinates is as follows:

$$y_a(\tau) = \exp\left(\exp\left(-2\sqrt{2}\left(\frac{1+\tau}{1-\tau}\right)\right)\right)$$
$$y_b(\tau) = \exp\left(\left(\frac{1-\tau}{2}\right)^{2\sqrt{2}}\right)$$
$$y_c(\tau) = \exp\left(\frac{(1-\tau)^{4\sqrt{2}}}{4^{2\sqrt{2}}}\right).$$

Here the subscripts a, b, and c correspond to the three choices of ϕ given in Eqs. (2)–(4). An advantage of using a logarithmic change of variables given in Eq. (3) or (4), as compared to the function given in Eq. (2), is that logarithmic functions essentially move collocation points associated with large values of t to the left. Because the exact solution changes slowly when t is large, this leftward movement of the collocation points is beneficial since more collocation points are situated where the solution is changing most rapidly. The disadvantage of a logarithmic change of variables is seen in the function $\log(1-\tau)$ where the growth is so slow near $\tau = +1$ that

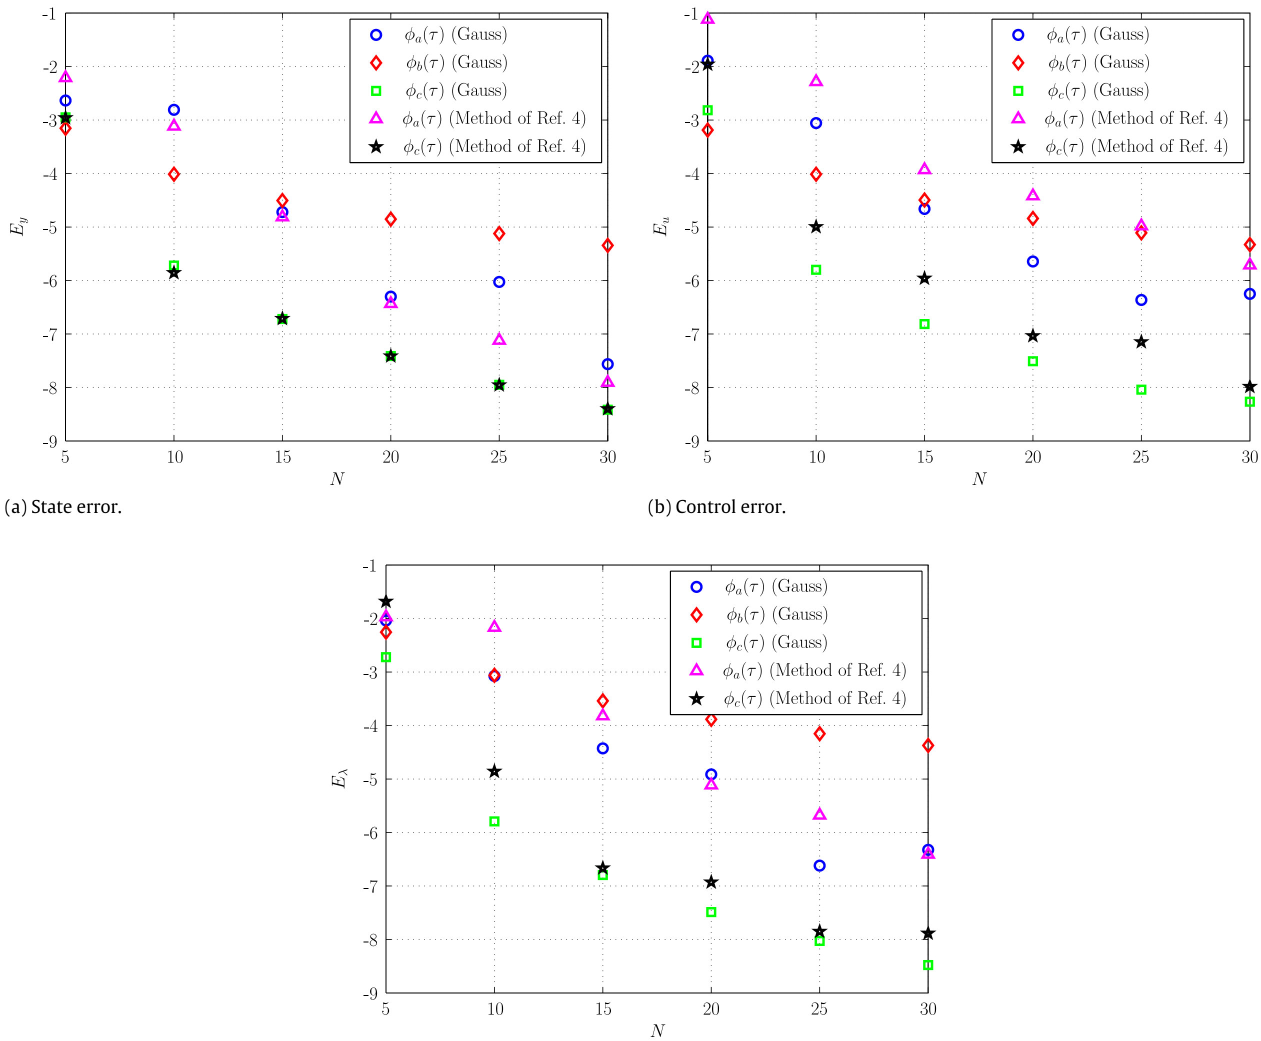

(a) State error.

(b) Control error.

(c) Costate error.

Fig. 1. Maximum Gauss pseudospectral state, control, and costate errors for Example 1.

the transformed solution possesses a singularity in a derivative at $\tau = +1$. In other words, the j-th derivative of a function of the form $(1-\tau)^\alpha$, where $\alpha > 0$ is not an integer, is singular at $\tau = +1$ for $j > \alpha$. In particular, $y_b(\tau)$ has two derivatives at $\tau = +1$ but not three, while $y_c(\tau)$ has five derivatives at $\tau = +1$ but not six. To achieve exponential convergence, $y(\tau)$ should be infinitely smooth. For this particular problem, the choice of Eq. (4) has the following nice properties: $y_c(\tau)$ is relatively smooth with five derivatives, although not infinitely smooth, and collocation points corresponding to large