



Geometric Methods for Adjoint Systems

Brian Kha Tran¹ · Melvin Leok¹

Received: 5 May 2022 / Accepted: 10 November 2023
© The Author(s) 2023

Abstract

Adjoint systems are widely used to inform control, optimization, and design in systems described by ordinary differential equations or differential-algebraic equations. In this paper, we explore the geometric properties and develop methods for such adjoint systems. In particular, we utilize symplectic and presymplectic geometry to investigate the properties of adjoint systems associated with ordinary differential equations and differential-algebraic equations, respectively. We show that the adjoint variational quadratic conservation laws, which are key to adjoint sensitivity analysis, arise from (pre)symplecticity of such adjoint systems. We discuss various additional geometric properties of adjoint systems, such as symmetries and variational characterizations. For adjoint systems associated with a differential-algebraic equation, we relate the index of the differential-algebraic equation to the presymplectic constraint algorithm of Gotay et al. (J Math Phys 19(11):2388–2399, 1978). As an application of this geometric framework, we discuss how the adjoint variational quadratic conservation laws can be used to compute sensitivities of terminal or running cost functions. Furthermore, we develop structure-preserving numerical methods for such systems using Galerkin Hamiltonian variational integrators (Leok and Zhang in IMA J. Numer. Anal. 31(4):1497–1532, 2011) which admit discrete analogues of these quadratic conservation laws. We additionally show that such methods are natural, in the sense that reduction, forming the adjoint system, and discretization all commute, for suitable choices of these processes. We utilize this naturality to derive a variational error analysis result for the presymplectic variational integrator that we use to discretize the adjoint DAE system. Finally, we discuss the application of adjoint systems in the context of optimal control problems, where we prove a similar naturality result.

Communicated by Anthony Bloch.

✉ Melvin Leok
mleok@ucsd.edu
Brian Kha Tran
b3tran@ucsd.edu

¹ Department of Mathematics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0112, USA

Contents

1	Introduction
1.1	Applications of the Adjoint Equations
1.2	Symplectic and Presymplectic Geometry
1.3	Main Contributions
1.4	Main Results
2	Adjoint Systems
2.1	Adjoint Equations on Vector Spaces
2.2	Adjoint Systems on Manifolds
2.2.1	Adjoint Systems with Augmented Hamiltonians
2.3	Adjoint Systems for DAEs via Presymplectic Mechanics
2.3.1	DAE Index and the Presymplectic Constraint Algorithm
2.3.2	Adjoint Systems for DAEs with Augmented Hamiltonians
3	Applications
3.1	Adjoint Sensitivity Analysis for Semi-Explicit Index 1 DAEs
3.2	Structure-Preserving Discretizations of Adjoint Systems
3.2.1	Naturality of the Adjoint DAE System Discretization
3.2.2	Numerical Example
3.3	Optimal Control of DAE Systems
4	Conclusion and Future Research Directions
	Appendix A: Proofs of Discrete Adjoint Variational Quadratic Conservation Laws
	Appendix B: Proof of Naturality of Adjoint System Discretization
	Appendix C: An Intrinsic Type II Variational Principle for Adjoint Systems
	References

1 Introduction

1.1 Applications of the Adjoint Equations

The solution of many nonlinear problems involves successive linearization, and as such, variational equations and their adjoints play a critical role in a variety of applications. Adjoint equations are of particular interest when the parameter space is significantly higher dimension than that of the output or objective. In particular, the simulation of adjoint equations arises in sensitivity analysis (Cacuci 1981; Cao et al. 2003), adaptive mesh refinement (Li and Petzold 2003), uncertainty quantification (Wang et al. 2012), automatic differentiation (Griewank 2003), superconvergent functional recovery (Pierce and Giles 2000), optimal control (Ross 2005), optimal design (Giles and Pierce 2000), optimal estimation (Nguyen et al. 2016), and deep learning viewed as an optimal control problem (Benning et al. 2019).

The study of geometric aspects of adjoint systems arose from the observation that the combination of any system of differential equations and its adjoint equations is described by a formal Lagrangian (Ibragimov 2006, 2007). This naturally leads to the question of when the formation of adjoints and discretization commutes (Sirkes and Tziperman 1997), and prior work on this includes the Ross–Fahroo lemma (Ross and Fahroo 2001), and the observation by (Sanz-Serna 2016) that the adjoints and discretization commute if and only if the discretization is symplectic.

1.2 Symplectic and Presymplectic Geometry

Throughout the paper, we will assume that all manifolds and maps are smooth, unless otherwise stated. Let (P, Ω) be a (finite-dimensional) symplectic manifold, i.e., Ω is a closed nondegenerate two-form on P . Given a Hamiltonian $H : P \rightarrow \mathbb{R}$, the Hamiltonian system is defined by

$$i_{X_H} \Omega = dH,$$

where the vector field X_H is a section of the tangent bundle to P . By nondegeneracy, the vector field X_H exists and is uniquely determined. For an open interval $I \subset \mathbb{R}$, we say that a curve $z : I \rightarrow P$ is a solution of Hamilton's equations if z is an integral curve of X_H , i.e., $\dot{z}(t) = X_H(z(t))$ for all $t \in I$.

A particularly important example for our purposes is when the symplectic manifold is the cotangent bundle of a manifold, $P = T^*M$, equipped with the canonical symplectic form $\Omega = dq \wedge dp$ in natural coordinates (q, p) on T^*M . A Hamiltonian system has the coordinate expression

$$\begin{aligned} \dot{q} &= \frac{\partial H(q, p)}{\partial p}, \\ \dot{p} &= -\frac{\partial H(q, p)}{\partial q}. \end{aligned}$$

By Darboux's theorem, any symplectic manifold is locally symplectomorphic to a cotangent bundle equipped with its canonical symplectic form. As such, any Hamiltonian system can be locally expressed in the above form (even when P is not a cotangent bundle), using Darboux coordinates.

We now consider the generalization of Hamiltonian systems where we relax the condition that Ω is nondegenerate, i.e., presymplectic geometry. Let (P, Ω) be a presymplectic manifold, i.e., Ω is a closed two-form on P with constant rank. As before, given a Hamiltonian $H : P \rightarrow \mathbb{R}$, we define the associated Hamiltonian system as

$$i_{X_H} \Omega = dH.$$

Note that since Ω is now degenerate, X_H is not guaranteed to exist and if it does, it need not be unique and in general is only partially defined on a submanifold of P . Again, we say a curve on P is a solution to Hamilton's equations if it is an integral curve of X_H . Using Darboux coordinates (q, p, r) adapted to (P, Ω) , where $\Omega = dq \wedge dp$ and $\ker(\Omega) = \text{span}\{\partial/\partial r\}$, the local expression for Hamilton's equations is given by

$$\begin{aligned} \dot{q} &= \frac{\partial H(q, p, r)}{\partial p}, \\ \dot{p} &= -\frac{\partial H(q, p, r)}{\partial q}, \end{aligned}$$

