

Polynomial mechanics and optimal control

Akshay Srinivasan¹, Madhusudhan Venkadesan²

Abstract—We describe a new algorithm for trajectory optimization of mechanical systems. Our method combines pseudo-spectral methods for function approximation with variational discretization schemes that exactly preserve conserved mechanical quantities such as momentum. We thus obtain a global discretization of the Lagrange-d’Alembert variational principle using pseudo-spectral methods. Our proposed scheme inherits the numerical convergence characteristics of spectral methods, yet preserves momentum-conservation and symplecticity after discretization. We compare this algorithm against two other established methods for two examples of underactuated mechanical systems; minimum-effort swing-up of a two-link and a three-link acrobat.

I. INTRODUCTION

Trajectory optimization methods are broadly classified into two categories: Direct, and Indirect [1]. Direct methods discretize both state and control trajectories to derive a finite-dimensional constrained optimization problem, while indirect methods solve the discretized nonlinear equations resulting from the necessary first-order variational conditions of Euler-Lagrange-Pontryagin [2]. Direct methods more convenient for the non-specialized practitioner especially because of the difficulty in deriving conditions equivalent to Pontryagin’s when the problem involves free parameters or inequality constraints [1].

Pseudo-spectral implementations of direct methods have seen increasing use in recent years because of their super-polynomial convergence [3] [4] [5] [6]. While these methods are applicable for a broad class of systems, they tend however not to preserve certain geometric structures that are associated with conservation laws for mechanical systems.

The formulation of *Discrete-Mechanics and Optimal Control* (DMOC) satisfies the latter requirement by using variational integrators to construct schemes which are both symplectomorphic and momentum-conserving [7] [8]. These methods however lack the convergence and approximation properties of the former pseudo-spectral methods and only exhibit a fixed-order algebraic convergence.

The results from the theory of Geometric integration [9] [10], are not readily converted to pseudo-spectral discretizations because the polynomial bases only yield quadratures for weighted integrals. Furthermore, the theory only provides

statements for the map between initial and end-time values; the global discretization is essentially one step of an ODE integrator.

In this paper, we propose a direct method which incorporates the benefits of both pseudo-spectral and DMOC methods into a single algorithm. We extend the analysis of DMOC, using variations over polynomials, to derive a discrete non-causal analogue of the Euler-Lagrange equation. We also prove that the resulting discretization is both symplectomorphic and momentum-conserving.

In the following sections, we make use of specialized notation which serves the dual purpose of being both intuitive and reasonably precise. The precise definitions can be found in the appendix (VI-A).

II. PROBLEM STATEMENT

1) *Dynamics*: We restrict our attention in this paper to Lagrangian systems. Given the Lagrangian $\mathcal{L}(q, \bar{v})$, the corresponding dynamics is generated by the condition that every trajectory-curve of the system $q : [0, t_f] \rightarrow \mathcal{M}$, satisfy the Lagrange-d’Alembert principle,

$$\begin{aligned} \mathcal{L} : T\mathcal{M} \rightarrow \mathbb{R}, \quad \mathcal{L}|_q : T\mathcal{M}_q \rightarrow \mathbb{R} \text{ is convex.} \\ \delta_{\bar{q}} \left[\int_0^{t_f} \mathcal{L}(q(t), D_t \bar{q}(t)) dt \right] + \int_0^{t_f} \underline{u}(t) \delta \bar{q}(t) dt = \underline{p} \bar{q}|_0^{t_f}, \\ \forall \delta \bar{q} \in (\mathbb{R}[t]_N)^{\dim(\mathcal{M})}, \\ \text{where, } \underline{p}(t) := \partial_v \mathcal{L}(q(t), \bar{q}(t)). \end{aligned} \quad (1)$$

Using variational arguments, it can be shown that every such (smooth) solution also satisfies the Euler-Lagrange equations,

$$D_t \begin{bmatrix} q(t) \\ \partial_v \mathcal{L}(q(t), \bar{q}(t)) \end{bmatrix} = \begin{bmatrix} \bar{q}(t) \\ \partial_q \mathcal{L}(\bar{q}(t), \bar{q}(t)) + \underline{u}(t) \end{bmatrix}. \quad (2)$$