values of t, where the solution changes slowly, are moved to the left [when compared to $t = (1+\tau)/(1-\tau)$] where the solution changes more rapidly. As a result, for $5 \le N \le 30$, the function of Eq. (4) yields a solution that is often two or more orders of magnitude more accurate than the other choices for ϕ.

5.2. Infinite-horizon control of an inverted pendulum

Consider the following optimal control problem. Mills, Willis, and Ninness (2009) Minimize the cost functional

$$J = \int_0^\infty \left(p^2 + 10\theta^2 + 10^{-4}v^2 + 10^{-4}\omega^2 + 0.1u^2\right) dt$$

subject to (43)

$$\dot{p} = v, \qquad \dot{\theta} = \omega,$$

$$\dot{v} = \frac{a_1 w_1 + w_2 \cos\theta}{d}, \qquad \dot{\omega} = \frac{w_1\cos\theta + a_2 w_2}{d},$$

with the initial conditions

$$\begin{aligned} p(0) &= 0, \qquad \theta(0) = 225\ \text{deg}, \\ v(0) &= 0, \qquad \omega(0) = 0, \end{aligned} \qquad (44)$$

where $w_1 = -k_1 u - \omega^2 \sin\theta - k_2 v$, $w_2 = g\sin\theta - k_3\omega$, $d = b - \cos^2\theta$, and $(a_1, a_2, k_1, k_2, g, b) = (5, 1, 1/2, -1/2, 9.81, 5)$. The example of Eq. (42)–(44) was solved for $N = (5, 10, 15, 20, 25, 30)$ using both the infinite-horizon pseudospectral methods described in Section 4 and the approach of Fahroo and Ross (2008) with the three strictly monotonic transformations of the domain $\tau \in [-1, +1)$ given in Eqs. (2)–(4). All solutions were obtained using an Apple MacBook Pro running Mac OS-X version 10.6.4 with 4GB of DDR3 1 GHz memory, MATLAB Version 7.10.0.499 (R2010a), and the NLP solver SNOPT (Gill et al., 2002, 2005) with default optimality and feasibility tolerances. Furthermore, a straight line interpolation between the initial conditions and zero was used as the initial guess for the state, while a guess of zero was used for the control.