$$0 = \frac{\partial H(q, p, r)}{\partial r}.$$

The third equation above is interpreted as a constraint equation which any solution curve must satisfy. We will assume that the constraint defines a submanifold of P . It is clear that in order for a solution vector field X_H to exist, it must be restricted to lie on this submanifold. However, in order for its flow to remain on the submanifold, it must be tangent to this submanifold, which further restricts where X can be defined. Alternating restriction in order to satisfy these two constraints yields the presymplectic constraint algorithm of Gotay et al. (1978). The presymplectic constraint algorithm begins with the observation that for any X satisfying the above system, so does $X + Z$, where $Z \in \ker(\Omega)$. In order to obtain such a vector field X , one considers the subset P_1 of P such that $Z_p(H) = 0$ for any $Z \in \ker(\Omega)$, $p \in P_1$. We will assume that the set P_1 is a submanifold of P . We refer to P_1 as the primary constraint manifold. In order for the flow of the resulting Hamiltonian vector field X to remain on P_1 , one further requires that X is tangent to P_1 . The set of points satisfying this property defines a subsequent secondary constraint submanifold P_2 . Iterating this process, one obtains a sequence of submanifolds

$$\cdots \rightarrow P_k \rightarrow \cdots \rightarrow P_1 \rightarrow P_0 \equiv P,$$

defined by

$$P_{k+1} = \{p \in P_k : Z_p(H_k) = 0 \text{ for all } Z \in \ker(\Omega_k)\}, \quad (1.1)$$

where

$$\begin{aligned} \Omega_{k+1} &= \Omega_k|_{P_{k+1}}, \\ H_{k+1} &= H_k|_{P_{k+1}}. \end{aligned}$$

If there exists a nontrivial fixed point in this sequence, i.e., a submanifold P_k of P such that $P_k = P_{k+1}$, we refer to P_k as the final constraint manifold. If such a fixed point exists, we denote by ν_P the minimum integer such that $P_{\nu_P} = P_{\nu_P+1}$, i.e., ν_P is the number of steps necessary for the presymplectic constraint algorithm to terminate. If such a final constraint manifold P_{ν_P} exists, there always exists a solution vector field X defined on and tangent to P_{ν_P} such that $i_X \Omega_{\nu_P} = dH_{\nu_P}$ and X is unique up to the kernel of Ω_{ν_P} . Furthermore, such a final constraint manifold is maximal in the sense that if there exists a submanifold N of P which admits a vector field X defined on and tangent to N such that $i_X \Omega|_N = dH|_N$, then $N \subset P_{\nu_P}$ (Gotay and Nester 1979).

1.3 Main Contributions

In this paper, we explore the geometric properties of adjoint systems associated with ordinary differential equations (ODEs) and differential-algebraic equations (DAEs). For a discussion of adjoint systems associated with ODEs and DAEs, see Sanz-Serna

(2016) and Cao et al. (2003), respectively. In particular, we utilize the machinery of symplectic and presymplectic geometry as a basis for understanding such systems.

In Sect. 2.1, we review the notion of adjoint equations associated with ODEs over vector spaces. We show that the quadratic conservation law, which is the key to adjoint sensitivity analysis, arises from the symplecticity of the flow of the adjoint system. In Sect. 2.2, we investigate the symplectic geometry of adjoint systems associated with ODEs on manifolds. We additionally discuss augmented adjoint systems, which are useful in the adjoint sensitivity of running cost functions. In Sect. 2.3, we investigate the presymplectic geometry of adjoint systems associated with DAEs on manifolds. We investigate the relation between the index of the base DAE and the index of the associated adjoint system, using the notions of DAE reduction and the presymplectic constraint algorithm. We additionally consider augmented systems for such adjoint DAE systems. For the various adjoint systems that we consider, we derive various quadratic conservation laws which are useful in adjoint sensitivity analysis of terminal and running cost functions. We additionally discuss symmetry properties and present variational characterizations of such systems that provide a useful perspective for constructing geometric numerical methods for these systems.

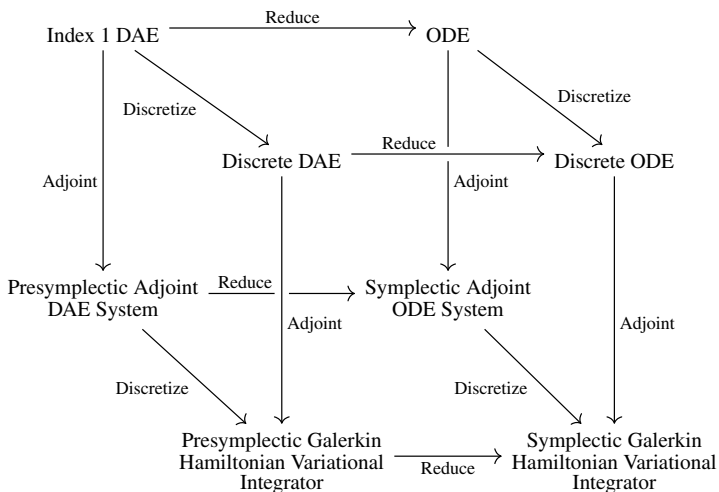
In Sect. 3, we discuss applications of the various adjoint systems to adjoint sensitivity and optimal control. In Sect. 3.1, we show how the quadratic conservation laws developed in Sect. 2 can be used for adjoint sensitivity analysis of running and terminal cost functions, subject to ODE or DAE constraints. In Sect. 3.2, we construct structure-preserving discretizations of adjoint systems using the Galerkin Hamiltonian variational integrator construction of Leok and Zhang (2011). For adjoint DAE systems, we introduce a presymplectic analogue of the Galerkin Hamiltonian variational integrator construction. We show that such discretizations admit discrete analogues of the aforementioned quadratic conservation laws and hence are suitable for the numerical computation of adjoint sensitivities. Furthermore, we show that such discretizations are natural when applied to DAE systems, in the sense that reduction, forming the adjoint system, and discretization all commute (for particular choices of these processes). As an application of this naturality, we derive a variational error analysis result for the resulting presymplectic variational integrator for adjoint DAE systems. Finally, in Sect. 3.3, we discuss adjoint systems in the context of optimal control problems, where we prove a similar naturality result, in that suitable choices of reduction, extremization, and discretization commute.

By developing a geometric theory for adjoint systems, the application areas that utilize such adjoint systems can benefit from the existing work on geometric and structure-preserving methods.

1.4 Main Results

In this paper, we prove that, starting with an index 1 DAE, appropriate choices of reduction, discretization, and forming the adjoint system commute. That is, the following

diagram commutes.



In order to prove this result, we develop along the way the definitions of the various vertices and arrows in the above diagram. Roughly speaking, the four “Adjoint” arrows are defined by forming the appropriate continuous or discrete action and enforcing the variational principle; the four “Reduce” arrows are defined by solving the algebraic variables in terms of the kinematic variables through the continuous or discrete constraint equations; the two “Discretize” arrows on the top face are given by a Runge–Kutta method, while the two “Discretize” arrows on the bottom face are given by the associated symplectic partitioned Runge–Kutta method. The above commutative diagram can be understood as an extension of the result of Sanz-Serna (2016) (that discretization and forming the adjoint of an ODE commute when the discretization is a symplectic Runge–Kutta method) by adding the reduction operation. In order to appropriately define this reduction operation, we will show that the presymplectic adjoint DAE system has index 1 if the base DAE has index 1, so that the reduction of the presymplectic adjoint DAE system results in a symplectic adjoint ODE system; the tool for this will be the presymplectic constraint algorithm.