2) *Control*: The optimal control problem is defined as finding a control sequence which incurs the least cost (application-specific), while respecting the dynamics defined by (1),

$$\begin{aligned} \underline{u}^* &= \arg \min_{\underline{u}} J(q_0, \underline{u}), \\ J(q_0, \underline{u}) &= \int_0^{t_f} l([q(t), D_t \bar{q}(t)], \underline{u}(t)) dt + \\ &\quad V_f([q(t_f), D_t \bar{q}(t_f)]), \\ &\text{s. t } q(0) = q_0, \\ & q|_{[0, t_f]}, \underline{u}|_{[0, t_f]} \text{ satisfies the conditions of (1).} \end{aligned} \quad (3)$$

*This work was funded by grants to M. Venkadesan from the Human Frontier Science Program (RGY0091), the National Centre for Biological Sciences and the Simons Foundation. A. Srinivasan was supported by grants from NSF and DARPA, awarded to Emanuel Todorov, Department of Computer Science & Engineering, University of Washington, Seattle.

¹ Department of Computer Science & Engineering, University of Washington, Seattle, WA-98195, USA, akshays@cs.washington.edu

² Department of Mechanical Engineering & Materials Science, Yale University, New Haven, CT 06511, USA, m.venkadesan@yale.edu

III. PRELIMINARIES

A. Orthogonal Polynomials

The pseudo-spectral approach uses a finite-degree polynomial basis to describe control and state trajectories. Although our implementation makes use of Chebyshev polynomials, the discretization itself is equally applicable to every orthogonal-polynomial base on a closed interval. We briefly summarize the general theory of orthogonal polynomials.

Let $\mathbb{R}[t]$ denote the vector space of polynomials with real-coefficients, and let $\mathbb{R}[t]_n$ denote the n -dimensional vector space of all polynomials less than or equal to degree $(n-1)$.

Definition 1 (Lagrange interpolation polynomial): Given *grid-points*, $\{t_i\}_{i=1}^N$, ℓ_k is defined to be the unique polynomial in $\mathbb{R}[t]_N$ such that $\ell_k(t_i) = \delta_{ik}$,

$$\ell_k(t; \{t_j\}_{j=1}^N) = \prod_{i \neq k} \frac{t - t_i}{t_k - t_i}. \quad (4)$$

Definition 2 (Orthogonal polynomials): Given the inner product on the function space $L^2([-1, 1])$,

$$\langle f, g \rangle_w = \int_{-1}^1 w(t) f(t) g(t) dt. \quad (5)$$

The set of orthogonal polynomials \mathcal{P}_w corresponding to this metric is an ordered orthogonal basis for $\mathbb{R}[t]$ under this inner product.

$$\langle P_i, P_j \rangle_w = \gamma_i \delta_{ij},$$

$$i > j \iff \deg(P_i) > \deg(P_j), \quad \forall P_i, P_j \in \mathcal{P}_w, \quad (6)$$

where $\gamma_i \in \mathbb{R}_+$ is the normalization constant for P_i under the inner product.

Orthogonal polynomials also satisfy—and are numerically computed using—recurrence relations of the form [11],

$$\begin{aligned} P_0 &= a_0, & P_1 &= a_1 t \\ P_n(t) &= (a_n t + b_n) P_{n-1}(t) - c_n P_{n-2}(t). \end{aligned} \quad (7)$$

The specific polynomial bases used in our work, Chebyshev and Legendre polynomials, correspond to the weights $w(t) = \frac{1}{\sqrt{1-t^2}}$ and $w(t) = 1$, respectively.

1) *Gauss quadrature, Pullback inner-product:* The orthogonal polynomial set \mathcal{P}_w can be used to approximate integrals using the N -point Gauss quadrature scheme of order $2N$ [11],

$$\int_{-1}^1 w(t) f(t) dt \approx \sum_{k=1}^N w_k f(t_k). \quad (8)$$

The quadrature points $\{t_k\}_{k=1}^N$ are chosen to be the roots of the N 'th orthogonal polynomial. The roots and corresponding quadrature weights can be computed numerically from the eigenvalues and eigenvectors of the Jacobi operator, using the Golub-Welsch algorithm [12].