A representative numerical solution (obtained using the Radau pseudospectral method with $N = 40$) is shown in Fig. 4(a). As expected, both the state and the control approach zero as $\tau \longrightarrow +1$. Because this problem does not have an analytic solution, we focus the remainder of our attention on a comparison

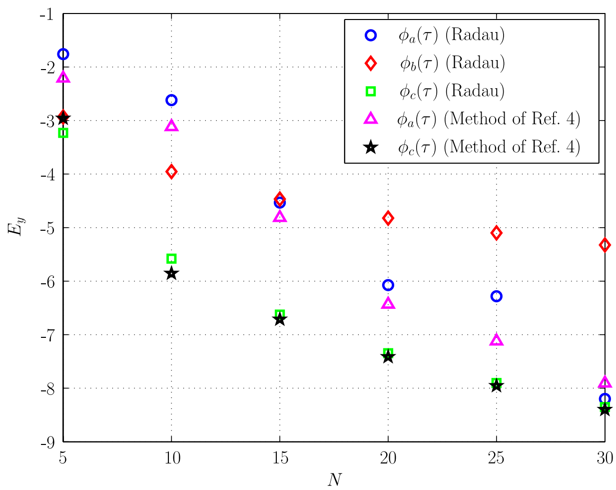

(a) State error.

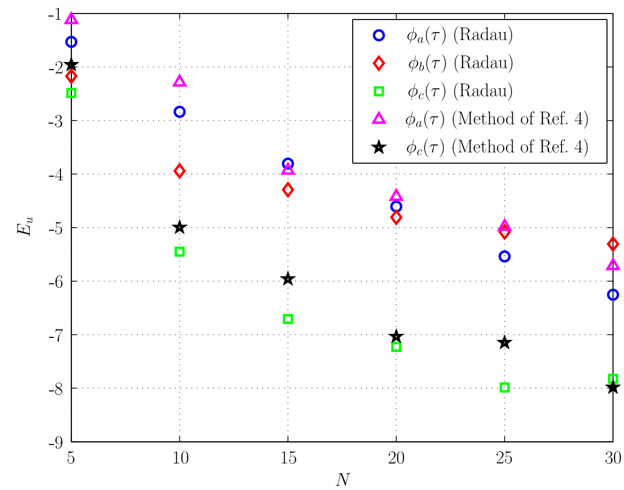

(b) Control error.

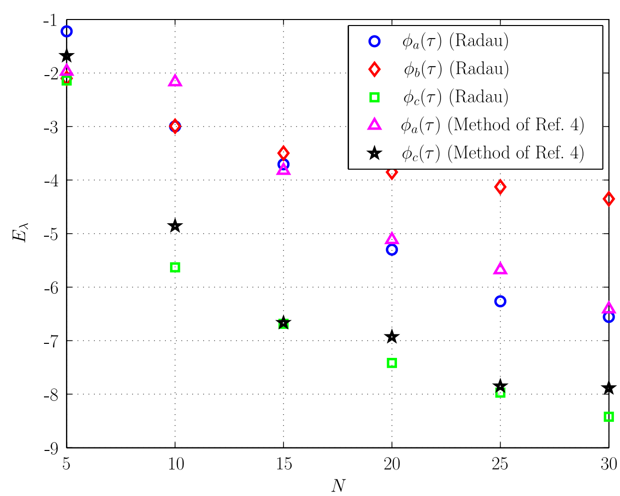

(c) Costate error.

Fig. 2. Maximum Radau pseudospectral state, control, and costate errors for Example 1.