In the process of defining the ingredients in the above diagram, we will additionally prove various properties of adjoint systems associated with ODEs and DAEs. The key properties that we will prove for such adjoint systems are the adjoint variational quadratic conservation laws, Propositions 2.3, 2.7, 2.11, 2.12. As we will show, these conservation laws can be used to compute adjoint sensitivities of running and terminal cost functions under the flow of an ODE or DAE. In order to prove these conservation laws, we will need to define the variational equations associated with an adjoint system. We will define them as the linearization of the base ODE or DAE; for the DAE case, we will show that the variational equations have the same index as the base DAE so that they have the same (local) solvability.

2 Adjoint Systems

2.1 Adjoint Equations on Vector Spaces

In this section, we review the notion of adjoint equations on vector spaces and their properties, as preparation for adjoint systems on manifolds.

Let Q be a finite-dimensional vector space and consider the ordinary differential equation on Q given by

$$\dot{q} = f(q), \tag{2.1}$$

where $f : Q \rightarrow Q$ is a differentiable vector field on Q . Let $Df(q)$ denote the linearization of f at $q \in Q$, $Df(q) \in L(Q, Q)$. Denoting its adjoint by $[Df(q)]^* \in L(Q^*, Q^*)$, the adjoint equation associated with (2.1) is given by

$$\dot{p} = -[Df(q)]^* p, \tag{2.2}$$

where p is a curve on Q^* .

Let q^A be coordinates for Q and let p_A be the associated dual coordinates for Q^* , so that the duality pairing is given by $\langle p, q \rangle = p_A q^A$. The linearization of f at q is given in coordinates by

$$(Df(q))_B^A = \frac{\partial f^A(q)}{\partial q^B},$$

where its action on $v \in Q$ in coordinates is

$$(Df(q)v)^A = \frac{\partial f^A(q)}{\partial q^B} v^B.$$

Its adjoint then acts on $p \in Q^*$ by

$$([Df(q)]^* p)_A = \frac{\partial f^B(q)}{\partial q^A} p_B.$$

Thus, the ODE and its adjoint can be expressed in coordinates as

$$\begin{aligned} \dot{q}^A &= f^A(q), \\ \dot{p}_A &= -\frac{\partial f^B(q)}{\partial q^A} p_B. \end{aligned}$$

Next, we recall that the combined system (2.1)–(2.2), which we refer to as the adjoint system, arises from a variational principle. Letting $\langle \cdot, \cdot \rangle$ denote the duality pairing between Q^* and Q , we define the Hamiltonian

$$H : Q \times Q^* \rightarrow \mathbb{R},$$

$$(q, p) \mapsto H(q, p) \equiv \langle p, f(q) \rangle.$$

The associated action, defined on the space of curves on $Q \times Q^*$ covering some interval (t_0, t_1) , is given by

$$S[q, p] = \int_{t_0}^{t_1} (\langle p, \dot{q} \rangle - H(q, p)) dt = \int_{t_0}^{t_1} (\langle p, \dot{q} \rangle - \langle p, f(q) \rangle) dt.$$

Proposition 2.1 *The variational principle $\delta S = 0$, subject to variations $(\delta q, \delta p)$ which fix the endpoints $\delta q(t_0) = 0, \delta q(t_1) = 0$, yields the adjoint system (2.1)–(2.2).*

Proof Compute the variation of S with respect to a compactly supported variation $(\delta q, \delta p)$,

$$\begin{aligned} \delta S[q, p] \cdot (\delta q, \delta p) &= \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} S[q + \epsilon \delta q, p + \epsilon \delta p] \\ &= \int_{t_0}^{t_1} \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} \langle p + \epsilon \delta p, \dot{q} + \epsilon \delta \dot{q} - f(q + \epsilon \delta q) \rangle dt \\ &= \int_{t_0}^{t_1} (\langle \delta p, \dot{q} - f(q) \rangle + \langle p, \delta \dot{q} - Df(q) \delta q \rangle) dt \\ &= \int_{t_0}^{t_1} (\langle \delta p, \dot{q} - f(q) \rangle + \langle -\dot{p} - [Df(q)]^* p, \delta q \rangle) dt. \end{aligned}$$

The fundamental lemma of the calculus of variations then yields (2.1)–(2.2). \square

Remark 2.1 Note that an analogous statement of Proposition 2.1 can also be stated using the Type II variational principle, where one instead considers the generating function

$$H_+(q_0, p_1) = \text{ext} \left[p(t_1)q(t_1) - \int_{t_0}^{t_1} \langle p, \dot{q} - H(q, p) \rangle dt \right],$$

and one extremizes over C^2 curves from $[t_0, t_1]$ to T^*Q such that $q(t_0) = q_0, p(t_1) = p_1$. The Type II variational principle again gives the above adjoint system, but with differing boundary conditions. These boundary conditions are typical in adjoint sensitivity analysis, where one fixes the initial position and the final momenta.

The variational principle utilized above is formulated so that the stationarity condition $\delta S = 0$ is equivalent to Hamilton's equations, where we view $Q \times Q^* \cong T^*Q$ with the canonical symplectic form on the cotangent bundle $\Omega = dq \wedge dp$ and with the corresponding Hamiltonian $H : T^*Q \rightarrow \mathbb{R}$ given as above. It then follows that the flow of the adjoint system is symplectic.

The symplecticity of the adjoint system is a key feature of the system. In fact, the symplecticity of the adjoint system implies that a certain quadratic invariant is preserved along the flow of the system. This quadratic invariant is the key ingredient

to the use of adjoint equations for sensitivity analysis. To state the quadratic invariant, consider the variational equation associated with equation (2.1),

$$\frac{d}{dt}\delta q = Df(q)\delta q, \tag{2.3}$$

which corresponds to the linearization of (2.1) at $q \in Q$. For solution curves p and δq to (2.2) and (2.3), respectively, over the same curve q , one has that the quantity $\langle p, \delta q \rangle$ is preserved along the flow of the system, since

$$\begin{aligned} \frac{d}{dt}\langle p, \delta q \rangle &= \langle \dot{p}, \delta q \rangle + \left\langle p, \frac{d}{dt}\delta q \right\rangle = \langle -[Df(q)]^*p, \delta q \rangle + \langle p, Df(q)\delta q \rangle \\ &= -\langle p, Df(q)\delta q \rangle + \langle p, Df(q)\delta q \rangle = 0. \end{aligned}$$

To see that symplecticity implies the preservation of this quantity, recall that symplecticity is the statement that, along a solution curve of the adjoint system (2.1)–(2.2), one has

$$\frac{d}{dt}\Omega(V, W) = 0,$$

where V and W are first variations to the adjoint system (i.e., that the flow of V and W on solutions are again solutions). Infinitesimally, first variations V and W correspond to solutions of the linearization of the adjoint system (2.1)–(2.2). At a solution (q, p) to the adjoint system, the linearization of the system is given by

$$\begin{aligned} \frac{d}{dt}\delta q &= Df(q)\delta q, \\ \frac{d}{dt}\delta p &= -[Df(q)]^*\delta p. \end{aligned}$$

Note that the first equation is just the variational equation (2.3), while the second equation is the adjoint equation (2.2), with p replaced by δp , since the adjoint equation is linear in p . The first variation vector field V corresponding to a solution $(\delta q, \delta p)$ of this linearized system is

$$V = \delta q \frac{\partial}{\partial q} + \delta p \frac{\partial}{\partial p}.$$

Now, we make two choices for the first variations V and W . For W , we take the solution $\delta q = 0, \delta p = p$ of the linearized system, which gives $W = p \partial/\partial p$. For V , we take the solution $\delta q = \delta q, \delta p = 0$ of the linearized system, which gives $V = \delta q \partial/\partial q$. Inserting these into Ω gives

$$\Omega(V, W) = p \frac{\partial}{\partial p} \lrcorner \left(\delta q \frac{\partial}{\partial q} \lrcorner (dq \wedge dp) \right) = \langle p, \delta q \rangle.$$

Thus, symplecticity $\frac{d}{dt}\Omega(V, W) = 0$ with this particular choice of first variations V, W gives the preservation of the quadratic invariant $\langle p, \delta q \rangle$.