Because the order of the quadrature is $2N$, it is exact for all $p \in \mathbb{R}[t]_{2N}$. As a result, Gauss quadrature defines the

discrete inner product $\langle P_i, P_j \rangle_w^N$ over $\mathbb{R}[t]_N$,

$$\forall i, j < N,$$

$$\begin{aligned} \langle P_i, P_j \rangle_w &= \int_{-1}^1 w(t) P_i(t) P_j(t) dt \\ &= \sum_{k=1}^N w_k P_i(t_k) P_j(t_k) = \gamma_i \delta_{ij} \\ &:= \langle P_i, P_j \rangle_w^N. \end{aligned} \quad (9)$$

This connection ties collocation using these points, to the Galerkin method [11] and, crucially, to the quadrature itself,

$$\int_{-1}^1 w(t) p(t) dt = \langle 1, p \rangle_w^N \quad \forall p \in \mathbb{R}[t]_N. \quad (10)$$

Given polynomials $p, q \in \mathbb{R}[t]_N$, the integral of pq is given exactly by $\langle p, q \rangle_w^N$. Since every N -point (or basis) representation of a polynomial in $\mathbb{R}[t]_N$ is related to every other by a linear transformation, the pullback of the Legendre inner-product, $\langle \cdot, \cdot \rangle^{N*}$, preserves $\langle p, q \rangle$, even though the pullback of the quadrature does not,

$$\langle p, q \rangle = \langle 1, pq \rangle^N = \langle p, q \rangle^{N*} \neq \langle 1, pq \rangle^{N*}. \quad (11)$$

Seen in the coordinate free sense, the quadrature is generically only of order N . The order $2N$ quadrature is achieved only when the grid-points are restricted to the roots of P_N . The use of the pullback, on the other hand, allows one to represent polynomials over arbitrary grid points, and yet be able to use the discretized L^2 inner-product. The ability to do this is important since collocation on Chebyshev grids gives a $O(\log(N))$ approximation to the best-uniform-approximation polynomial [13], while that on the Legendre-grid has no such guarantee.

The loss of accuracy in quadrature is not a concern when using Chebyshev polynomials, because the resulting quadrature (associated with Clenshaw-Curtis) is known to be nearly as accurate as Gauss-Lobatto [14]. In the general case however it may perform poorly because of Runge's phenomenon [15].

B. Symplectic maps

The flow generated by Lagrange-d'Alembert systems (1) are symplectic when the control sequence is fixed in time. Discretizations which preserve this property of the system often exhibit desirable numerical properties [9].

There are many equivalent definitions for a diffeomorphism to be symplectic. The following is prevalently used on Darboux coordinates,

Definition 3 (Symplectic map): A diffeomorphism $f : (p, q) \mapsto (P, Q)$ is said to be symplectic if it leaves the symplectic form, J , invariant,

$$J = (\partial f)^T J \partial f, \quad J := \begin{bmatrix} 0 & -\text{id} \\ \text{id} & 0 \end{bmatrix}.$$

We shall however make use of the following equivalent condition in the upcoming sections [9, p. 196],

Lemma 1 (Total differential): A map $\phi : (p, q) \mapsto (P, Q)$ is symplectic if and only if, $Pd\bar{Q} - p\bar{d}q$, is a total differential in $\underline{dp}, \bar{d}q$.

IV. DISCRETE VARIATIONAL CONSTRAINT

A. Discrete Lagrange-d'Alembert

Let the Lagrangian of the system be given by $\mathcal{L}(q, \bar{v})$. Using the notation $\langle \underline{f}, \bar{g} \rangle := \int \underline{f}(t) \bar{g}(t) dt$ for the standard inner product on $L^2, \bar{1}$ is written as,

$$\begin{aligned} \delta_{\bar{q}} \langle \mathcal{L} \circ q, \bar{1} \rangle + \langle \underline{u}, \delta \bar{q} \rangle &= \underline{p} \delta \bar{q}|_0^{t_f}, \forall \delta \bar{q}, \\ \text{where, } \mathcal{L} \circ q(t) &:= \mathcal{L}(q(t), \bar{D}_t q(t)), \\ p(t) &:= \partial_v \mathcal{L}(q(t), \bar{D}_t q(t)). \end{aligned} \quad (12)$$

Using the inner-product on Legendre polynomials and collocation points $\{t_i\}$, this condition is approximated as,

$$\delta \langle \mathcal{L} \circ q, \bar{1} \rangle^{N*} + \langle \underline{u}, \delta \bar{q} \rangle^{N*} = \underline{p} \delta \bar{q}|_0^{t_f}, \forall \delta \bar{q} \in (\mathbb{R}[t]_N)^{\dim(\mathcal{M})}. \quad (13)$$

Expanding the variation, we find,

$$\langle \mathcal{L}_q + D^\dagger \mathcal{L}_v + \underline{u}, \delta \bar{q} \rangle^{N*} = \underline{p} \delta \bar{q}|_0^{t_f}, \forall \delta \bar{q} \in (\mathbb{R}[t]_N)^{\dim(\mathcal{M})}, \quad (14)$$

where $\underline{p}(t) := \sum_{i=1}^N \mathcal{L}_v^i \ell_i(t, \{t_k\})$.