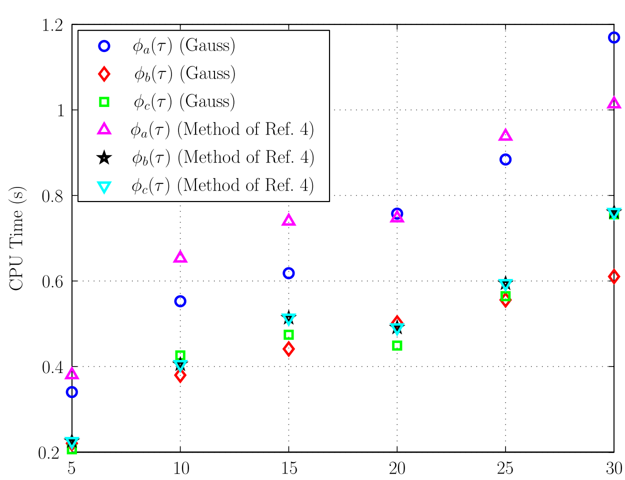

(a) CPU times using Gauss pseudospectral method and the method of Fahroo and Ross (2008).

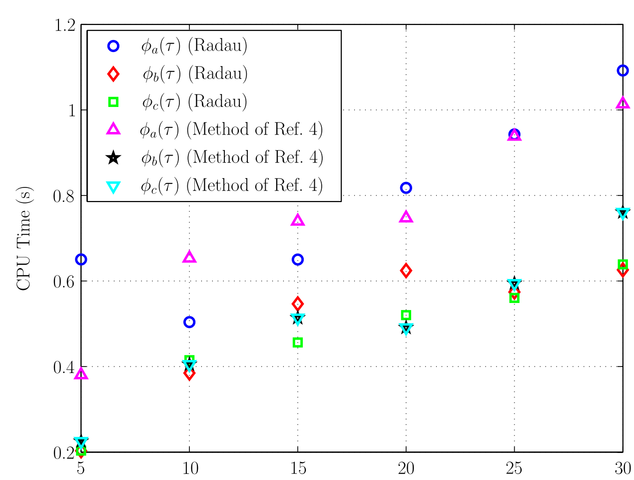

(b) CPU times using Radau pseudospectral method and the method of Fahroo and Ross (2008).

Fig. 3. CPU times required to solve Example 1.

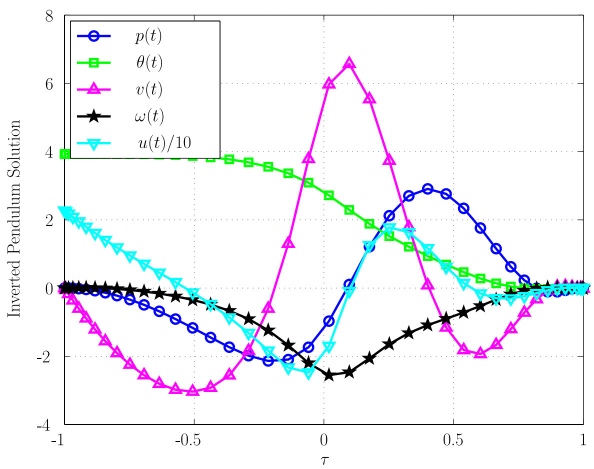

(a) Representative solution to Example 2 using the Radau pseudospectral method with $n = 40$.

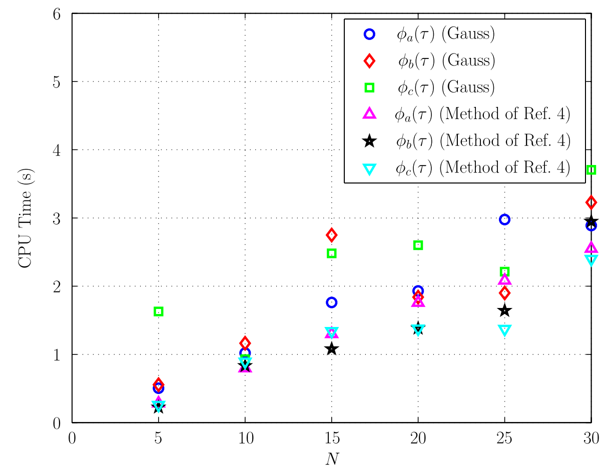

(b) CPU times using both the infinite-horizon Gauss pseudospectral method and the method of Fahroo and Ross (2008).