2.2 Adjoint Systems on Manifolds

We now extend the notion of the adjoint system to the case where the configuration space of the base ODE is a manifold. We will provide a symplectic characterization of these adjoint systems, prove the associated adjoint variational quadratic conservation laws, and additionally discuss symmetries and variational principles associated with these systems.

Let M be a manifold and consider the ODE on M given by

$$\dot{q} = f(q), \quad (2.4)$$

where f is a vector field on M . Letting $\pi : TM \rightarrow M$ denote the tangent bundle projection, we recall that a vector field f is a map $f : M \rightarrow TM$ which satisfies $\pi \circ f = \mathbf{1}_M$, i.e., f is a section of the tangent bundle.

Analogous to the adjoint system on vector spaces, we will define the adjoint system on a manifold as an ODE on the cotangent bundle T^*M which covers (2.4), such that the time evolution of the momenta in the fibers of T^*M are given by an adjoint linearization of f .

To do this, in analogy with the vector space case, consider the Hamiltonian $H : T^*M \rightarrow \mathbb{R}$ given by $H(q, p) = \langle p, f(q) \rangle_q$ where $\langle \cdot, \cdot \rangle_q$ is the duality pairing of T_q^*M with T_qM . When there is no possibility for confusion of the base point, we simply denote this duality pairing as $\langle \cdot, \cdot \rangle$. Recall that the cotangent bundle T^*M possesses a canonical symplectic form $\Omega = -d\Theta$ where Θ is the tautological one-form on T^*M . With coordinates $(q, p) = (q^A, p_A)$ on T^*M , this symplectic form has the coordinate expression $\Omega = dq \wedge dp \equiv dq^A \wedge dp_A$.

We define the adjoint system as the ODE on T^*M given by Hamilton's equations, with the above choice of Hamiltonian H and the canonical symplectic form. Thus, the adjoint system is given by the equation

$$i_{X_H}\Omega = dH,$$

whose solution curves on T^*M are the integral curves of the Hamiltonian vector field X_H . As is well-known, for the particular choice of Hamiltonian $H(q, p) = \langle p, f(q) \rangle$, the Hamiltonian vector field X_H is given by the cotangent lift \widehat{f} of f , which is a vector field on T^*M that covers f (for a discussion of the geometry of the cotangent bundle and lifts, see Yano and Ishihara 1973; for a discussion of cotangent lifts in the context of optimal control, see Bullo and Lewis 2014). With coordinates $z = (q, p)$ on T^*M , the adjoint system is the ODE on T^*M given by

$$\dot{z} = \widehat{f}(z). \quad (2.5)$$

To be more explicit, recall that the cotangent lift of f is constructed as follows. Let $\Phi_\epsilon : M \rightarrow M$ denote the one-parameter family of diffeomorphisms generated

by f . Then, we consider the cotangent lifted diffeomorphisms given by $(\Phi_{-\epsilon})^* : T^*M \rightarrow T^*M$. This covers Φ_ϵ in the sense that $\pi_{T^*M} \circ (\Phi_{-\epsilon})^* = \Phi_\epsilon \circ \pi_{T^*M}$ where $\pi_{T^*M} : T^*M \rightarrow M$ is the cotangent projection. The cotangent lift \widehat{f} is then defined to be the infinitesimal generator of the cotangent lifted flow,

$$\widehat{f}(z) = \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} (\Phi_{-\epsilon})^*(z).$$

We can directly verify that \widehat{f} is the Hamiltonian vector field for H , which follows from

$$i_{\widehat{f}}\Omega = -i_{\widehat{f}}d\Theta = -\mathcal{L}_{\widehat{f}}\Theta + d(i_{\widehat{f}}\Theta) = d(i_{\widehat{f}}\Theta) = dH,$$

where $\mathcal{L}_{\widehat{f}}\Theta = 0$ follows from the fact that cotangent lifted flows preserve the tautological one-form and $H = i_{\widehat{f}}\Theta$ follows from a direct computation (where $i_{\widehat{f}}\Theta$ is interpreted as a function on the cotangent bundle which maps (q, p) to $\langle \Theta(q, p), \widehat{f}(q, p) \rangle$)

The adjoint system (2.5) covers (2.4) in the following sense.

Proposition 2.2 *Integral curves to the adjoint system (2.5) lift integral curves to the system (2.4).*

Proof Let $z = (q, p)$ be coordinates on T^*M . Let $(\dot{q}, \dot{p}) \in T_{(q,p)}T^*M$. Then, $T\pi_{T^*M}(\dot{q}, \dot{p}) = \dot{q}$ where $T\pi_{T^*M}$ is the pushforward of the cotangent projection. Furthermore,

$$\begin{aligned} T\pi_{T^*M}\widehat{f}(q, p) &= T\pi_{T^*M}\left. \frac{d}{d\epsilon} \right|_{\epsilon=0} (\Phi_{-\epsilon})^*(q, p) = \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} (\pi_{T^*M} \circ (\Phi_{-\epsilon})^*)(q, p) \\ &= \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} (\Phi_\epsilon \circ \pi_{T^*M})(q, p) = \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} \Phi_\epsilon(q) = f(q). \end{aligned}$$

Thus, the pushforward of the cotangent projection applied to (2.5) gives (2.4). It then follows that integral curves of (2.5) lift integral curves of (2.4). \square

Remark 2.2 This can also be seen explicitly in coordinates. Recalling that $i_{\widehat{f}}\Omega = dH$, one has

$$dH = d(p_A f^A(q)) = f^A(q)dp_A + p_B \frac{\partial f^B(q)}{\partial q^A} dq^A,$$

and, on the other hand, denoting $\widehat{f}(q, p) = X^A(q, p)\partial/\partial q^A + Y_A(q, p)\partial/\partial p_A$,

$$i_{\widehat{f}}\Omega = (X^A(q, p)\partial_{q^A} + Y_A(q, p)\partial_{p_A}) \lrcorner (dq^B \wedge dp_B) = X^A(q, p)dp_A - Y_A(q, p)dq^A.$$

Equating these two gives the coordinate expression for the cotangent lift \widehat{f} ,

$$\widehat{f}(q, p) = f^A(q) \frac{\partial}{\partial q^A} - p_B \frac{\partial f^B(q)}{\partial q^A} \frac{\partial}{\partial p_A}.$$

Thus, the system $\dot{z} = \widehat{f}(z)$ can be expressed in coordinates as

$$\dot{q}^A = f^A(q), \tag{2.6a}$$

$$\dot{p}_A = -p_B \frac{\partial f^B(q)}{\partial q^A}, \tag{2.6b}$$

which clearly covers the original ODE $\dot{q}^A = f^A(q)$. Also, note that this coordinate expression for the adjoint system recovers the coordinate expression for the adjoint system in the vector space case.