Next, we prove that the map associated with trajectories satisfying (14), is both symplectomorphic and momentum-conserving.

B. Symplecticity, Momentum conservation.

Consider the dynamics described by (1). Given the end points $q_0, q_f \in \mathcal{M}$, there exists a unique polynomial $q^* \in (\mathbb{R}[t]_N)^{\dim(\mathcal{M})}$ that minimizes the discrete action defined in (13), the corresponding momentum polynomial is given by $\underline{p}^*(t) := \sum_{i=1}^N \mathcal{L}_v^i \ell_i(t, \{t_k\})$. The scheme (14) is symplectomorphic if the map $(\delta q(\tau'), \delta p(\tau')) \mapsto (\delta q(\tau), \delta p(\tau))$, is symplectic for all τ .

We define the discretized action over the interval $[\tau', \tau]$, for q^* to be,

$$S_{\tau', \tau}(q_{\tau'}^*, q_\tau^*) = \langle \mathcal{L} \circ q^*, \bar{1} \rangle_{\tau', \tau}^{N*}, \quad \tau', \tau \in [0, t_f]. \quad (15)$$

The pullback of the inner-product in the above equation is generated by linear affine transforms between the intervals $[\tau, \tau']$ and $[0, t_f]$. Note that because the polynomials are defined by their values on the original grid points, the discretized action $S_{\tau', \tau}$ as defined above, also depends on the values over grid points outside $[\tau', \tau]$. This lack of causality in the action leads to equivalent non-causal notions of symplecticity and momentum conservation.

Taking discrete variations of $S_{\tau', \tau}$ around q^* ,

$$\delta S_{\tau', \tau} = \langle \mathcal{L}_q + \underline{u}, \delta \bar{q} \rangle_{\tau', \tau}^{N*} + \langle \mathcal{L}_v, D \delta \bar{q} \rangle_{\tau', \tau}^{N*}. \quad (16)$$

Exploiting the fact that the inner-product, $\langle \mathcal{L}_v, D \delta \bar{q} \rangle_{\tau', \tau}^{N*}$, is exact on $\mathbb{R}[t]_N$,

$$\begin{aligned} \langle \underline{p}, D \delta \bar{q} \rangle_{\tau', \tau}^{N*} &= \int_{\tau'}^{\tau} \underline{p}(t) D_t \delta \bar{q}(t) dt, \\ &= \underline{p} \delta \bar{q}|_{\tau'}^{\tau} - \int_{\tau'}^{\tau} D_t \underline{p}(t) \delta \bar{q}(t) dt, \\ &= \underline{p} \delta \bar{q}|_{\tau'}^{\tau} - \langle D \underline{p}, \delta \bar{q} \rangle_{\tau', \tau}^{N*}. \end{aligned} \quad (17)$$

Hence,

$$\delta S_{\tau', \tau} = \langle \mathcal{L}_q + \underline{u} - D \mathcal{L}_v, \delta \bar{q} \rangle_{\tau', \tau}^{N*} + \underline{p} \delta \bar{q}|_{\tau'}^{\tau}. \quad (18)$$

For the constraint (14) to be satisfied, we require,

$$\begin{aligned} \delta S_{0, t_f} &:= \langle \mathcal{L}_q + \underline{u} - D \mathcal{L}_v, \delta \bar{q} \rangle^{N*} + \underline{p} \delta \bar{q}|_0^{t_f} = \underline{p} \delta \bar{q}|_0^{t_f}, \\ &\forall \delta \bar{q} \in (\mathbb{R}[t]_N)^{\dim(\mathcal{M})}, \\ &\Rightarrow (\mathcal{L}_q + \underline{u} - D \mathcal{L}_v) = \underline{0}. \end{aligned} \quad (19)$$

It follows hence from (18) that, $\delta S_{\tau', \tau} = \underline{p} \delta \bar{q}|_{\tau'}^{\tau}, \forall \tau', \tau$. This condition along with Lemma-1 proves that the discretization (14) is symplectic. \square

C. Algorithm

In the collocation ‘‘co-ordinates’’, let the metric tensor corresponding to the Legendre inner-product be given by G , the Lagrange derivative matrix on the grid by D , and the dual forms evaluating the polynomial at $t = 0, t_f$ by $L(0), L(t_f)$ respectively. Then the condition (14) becomes,

$$G(\mathcal{L}_q + \underline{u}) + (D^T G - [L^T(t_f)L(t_f) - L^T(0)L(0)])\mathcal{L}_v = 0, \quad (20)$$

where we use the fact that $D^\dagger := G^{-1}D^T G$.