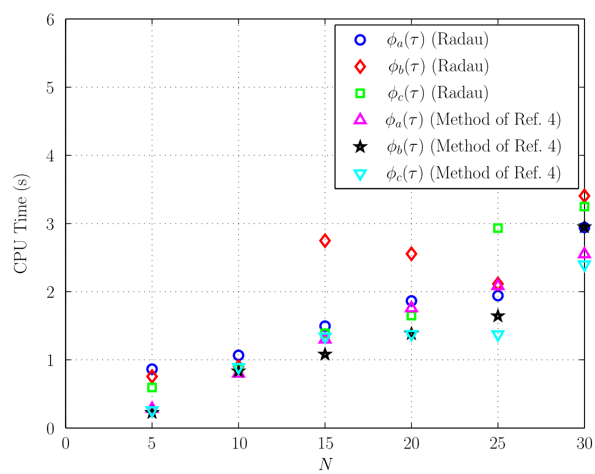

(c) CPU times using both the infinite-horizon Radau pseudospectral method and the method of Fahroo and Ross (2008).

Fig. 4. Representative solution and CPU times for Example 2.

of the computation time required to solve the problem using each method. Specifically, Fig. 4 shows the CPU time as a function of N for all three transformations $\phi_a(\tau)$, $\phi_b(\tau)$, and $\phi_c(\tau)$ using both methods developed in Section 4 and using the method of Fahroo and Ross (2008). First, it is seen that, regardless of method or transformation, the CPU time grows essentially linearly as a function of N. Second, it is observed that, for some values of N, the CPU time required by the two methods of this paper is larger than the CPU time required by the method Fahroo and Ross (2008). Even in these cases, however, it is noted that the trend is not consistent. For example, the CPU time using the Gauss pseudospectral method with the transformation $\phi_b(\tau)$ and $N = 15$ is approximately twice that required by the method of Fahroo and Ross (2008), whereas for $N = 10$ the CPU times for all methods and transformations is nearly the same. In the case of the Radau pseudospectral method, it is seen that the CPU times are nearly the same except for certain transformations (for example, the CPU time using Radau with $N = 15$ and $\phi_b(\tau)$ is twice as large as all other CPU times for $N = 15$). Finally, because the CPU times are all the same order of magnitude, any computational performance penalty is compensated by the increased accuracy that may be offered by the Gauss and Radau schemes of our paper.

6. Conclusions

Two pseudospectral methods have been presented for the numerical solution of nonlinear infinite horizon optimal control problems using global collocation at Legendre–Gauss and Legendre–Gauss–Radau points. It was shown that the nonlinear programming problems which arise from a change of variables followed by either Gauss or Radau collocation includes an approximation to the state at $t = +\infty$. The Legendre–Gauss and Legendre–Gauss–Radau transformed adjoint systems connecting the KKT conditions of the nonlinear programming problem to the Pontryagin minimum principle were then derived. These transformed adjoint systems resulted in approximations for the costate at $t = +\infty$. Finally, it was shown that either of the methods developed in this paper can be written equivalently in either a differential or integral form. The results of this paper indicate that the use of Legendre–Gauss and Legendre–Gauss–Radau points lead to accurate approximations to a continuous nonlinear infinite-horizon optimal control problem in such a manner that the solution is obtained on the entire infinite-horizon. By tuning the change of variables used to map the infinite time domain $[0, \infty)$ to a finite interval, it was possible to improve the accuracy in the discrete approximation by several orders of magnitude.

References

Benson, D. A. (2004). A Gauss pseudospectral transcription for optimal control. *Ph.D. dissertation*. MIT.

Benson, D. A., Huntington, G. T., Thorvaldsen, T. P., & Rao, A. V. (2006). Direct trajectory optimization and costate estimation via an orthogonal collocation method. *Journal of Guidance, Control, and Dynamics*, *29*(6), 1435–1440.

Elnagar, G., Kazemi, M., & Razzaghi, M. (1995). The pseudospectral Legendre method for discretizing optimal control problems. *IEEE Transactions on Automatic Control*, *40*(10), 1793–1796.

Elnagar, G. N., & Razzaghi, M. (1997). A collocation-type method for linear quadratic optimal control problems. *Optimal Control Applications and Methods*, *18*(3), 227–235.