Analogous to the vector space case, the adjoint system possesses a quadratic invariant associated with the variational equations of (2.4). The variational equation is given by considering the tangent lifted vector field on TM , $\tilde{f} : TM \rightarrow TTM$, which is defined in terms of the flow Φ_ϵ generated by f by

$$\tilde{f}(q, \delta q) = \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} T\Phi_\epsilon(q, \delta q),$$

where $(q, \delta q)$ are coordinates on TM . That is, \tilde{f} is the infinitesimal generator of the tangent lifted flow. The variational equation associated with (2.4) is the ODE associated with the tangent lifted vector field. In coordinates,

$$\frac{d}{dt}(q, \delta q) = \tilde{f}(q, \delta q). \tag{2.7}$$

Proposition 2.3 For integral curves (q, p) of (2.5) and $(q, \delta q)$ of (2.7), which cover the same curve q ,

$$\frac{d}{dt} \left\langle (q(t), p(t)), (q(t), \delta q(t)) \right\rangle_{q(t)} = 0. \tag{2.8}$$

Proof Note that $(q(t), p(t)) \in T_{q(t)}^*M$ and $(q(t), \delta q(t)) \in T_{q(t)}M$, so the duality pairing is well-defined. Then,

$$\begin{aligned} \left\langle (q(t), p(t)), (q(t), \delta q(t)) \right\rangle_{q(t)} &= \left\langle (\Phi_{-t})^*(q(0), p(0)), T\Phi_t(q(0), \delta q(0)) \right\rangle_{q(t)} \\ &= \left\langle (q(0), p(0)), T\Phi_{-t} \circ T\Phi_t(q(0), \delta q(0)) \right\rangle_{q(0)} \\ &= \left\langle (q(0), p(0)), T(\Phi_{-t} \circ \Phi_t)(q(0), \delta q(0)) \right\rangle_{q(0)} \\ &= \left\langle (q(0), p(0)), (q(0), \delta q(0)) \right\rangle_{q(0)}, \end{aligned}$$

so the pairing is constant. □

Remark 2.3 In the vector space case, we saw that the preservation of the quadratic invariant is implied by symplecticity. The above result is analogously implied by

symplecticity, noting that the flow of the adjoint system is symplectic since \widehat{f} is a Hamiltonian vector field.

Another conserved quantity for the adjoint system (2.5) is the Hamiltonian, since the adjoint system corresponds to a time-independent Hamiltonian flow, $\frac{d}{dt}H = \Omega(X_H, X_H) = 0$.

Additionally, conserved quantities for adjoint systems are generated, via cotangent lift, by symmetries of the original ODE (2.4), where we say that a vector field g is a symmetry of the ODE $\dot{x} = h(x)$ if $[g, h] = 0$.

Proposition 2.4 *Let g be a symmetry of (2.4), i.e., $[g, f] = 0$. Then, its cotangent lift \widehat{g} is a symmetry of (2.5), and additionally, the function*

$$\langle \Theta, \widehat{g} \rangle$$

*on T^*M is preserved along the flow of \widehat{f} , i.e., under the flow of the adjoint system (2.5).*

Proof We first show that \widehat{g} is a symmetry of (2.5), i.e., that $[\widehat{g}, \widehat{f}] = 0$. To see this, we recall that the cotangent lift of the Lie bracket of two vector fields equals the Lie bracket of their cotangent lifts,

$$[\widehat{g}, \widehat{f}] = \widehat{[g, f]}.$$

Then, since $[g, f] = 0$ by assumption, $[\widehat{g}, \widehat{f}] = \widehat{[g, f]} = \widehat{0} = 0$.

To see that $\langle \Theta, \widehat{g} \rangle$ is preserved along the flow of \widehat{f} , we have

$$\mathcal{L}_{\widehat{f}}\langle \Theta, \widehat{g} \rangle = \langle \mathcal{L}_{\widehat{f}}\Theta, \widehat{g} \rangle + \langle \Theta, \mathcal{L}_{\widehat{f}}\widehat{g} \rangle = \langle 0, \widehat{g} \rangle + \langle \Theta, [\widehat{f}, \widehat{g}] \rangle = 0,$$

where we used that $\mathcal{L}_{\widehat{f}}\Theta = 0$ since \widehat{f} is a cotangent lifted vector field. □

Remark 2.4 The above proposition states when $[f, g] = 0$, the Hamiltonian for the adjoint system associated with g , $\langle \Theta, \widehat{g} \rangle$, is preserved along the Hamiltonian flow corresponding to the Hamiltonian for the adjoint system associated with f , $\langle \Theta, \widehat{f} \rangle$, and vice versa. Note $\langle \Theta, \widehat{g} \rangle$ can be interpreted as the momentum map corresponding to the action on T^*M given by the flow of \widehat{g} .

The above proposition shows that (at least some) symmetries of the adjoint system (2.5) can be found by cotangent lifting symmetries of the original ODE (2.4). Additionally, the above proposition states that such cotangent lifted symmetries give rise to conserved quantities.

In light of the above proposition, it is natural to ask the following question. Given a symmetry G of the adjoint system (2.5) (i.e., $[G, \widehat{f}] = 0$), does it arise from a cotangent lifted symmetry in the sense of Proposition 2.4? In general, the answer is no. However, for a projectable vector field G which is a symmetry of the adjoint system, its projection by $T\pi_{T^*M}$ to a vector field on M does satisfy the assumptions of Proposition 2.4. This gives the following partial converse to the above proposition.

Proposition 2.5 *Let G be a projectable vector field on the bundle $\pi_{T^*M} : T^*M \rightarrow M$ which is a symmetry of (2.5), i.e., $[G, \widehat{f}] = 0$. Then, the pushforward vector field $g = T\pi_{T^*M}(G)$ on M satisfies the assumptions of Proposition 2.4 and $T\pi_{T^*M}\widehat{g} = T\pi_{T^*M}G$.*

Proof Since G is a projectable vector field on the cotangent bundle, $g = T\pi_{T^*M}G$ defines a well-defined vector field on M . Thus,

$$[g, f] = [T\pi_{T^*M}G, T\pi_{T^*M}\widehat{f}] = T\pi_{T^*M}[G, \widehat{f}] = T\pi_{T^*M}0 = 0,$$

so g is a symmetry of (2.4). Furthermore, we also have

$$T\pi_{T^*M}\widehat{g} = T\pi_{T^*M}(T\widehat{\pi_{T^*M}G}) = T\pi_{T^*M}G.$$

□

The preceding proposition shows that, for the class of projectable symmetries of the adjoint system (2.5), it is always possible to find an associated symmetry of the original ODE (2.4) which, by Proposition 2.4, corresponds to a Hamiltonian symmetry. Note that this implies that we can associate a conserved quantity $\langle \Theta, \widehat{g} \rangle$ to G , where $g = T\pi_{T^*M}G$. Furthermore, since $T\pi_{T^*M}\widehat{g} = T\pi_{T^*M}G$ and the canonical form Θ is a horizontal one-form, this implies that $\langle \Theta, G \rangle$ equals $\langle \Theta, \widehat{g} \rangle$ and hence, is conserved.

These two propositions show that symmetries of an ODE can be identified with equivalence classes of projectable symmetries of the associated adjoint system, where two projectable symmetries are equivalent if their difference lies in the kernel of $T\pi_{T^*M}$.