Incorporating all the constraints from (3), the optimal control problem can now be approximated as the following finite-dimensional nonlinear optimization problem,

$$\begin{aligned} \min_{q, \underline{u}} & \mathbf{1}^T G \mathbf{l}, \\ G(\mathcal{L}_q + \underline{u}) + (D^T G - [L^T(t_f)L(t_f) - L^T(0)L(0)])\mathcal{L}_v &= 0, \\ L(0)q &= q_0, \\ \text{where } \mathbf{l}[i] &= l(q(t_i), (Dq)(t_i), u(t_i)). \end{aligned} \quad (21)$$

V. NUMERICAL EXAMPLES

We illustrate the performance of the algorithm described in this paper, using the classical Acrobot, and a 3-link analogue that we call the *3crobot*: a 3-link pendulum with free pivots on all but the last joint that has a torque actuator. The models in both-cases are non-dimensionalized using the mass & length of the first link, and time expressed in units such that $\mathbf{g} = 1.0$.

The common goal in both control problems is to start off from the lowest-energy state and swing up to the upright position, while minimizing $\int_0^{t_f} |\tau|^2 dt$, and optimizing $t_f \in [0, 10]$. The total duration for the optimization is bounded, but not exactly specified, thereby introducing a parameter to be optimized in addition to the control sequence.

The scheme presented in this paper is abbreviated by PMOC. The acronym DAE-EL refers to the pseudo-spectral discretization of the Euler-Lagrange condition : $D \mathcal{L}_v = \mathcal{L}_q + \underline{u}$. ODE-EL refers to the pseudo-spectral discretization of the resultant first-order ODE from Euler-Lagrange: $D \dot{q} = (\mathcal{L}_{vv})^{-1}(\mathcal{L}_q + \underline{u} - \mathcal{L}_{qv} \dot{q})$. The problems are discretized using Chebyshev polynomials, and the resulting nonlinear programs (21) are solved using SNOPT [16]¹.

¹Our implementation assumes all the constraints to be nonlinear; an assumption that impacts all the schemes considered here equally.