Fahroo, F., & Ross, I. M. (2008). Pseudospectral methods for infinite-horizon nonlinear optimal control problems. *Journal of Guidance, Control, and Dynamics*, *31*(4), 927–936.

Garg, D., Patterson, M. A., Darby, C. L., Francolin, C., Huntington, G. T., Hager, W. W., et al. (2009). Direct trajectory optimization and costate estimation of finite-horizon and infinite-horizon optimal control problems using a radau pseudospectral method. *Computational Optimization and Applications*, doi:10.1007/s10589-009-9291-0.

Garg, D., Patterson, M. A., Hager, W. W., Rao, A. V., Benson, D. A., & Huntington, G. T. (2010). A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*, *46*(11), 1843–1851.

Gill, P. E., Murray, W., & Saunders, M. A. (2002). SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, *12*, 979–1006.

Gill, P. E., Murray, W., & Saunders, M. A. (2005). SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Review*, *47*, 99–131.

Hager, W. W. (2000). Runge–Kutta methods in optimal control and the transformed adjoint system. *Numerische Mathematik*, *87*, 247–282.

Huntington, G. T. (2007). Advancement and analysis of a Gauss pseudospectral transcription for optimal control. *Ph.D. dissertation*. Department of Aeronautics and Astronautics. Massachusetts Institute of Technology.

Mills, A., Willis, A., & Ninness, B. (2009). Nonlinear model predictive control of an inverted pendulum. In *Proceedings of the 2009 American control conference. St. Louis, Missouri. July 10–12* (pp. 2335–2340).

Rao, A. V., Benson, D. A., Darby, C., Patterson, M. A., Francolin, C., Sanders, I., et al. (2010). Algorithm 902: GPOPS, a MATLAB software for solving multiple-phase optimal control problems using the Gauss pseudospectral method. *ACM Transactions on Mathematical Software*, *37*(2), 22:1–22:39.

Williams, P. (2004). Jacobi pseudospectral method for solving optimal control problems. *Journal of Guidance, Control, and Dynamics*, *27*(2), 293–297.

Divya Garg received a Bachelor of Technology degree in Mechanical Engineering from the Malaviya National Institute of Technology, Jaipur, India, in 2007. Since 2007 she has been a graduate student in the Department of Mechanical and Aerospace Engineering at the University of Florida where in 2008 she earned a Master of Science degree and is currently a Ph.D. student.

William W. Hager received a B.S. degree in mathematics from Harvey Mudd College, Claremont, CA, in 1970 and Master's and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in 1971 and 1974, respectively. He was awarded a National Science Foundation Fellowship for graduate studies in mathematics. He was an Assistant Professor at the University of South Florida, Tampa, from 1974 to 1976, an Assistant Professor at Carnegie Mellon University, Pittsburgh, PA, from 1976 to 1980, and an Associate Professor from 1980 to 1986 and a Professor from 1986 to 1988 at Pennsylvania State University, University Park. Since 1988, he has been a Professor at the University of Florida, Gainesville, and since 1992, he has been Co-Director of the Center for Applied Optimization at the University of Florida. He was Co-Chair of the 2001 SIAM Conference on Control and Its Applications. He was Program Director of the SIAM Activity Group on Control and System Theory from 1998 to 2001, and he has been Editor-in-Chief of Computational Optimization and Applications since 1992.

Anil V. Rao earned an A.B. in mathematics and a B.S. in mechanical engineering from Cornell University in 1988, an M.S.E. in aerospace engineering from the University of Michigan in 1989, and an M.A. and Ph.D. in mechanical and aerospace engineering from Princeton University in 1992 and 1996, respectively. After earning his Ph.D., Dr. Rao was a Senior Member of the Technical Staff at The Aerospace Corporation in Los Angeles, California and a Senior Member of the Technical Staff at The Charles Stark Draper Laboratory in Cambridge, Massachusetts. Since 2006 he has been an Assistant Professor of Mechanical and Aerospace Engineering at The University of Florida. He was Technical Chair of the 2009 Astrodynamics Specialist Conference, is an Associate Editor for The Journal of the Astronautical Sciences and the Journal of Optimization Theory and Applications, and is an Associate Fellow of the American Institute of Aeronautics and Astronautics.