We also recall that the adjoint system (2.5) formally arises from a variational principle. To do so, let Θ be the tautological one-form on T^*M . The action is defined to be

$$S[\psi] = \int_I [\psi^*\Theta - (H \circ \psi)dt], \tag{2.9}$$

where $\psi(t) = (q(t), p(t))$ is a curve on T^*M over the interval $I = (t_0, t_1)$. We consider the variational principle $\delta S[\psi] = 0$, subject to variations which fix the endpoints $q(t_0), q(t_1)$.

Proposition 2.6 *Let ψ be a curve on T^*M over the interval I . Then, ψ is a stationary point of S with respect to variations which fix $q(t_0), q(t_1)$ if and only if (2.5) holds.*

The proof of the above proposition is standard in the literature, so we will omit it.

Remark 2.5 It should be noted that although the fixed endpoint conditions on $q(t_0)$ and $q(t_1)$ in the variational principle formally obtains the correct equations of motion for the adjoint system, these boundary conditions are incompatible with the adjoint system, since it covers an ODE on the base manifold. From a theoretical perspective, this is not an obstruction to Proposition 2.6 since the equations of motion are obtained after

enforcing the variational principle. However, from a numerical perspective, a variational principle with Type II boundary conditions fixing $q(t_0)$ and $p(t_1)$ is preferable for constructing variational integrators for adjoint systems. In Appendix C, we develop an intrinsic Type II variational principle to incorporate these boundary conditions.

Remark 2.6 In coordinates, the above action (2.9) takes the form

$$S = \int_{t_0}^{t_1} (\langle p, \dot{q} \rangle - \langle p, f(q) \rangle) dt,$$

which is the same coordinate expression as the action in the vector space case.

2.2.1 Adjoint Systems with Augmented Hamiltonians

In this section, we consider a class of modified adjoint systems, where some function on the base manifold M is added to the Hamiltonian of the adjoint system. More precisely, let $H : T^*M \rightarrow \mathbb{R}$, $H(q, p) = \langle p, f(q) \rangle$ be the Hamiltonian of the previous section, corresponding to the ODE $\dot{q} = f(q)$. Let $L : M \rightarrow \mathbb{R}$ be a function on M . We identify L with its pullback through $\pi_{T^*M} : T^*M \rightarrow M$. Then, we define the augmented Hamiltonian

$$H_L \equiv H + L : T^*M \rightarrow \mathbb{R}$$

$$(q, p) \mapsto H(q, p) + L(q) = \langle p, f(q) \rangle + L(q).$$

We define the augmented adjoint system as the Hamiltonian system associated with H_L relative to the canonical symplectic form Ω on T^*M ,

$$i_{X_{H_L}} \Omega = dH_L. \tag{2.10}$$

Remark 2.7 The motivation for such systems arises from adjoint sensitivity analysis and optimal control. For adjoint sensitivity analysis of a running cost function, one is concerned with the sensitivity of some functional

$$\int_0^t L(q) dt$$

along the flow of the ODE $\dot{q} = f(q)$. In the setting of optimal control, the goal is to minimize such a functional, constrained to curves satisfying the ODE (see, for example, Aguiar et al. 2021). We will discuss such applications in more detail in Sect. 3.

In coordinates, the augmented adjoint system (2.10) takes the form

$$\dot{q}^A = \frac{\partial H}{\partial p_A} = f^A(q), \tag{2.11a}$$

$$\dot{p}_A = -\frac{\partial H}{\partial q^A} = -p_B \frac{\partial f^B(q)}{\partial q^A} - \frac{\partial L(q)}{\partial q^A}. \tag{2.11b}$$

We now prove various properties of the augmented adjoint system, analogous to the previous section. To start, first note that we can decompose the Hamiltonian vector field X_{H_L} as follows. Let \widehat{f} be the cotangent lift of f . Let $X_L \equiv X_{H_L} - \widehat{f}$. Then, observe that

$$i_{X_L} \Omega = i_{X_{H_L}} \Omega - i_{\widehat{f}} \Omega = dH_L - dH = dL.$$

Thus, we have the decomposition $X_{H_L} = \widehat{f} + X_L$, where \widehat{f} and X_L are the Hamiltonian vector fields for H and L , respectively. In coordinates,

$$X_L = -\frac{\partial L}{\partial q^A} \frac{\partial}{\partial p_A}.$$

From the coordinate expression, we see that X_L is a vertical vector field over the bundle $T^*M \rightarrow M$. We can also see this intrinsically, since dL is a horizontal one-form on T^*M , X_L satisfies $i_{X_L} \Omega = dL$, and Ω restricts to an isomorphism from vertical vector fields on T^*M to horizontal one forms on T^*M . Thus, it is immediate to see intrinsically that an analogous statement to Proposition 2.2 holds, since the flow of \widehat{f} lifts the flow of f , while the flow of X_L is purely vertical. That is, since $T\pi_{T^*M} X_L = 0$,

$$T\pi_{T^*M} X_{H_L} = T\pi_{T^*M} \widehat{f} = f.$$

We can of course also see that the augmented adjoint system lifts the original ODE from the coordinate expression for the augmented adjoint system, (2.11a)–(2.11b).

We now prove analogous statements to Propositions 2.3 and 2.4, modified appropriately for the presence of L in the augmented Hamiltonian.

Proposition 2.7 *Let (q, p) be an integral curve of the augmented adjoint system (2.10) and let $(q, \delta q)$ be an integral curve of the variational equation (2.7), covering the same curve q . Then,*

$$\frac{d}{dt} \langle p, \delta q \rangle = -\langle dL, \delta q \rangle.$$

Remark 2.8 Note that the variational equation associated with the above system is the same as in the nonaugmented case, equation (2.7), since augmenting L to the Hamiltonian system only shifts the Hamiltonian vector field in the vertical direction.

Proof We will prove this in coordinates. We have the equations

$$\dot{p}_A = -p_B \frac{\partial f^B}{\partial q^A} - \frac{\partial L}{\partial q^A},$$

$$\frac{d}{dt} \delta q^B = \frac{\partial f^B}{\partial q^A} \delta q^A.$$

Then,

$$\begin{aligned} \frac{d}{dt} \langle p, \delta q \rangle &= \frac{d}{dt} p_A \delta q^A = \dot{p}_A \delta q^A + p_B \frac{d}{dt} \delta q^B \\ &= -p_B \frac{\partial f^B}{\partial q^A} \delta q^A - \frac{\partial L}{\partial q^A} \delta q^A + p_B \frac{\partial f^B}{\partial q^A} \delta q^A \\ &= -\frac{\partial L}{\partial q^A} \delta q^A = -\langle dL, \delta q \rangle. \end{aligned}$$

□

Remark 2.9 Interestingly, the above proposition states that in the augmented case, $\langle p, \delta q \rangle$ is no longer preserved but rather, its change measures the change of L with respect to the variation δq . This may at first seem contradictory since both the augmented and nonaugmented Hamiltonian vector fields, X_{H_L} and X_H , preserve Ω , and as we noted previously in Remark 2.3, the preservation of the quadratic invariant is implied by symplecticity. However, upon closer inspection, there is no contradiction because the two cases have different first variations, where recall a first variation is a symmetry vector field of the Hamiltonian system and symplecticity can be stated as

$$\frac{d}{dt} \Omega(V, W) = 0,$$

for first variation vector fields V and W . In the nonaugmented case, the equations satisfied by the first variation of the momenta p can be identified with p itself, since the adjoint equation for p is linear in p . On the other hand, in the augmented case, the adjoint equation for p , (2.11b), is no longer linear in p ; rather, it is affine in p . Furthermore, the failure of this equation to be linear in p is given precisely by $-dL$. Thus, in the augmented case, first variations in p can no longer be identified with p , and this leads to the additional term $-\langle dL, \delta q \rangle$ in the above proposition.