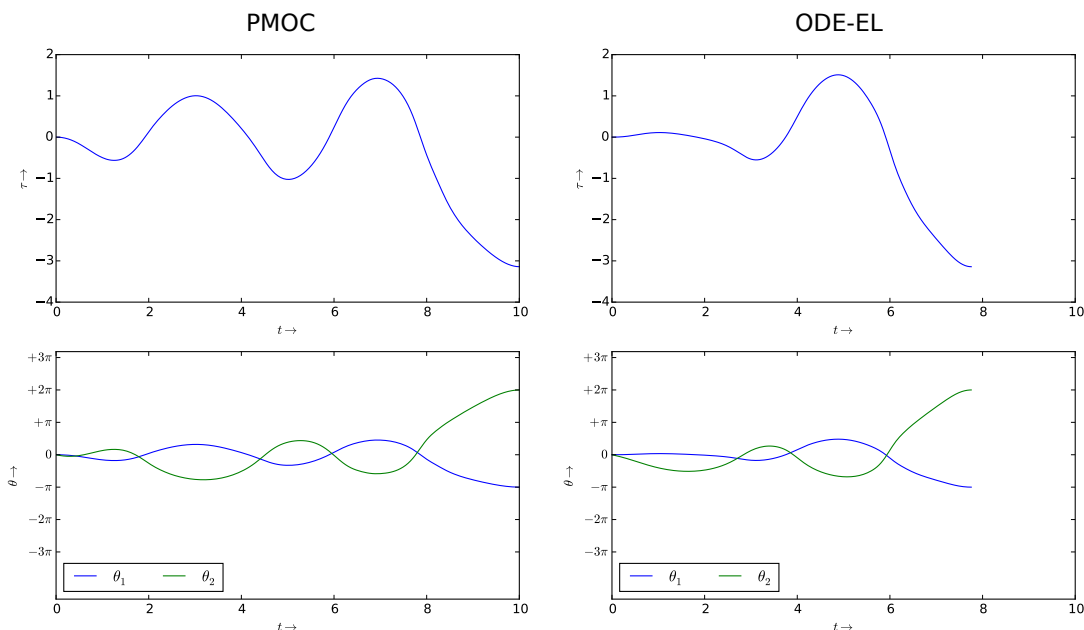


Fig. 1. Acrobot: Locally optimal solutions in $\mathbb{R}[t]_{64}$ found by (Left) PMOC & (Right) ODE-EL; (Top) Optimal control sequence; (Bottom) Corresponding trajectory found by the optimizer. The solution found by ODE-EL takes a shorter time to swing-up and also incurs a higher cost.

A. Acrobot

Taking inspiration from [17], we obtain the initial guess for the problem by using a sinusoidal waveform for $\tau(t)$. The number of major iterations taken by the SQP solver are listed in Table I.

We see that both PMOC and ODE-EL, converge to locally optimal solutions, but DAE-EL fails to find a feasible solution. PMOC converges with fewer major iterations than ODE-EL.

B. 3crobot

The schematic for the 3crobot is illustrated in Fig.2. We consider two versions of the control problem: one where the lengths of the links, l_2, l_3 are fixed at 0.5 each, and the other where it is required to find the optimal values for l_2, l_3 , such that $l_2 + l_3 = 1.0$. Such situations arise in coupled optimal control and design problems.

The initial (infeasible) guess was found by applying a constant torque around θ_1 , while the remaining joints were stabilized around 0 with a proportional controller. The number of major iterations taken by the SQP solver are listed in Table II.

In both cases, PMOC finds a local optimum faster than ODE-EL. ODE-EL had trouble satisfying the feasibility tolerance and exceeded the resource limit, but it appears to have landed in a neighborhood of the solutions found by PMOC. DAE-EL failed to converge in both cases.

VI. SUMMARY

Optimal control problems are generally non-convex, and underactuated problems of the kind presented in this paper




Fig. 2. 3crobot: The system is composed of 3 links of lengths 1.0, l_2, l_3 respectively, connected together by pivot joints as shown in the figure. An actuator situated on the joint farthest from the ground, can apply arbitrary torques τ on the joint. The control task is to start at rest from $\theta_{\{1,2,3\}} = 0$ and swing-up to $\theta_1 = \pi, \theta_{\{2,3\}} = 0$ (or equivalent co-ordinates) while minimizing the cost $\int_0^{t_f} |\tau|^2 dt, t_f \in [1, 10]$ s.

Algorithm	Major iterations	Cost
PMOC	218	0.63
DAE-EL	No feasible solution found	-
ODE-EL	688	0.80

TABLE I
PERFORMANCE ON THE ACROBOT PROBLEM

further accentuate the difficulty of finding the global or even a local minimum. Every smooth optimal control algorithm is susceptible to local optima, and this is partially addressed in practice using multiple starts. The challenge however is in designing algorithms that, more often than not, find




Fig. 3. 3-crobot: Locally optimal solutions in $\mathbb{R}[t]_{64}$ found by PMOC; (Left) Solution for fixed $l_2 = l_3 = 0.5$; (Right) Solution for link-lengths s.t $l_2 + l_3 = 1$: optimal parameter was found to be $(l_2^*, l_3^*) \approx (0.3, 0.7)$; (Top) Optimal control sequence; (Bottom) Corresponding trajectory found by the optimizer.

Algorithm	Major iterations	Cost
PMOC	498	0.61
DAE-EL	Singular basis	-
ODE-EL	> 1758	-
PMOC (l)	358	0.31
DAE-EL (l)	Singular basis	-
ODE-EL (l)	> 1758	-

TABLE II
PERFORMANCE ON THE 3CROBOT PROBLEM

feasible solutions. As seen from our numerical examples, even established methods sometimes fail to find feasible solutions.

We show that the algorithm that we propose enjoys computational advantages over the other candidates considered here; faster convergence and consistently finds feasible solutions. This comparison is by no means exhaustive, and only meant to be representative of the current state of the art in trajectory optimization. We also prove that the discretization used in our algorithm is both symplectic and momentum-conserving, and incorporates the beneficial aspects of pseudo-spectral methods. While our algorithm appears to also find “better” optima, we caution the reader against paying heed to this aspect of the results. With judicious use of multiple starts, it is possible that other algorithms may find comparable optima.