To prove an analogous statement to Proposition 2.4, we need the additional assumption that the symmetry vector field g leaves L invariant, $\mathcal{L}_g L = 0$.

Proposition 2.8 *Let g be a symmetry of the ODE $\dot{q} = f(q)$, i.e., $[g, f] = 0$. Additionally, assume that g is a symmetry of L , i.e., $\mathcal{L}_g L = 0$. Then, its cotangent lift \widehat{g} is a symmetry of the augmented adjoint system, $[\widehat{g}, X_{H_L}] = 0$ and additionally, the function*

$$\langle \Theta, \widehat{g} \rangle$$

on T^*M is preserved along the flow of X_{H_L} .

Proof To see that $[\widehat{g}, X_{H_L}] = 0$, note that with the decomposition $X_{H_L} = \widehat{f} + X_L$, we have

$$[\widehat{g}, X_{H_L}] = [\widehat{g}, \widehat{f}] + [\widehat{g}, X_L] = [\widehat{g}, X_L],$$

where we used that $[\widehat{g}, \widehat{f}] = [\widehat{g}, f] = 0$. To see that $[\widehat{g}, X_L] = 0$, we note that $[\widehat{g}, X_L]$ can be expressed as

$$[\widehat{g}, X_L] = \mathcal{L}_{\widehat{g}}X_L = \mathcal{L}_{\widehat{g}}(\Omega^{-1}(dL)),$$

where we interpret $\Omega : T(T^*M) \rightarrow T^*(T^*M)$. Then, note that \widehat{g} preserves Ω since \widehat{g} is a cotangent lift and it also preserves L (where, since we identify L with its pullback through π_{T^*M} , this is equivalent to g preserving L). More precisely, since we are identifying L with its pullback $(\pi_{T^*M})^*L$, we have

$$\mathcal{L}_{\widehat{g}}((\pi_{T^*M})^*L) = \langle (\pi_{T^*M})^*dL, \widehat{g} \rangle = \langle dL, T\pi_{T^*M}\widehat{g} \rangle = \langle dL, g \rangle = \mathcal{L}_g L = 0.$$

Hence, $\mathcal{L}_{\widehat{g}}(\Omega^{-1}(dL)) = 0$. One can also verify this in coordinates, and a direct computation yields

$$[\widehat{g}, X_L] = \frac{\partial}{\partial q^A} \left(g^B(q) \frac{\partial L}{\partial q^B} \right) \frac{\partial}{\partial p_A},$$

which vanishes since $\mathcal{L}_g L = 0$.

Now, to show that $\langle \Theta, \widehat{g} \rangle$ is preserved along the flow of X_{H_L} , compute

$$\mathcal{L}_{X_{H_L}} \langle \Theta, \widehat{g} \rangle = \mathcal{L}_{\widehat{f}} \langle \Theta, \widehat{g} \rangle + \mathcal{L}_{X_L} \langle \Theta, \widehat{g} \rangle = \mathcal{L}_{X_L} \langle \Theta, \widehat{g} \rangle,$$

where we used that $\mathcal{L}_{\widehat{f}} \langle \Theta, \widehat{g} \rangle = 0$ by Proposition 2.4. Now, we have

$$\begin{aligned} \mathcal{L}_{X_{H_L}} \langle \Theta, \widehat{g} \rangle &= \mathcal{L}_{X_L} \langle \Theta, \widehat{g} \rangle = \langle \mathcal{L}_{X_L} \Theta, \widehat{g} \rangle + \langle \Theta, \mathcal{L}_{X_L} \widehat{g} \rangle \\ &= \langle \mathcal{L}_{X_L} \Theta, \widehat{g} \rangle + \langle \Theta, \underbrace{[X_L, \widehat{g}]}_{=0} \rangle \\ &= \langle i_{X_L} d\Theta + d(i_{X_L} \Theta), \widehat{g} \rangle = \langle -i_{X_L} \Omega, \widehat{g} \rangle + \langle d(i_{X_L} \Theta), \widehat{g} \rangle \\ &= -\langle dL, \widehat{g} \rangle + \langle d(i_{X_L} \Theta), \widehat{g} \rangle. \end{aligned}$$

The first term above vanishes since $\mathcal{L}_g L = 0$. Furthermore, $\langle d(i_{X_L} \Theta), \widehat{g} \rangle = 0$ since X_L is a vertical vector field while Θ is a horizontal one-form. Hence, $\mathcal{L}_{X_{H_L}} \langle \Theta, \widehat{g} \rangle = 0$. □

2.3 Adjoint Systems for DAEs via Presymplectic Mechanics

In this section, we generalize the notion of adjoint system to the case where the base equation is a (semi-explicit) DAE. We will prove analogous results to the ODE case.

However, more care is needed than the ODE case, since the DAE constraint introduces issues with solvability. As we will see, the adjoint system associated with a DAE is a presymplectic system, so we will approach the solvability of such systems through the presymplectic constraint algorithm.

We consider the following setup for a differential-algebraic equation. Let M_d and M_a be two manifolds, where we regard M_d as the configuration space of the “dynamical” or “differential” variables and M_a as the configuration space of the “algebraic” variables. Let $\pi_\Phi : \Phi \rightarrow M_d \times M_a$ be a vector bundle over $M_d \times M_a$. Furthermore, let $\pi_d : M_d \times M_a \rightarrow M_d$ be the projection onto the first factor and let $\pi_{\overline{TM}_d} : \overline{TM}_d \rightarrow M_d \times M_a$ be the pullback bundle of the tangent bundle $\pi_{TM_d} : TM_d \rightarrow M_d$ by π_d , i.e., $\overline{TM}_d = \pi_d^*(TM_d)$. Then, a (semi-explicit) DAE is specified by a section $f \in \Gamma(\overline{TM}_d)$ and a section $\phi \in \Gamma(\Phi)$, via the system

$$\dot{q} = f(q, u), \tag{2.12a}$$

$$0 = \phi(q, u), \tag{2.12b}$$

where (q, u) are coordinates on $M_d \times M_a$. We refer to \overline{TM}_d as the differential tangent bundle, with coordinates (q, u, v) and to Φ as the constraint bundle.

Remark 2.10 For the local solvability of (2.12a)–(2.12b), regard ϕ locally as a map $\mathbb{R}^{\dim(M_d)} \times \mathbb{R}^{\dim(M_a)} \rightarrow \mathbb{R}^{\text{rank}(\Phi)}$. If $\partial\phi/\partial u$ is an isomorphism at a point (q_0, u_0) where $\Phi(q_0, u_0) = 0$, then by the implicit function theorem, one can locally solve $u = u(q)$ about (q_0, u_0) such that $\phi(q, u(q)) = 0$, and subsequently solve the unconstrained differential equation $\dot{q} = f(q, u(q))$ locally. This is the case for semi-explicit index 1 DAEs.