Performance of any of these methods is likely to be problem-dependent. For example, because there exist many different costs that generate the same optimal behavior of the system [18], the cost function itself could presumably

be tuned so as to favor the performance of one algorithm over the others. With the exception of such specially tuned cost functions, we find that PMOC is especially effective on complex mechanical systems with commonly used cost functions such as torque-squared, minimum-time and so on.

The use of variational integrators has been extensively pursued in the DMOC literature [7] [8] [10]. Unlike DMOC, the algorithm presented in this paper sacrifices causality in order to better approximate the dynamics using pseudo-spectral methods. While this leads to super-polynomial (as opposed to fixed-order algebraic) convergence rate requiring fewer grid points, it also leads to dense Jacobians and increased sensitivity to discontinuities. This trade-off is reminiscent of finite-element versus spectral methods, and suggests the future development of adaptive-variational schemes resembling hp-adaptive pseudo-spectral methods for optimal control [19].

The poor performance of DAE-EL in our examples underscores the lack of our present understanding on how geometric discretizations affect numerical methods for optimal control. It is known that ODE-stepping schemes based on the Gauss-Lobatto quadrature are both symplectic & momentum-conserving when working in the DAE-EL (Hamiltonian) form [9, p. 192]. Similar grid densities between Chebyshev and Legendre basis predisposed us to expect similar performance for DAE-EL and PMOC, but this clearly does not seem to be case for these examples (see Table I). This behaviour can partly be attributed to the ill-conditioning of the derivative matrix D_{ij} , a property which is ameliorated in PMOC by the use of the conjugate operator D^\dagger . The contrast in performance also seems to make the case for

symplectic discretizations, since DAE-EL like PMOC, is also momentum-conserving.

The results from this paper suggest that for a discretization of a given order, those that are geometry preserving offer advantages in terms of the rates and region of convergence for optimal control problems where finding feasible solutions is challenging. Future investigations will examine how the choice of the polynomial basis for pseudo-spectral interpolation affects the performance of the numerical method.

REFERENCES

- [1] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*. Siam, 2010, vol. 19.
- [2] L. S. Pontryagin, *Mathematical theory of optimal processes*. CRC Press, 1987.
- [3] F. Fahroo and I. M. Ross, "Direct trajectory optimization by a chebyshev pseudospectral method," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 1, pp. 160–166, 2002.
- [4] O. von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, vol. 37, no. 1, pp. 357–373, 1992.
- [5] D. A. Benson, G. T. Huntington, T. P. Thorvaldsen, and A. V. Rao, "Direct trajectory optimization and costate estimation via an orthogonal collocation method," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 6, pp. 1435–1440, 2006.
- [6] L. T. Biegler, "Solution of dynamic optimization problems by successive quadratic programming and orthogonal collocation," *Computers & chemical engineering*, vol. 8, no. 3, pp. 243–247, 1984.
- [7] S. Ober-blöbaum, O. Junge, and J. E. Marsden, "Discrete mechanics and optimal control: An analysis," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 17, pp. 322–352, 2011.
- [8] O. Junge, J. E. Marsden, and S. Ober-Blöbaum, "Discrete mechanics and optimal control," in *IFAC Congress, Praha*, 2005.
- [9] E. Hairer, C. Lubich, and G. Wanner, *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. Springer, 2006, vol. 31.
- [10] J. E. Marsden and M. West, "Discrete mechanics and variational integrators," *Acta Numerica*, vol. 10, no. 1, pp. 357–514, 2001.
- [11] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*. Dover, 2001.
- [12] G. H. Golub and J. H. Welsch, "Calculation of gauss quadrature rules," *Mathematics and Computation*, vol. 23, pp. 221–230, 1969.
- [13] S. Sachdeva and N. K. Vishnoi, "Faster algorithms via approximation theory," *Theoretical Computer Science*, vol. 9, no. 2, pp. 125–210, 2013.
- [14] L. N. Trefethen, "Is gauss quadrature better than clenshaw-curtis?" *SIAM review*, vol. 50, no. 1, pp. 67–87, 2008.
- [15] —, *Spectral methods in MATLAB*. Siam, 2000, vol. 10.
- [16] P. E. Gill, W. Murray, and M. A. Saunders, "Snopt: An sqp algorithm for large-scale constrained optimization," *SIAM review*, vol. 47, no. 1, pp. 99–131, 2005.
- [17] R. Tedrake, "Underactuated robotics: Learning, planning, and control for efficient and agile machines course notes for mit 6.832."
- [18] A. Y. Ng, S. J. Russell, *et al.*, "Algorithms for inverse reinforcement learning," in *Icml*, 2000, pp. 663–670.
- [19] C. L. Darby, W. W. Hager, and A. V. Rao, "An hp-adaptive pseudospectral method for solving optimal control problems," *Optimal Control Applications and Methods*, vol. 32, no. 4, pp. 476–502, 2011.

APPENDIX

A. Notation

We use $\langle \cdot, \cdot \rangle$ to denote the standard inner-product on L^2 . An N-point discretization via Gauss-Legendre quadrature is represented by $\langle \cdot, \cdot \rangle^N$. A '*' symbol in the super-script: $\langle \cdot, \cdot \rangle^{N*}$, is used to emphasize the grid-points at which the inner-product is computed; this being important for reasons of approximation. We denote by D^\dagger the conjugate linear operator under the inner product: $\langle f, Dg \rangle = \langle D^\dagger f, g \rangle$. The

matrix corresponding to a inner-product in some particular basis is termed the "Metric tensor", and denoted by G .

We employ a specialized notation for denoting vectors and duals over the state space of the system. A dual form is marked by a bar underneath the symbol: ' \underline{y} ', whilst a vector with one above it: ' \bar{x} '. The canonical pairing between a vector and a dual is denoted without special operators: $\underline{y}\bar{x} = \bar{x}y$. We also assume that operators on the function space act element-wise on a "stack" of elements - derivatives of a vector of polynomials, for instance. This also implies that the canonical pairing commutes with inner-products: $\langle \bar{f}, \underline{g} \rangle = \sum_i \langle \bar{f}_i, \underline{g}_i \rangle$. Partial derivatives of functions are assumed to be dual vectors (or "stacks" thereof) throughout the paper.

Generic vectors without special connotation or type will be represented in bold: ' \mathbf{x} '. Matrices (and metric tensors) are denoted in capitals 'B'. Generic vectors and matrices will appear in co-ordinate bound expressions, whilst the dual/tangent vectors. Note that unlike the mechanics literature, "co-ordinates" here refers to the polynomial base (or equivalently grid-points), and not the state space of the system.

B. Computing the Pullback

Let the collocation points, weights, and norm-squares (defined in (6), (9)) corresponding to the N-point Gauss-Lobatto quadrature be $\{t_i^l\}, \{w_i^l\}, \{\gamma_i^l\}$ respectively, and let those corresponding to the Orthogonal polynomial of interest \mathcal{B} , be $\{t_i^b\}, \{w_i^b\}, \{\gamma_i^b\}$. Given polynomials $x, y \in \mathbb{R}[t]_N$ let,

$$\mathbf{x}_i^b := x(t_i^b), \mathbf{x}_i^l := x(t_i^l), \mathbf{y}_i^b := y(t_i^b), \mathbf{y}_i^l := y(t_i^l),$$

$$B_{ij} := B_j(t_i^b), B_{ij}^l := B_j(t_i^l).$$

$$\begin{aligned} (x, y) &:= \int_{-1}^1 x(t)y(t)dt = \sum_i x(t_i^l)y(t_i^l)w_i^l \\ &= (\mathbf{x}^l)^T \text{diag}(\mathbf{w}^l)\mathbf{y}^l. \end{aligned} \quad (22)$$

Using the orthogonality of the basis polynomials (8), we note that,

$$\mathbf{p}_i^b := p(t_i^b), \mathbf{p}_i^l := p(t_i^l),$$

$$\mathbf{p}^l = B^l \text{diag}(1/\gamma^b)B^T \text{diag}(\mathbf{w}^b)\mathbf{p}^b, \quad (23)$$

Hence,

$$\begin{aligned} (p, q) &= \mathbf{p}^l \text{diag}(\mathbf{w}^l)\mathbf{q}^l = (\mathbf{p}^b)^T A^T \text{diag}(\mathbf{w}^l)A\mathbf{q}^b, \\ A &:= B^l \text{diag}(1/\gamma^b)B^T \text{diag}(\mathbf{w}^b). \end{aligned} \quad (24)$$

The pullback of the inner-product is therefore,

$$G^* = A^T \text{diag}(\mathbf{w}^l)A, \quad A := B^l \text{diag}(1/\gamma^b)B^T \text{diag}(\mathbf{w}^b). \quad (25)$$