In order for the $\text{rank}(\Phi) \times \dim(M_a)$ matrix $\partial\phi/\partial u(q_0, u_0)$ to be an isomorphism, it is necessary that $\text{rank}(\Phi) = \dim(M_a)$. However, we will make no such assumption, so as to treat the theory in full generality, allowing for, e.g., nonunique solutions. We will, however, assume that the $D\phi$ has constant rank (this corresponds to a fixed index DAE, which is the case if $\partial\phi/\partial u$ is a pointwise isomorphism) since we utilize the results of presymplectic geometry for constant rank presymplectic manifolds, as discussed in Sect. 1.2.

Now, let $\overline{T^*M}_d$ be the pullback bundle of the cotangent bundle T^*M_d by π_d , with coordinates (q, u, p) , which we refer to as the differential cotangent bundle. Furthermore, let Φ^* be the dual vector bundle to Φ , with coordinates (q, u, λ) . Let $\overline{T^*M}_d \oplus \Phi^*$ be the Whitney sum of these two vector bundles over $M_d \times M_a$ with coordinates (q, u, p, λ) , which we refer to as the generalized phase space bundle. We define a Hamiltonian on the generalized phase space,

$$H : \overline{T^*M}_d \oplus \Phi^* \rightarrow \mathbb{R},$$


$$H(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle.$$

Let Ω_d denote the canonical symplectic form on T^*M_d , with coordinate expression $\Omega_d = dq \wedge dp$. We define a presymplectic form Ω_0 on $\overline{T^*M}_d \oplus \Phi^*$ as follows: the

pullback bundle admits the map $\tilde{\pi}_d : \overline{T^*M}_d \rightarrow T^*M_d$ which covers π_d and acts as the identity on fibers and furthermore, the generalized phase space bundle admits the projection $\Pi : \overline{T^*M}_d^* \oplus \Phi^* \rightarrow \overline{T^*M}_d^*$, since the Whitney sum has the structure of a double vector bundle. Hence, we can pullback Ω_d along the sequence of maps

$$\overline{T^*M}_d \oplus \Phi^* \xrightarrow{\Pi} \overline{T^*M}_d \xrightarrow{\tilde{\pi}_d} T^*M_d,$$

which allows us to define a two-form $\Omega_0 \equiv \Pi^* \circ \tilde{\pi}_d^*(\Omega_d)$ on the generalized phase space bundle. Clearly, Ω_0 is closed as the pullback of a closed form. In general, Ω_0 will be degenerate except in the trivial case where M_d is empty and the fibers of Φ are the zero vector space. Hence, Ω_0 is a presymplectic form. Note that since Π acts by projection and $\tilde{\pi}_d$ acts as the identity on fibers, the coordinate expression for Ω_0 on $\overline{T^*M}_d \oplus \Phi^*$ with coordinates (q, u, p, λ) is the same as the coordinate expression for Ω_d , $\Omega_0 = dq \wedge dp$. The various spaces and their coordinates are summarized in the diagram below.



We now define the adjoint system associated with the DAE (2.12a)–(2.12b) as the Hamiltonian system

$$i_X \Omega_0 = dH. \tag{2.13}$$

Given a (generally, partially defined) vector field X on the generalized phase space satisfying (2.13), we say a curve $(q(t), u(t), p(t), \lambda(t))$ is a solution curve of (2.13) if it is an integral curve of X .

Let us find a coordinate expression for the above system. Expressing the coordinates with indices $(q^i, u^a, p_j, \lambda_A)$, the left-hand side of (2.13) along a solution curve has the expression

$$\begin{aligned} i_X \Omega_0 &= \left(\dot{q}^i \frac{\partial}{\partial q^i} + \dot{u}^a \frac{\partial}{\partial u^a} + \dot{p}_j \frac{\partial}{\partial p_j} + \dot{\lambda}_A \frac{\partial}{\partial \lambda_A} \right) \lrcorner dq^k \wedge dp_k \\ &= \dot{q}^i dp_i - \dot{p}_j dq^j. \end{aligned}$$

On the other hand, the right-hand side of (2.13) has the expression

$$dH = d\left(p_i f^i(q, u) + \lambda_A \phi^A(q, u) \right)$$

$$\begin{aligned}
 &= f^i(q, u)dp_i + \left(p_i \frac{\partial f^i}{\partial q^j} + \lambda_A \frac{\partial \phi^A}{\partial q^j} \right) dq^j + \phi^A(q, u)d\lambda_A \\
 &\quad + \left(p_i \frac{\partial f^i}{\partial u^a} + \lambda_A \frac{\partial \phi^A}{\partial u^a} \right) du^a.
 \end{aligned}$$

Equating these expressions gives the coordinate expression for the adjoint DAE system,

$$\dot{q}^i = f^i(q, u), \tag{2.14a}$$

$$\dot{p}_j = -p_i \frac{\partial f^i}{\partial q^j} - \lambda_A \frac{\partial \phi^A}{\partial q^j}, \tag{2.14b}$$

$$0 = \phi^A(q, u), \tag{2.14c}$$

$$0 = p_i \frac{\partial f^i}{\partial u^a} + \lambda_A \frac{\partial \phi^A}{\partial u^a}. \tag{2.14d}$$

Remark 2.11 As mentioned in Remark 2.10, in the index 1 case, one can locally solve the original DAE (2.14a) and (2.14c). Viewing such a solution (q, u) as fixed, one can subsequently locally solve for λ in equation (2.14d) as a function of p , since $\partial\phi/\partial u$ is locally invertible. Substituting this into (2.14b) gives an ODE solely in the variable p , which can be solved locally.

Stated another way, if the original DAE (2.12a)–(2.12b) is an index 1 system, then the adjoint DAE system (2.14a)–(2.14d) is an index 1 system with dynamical variables (q, p) and algebraic variables (u, λ) . To see this, if one denotes the constraints for the adjoint system (2.14c) and (2.14d) as

$$0 = \tilde{\phi}(q, u, p, \lambda) \equiv \begin{pmatrix} \phi^A(q, u) \\ p_i \frac{\partial f^i}{\partial u^a} + \lambda_A \frac{\partial \phi^A}{\partial u^a} \end{pmatrix},$$

then the matrix derivative of $\tilde{\phi}$ with respect to the algebraic variables (u, λ) can be locally expressed in block form as

$$\begin{pmatrix} \partial\phi/\partial u & A \\ 0 & \partial\phi/\partial u \end{pmatrix},$$

where the block A has components given by the derivative of the right-hand side of (2.14d) with respect to u . It is clear from the block triangular form of this matrix that it is pointwise invertible if $\partial\phi/\partial u$ is.

Remark 2.12 It is clear from the coordinate expression (2.14a)–(2.14d) that a solution curve of the adjoint DAE system, if it exists, covers a solution curve of the original DAE system.

We now prove several results regarding the structure of the adjoint DAE system.

First, we show that the constraint equations (2.14c)–(2.14d) can be interpreted as the statement that the Hamiltonian H has the same time dependence as the “dynamical” Hamiltonian,

$$H_d : \overline{T^*M_d} \oplus \Phi^* \rightarrow \mathbb{R},$$

$$H_d(q, u, p, \lambda) = \langle p, f(q, u) \rangle,$$

when evaluated along a solution curve.

Proposition 2.9 *For a solution curve (q, u, p, λ) of (2.13),*

$$\frac{d}{dt} H(q(t), u(t), p(t), \lambda(t)) = \frac{d}{dt} H_d(q(t), u(t), p(t), \lambda(t)).$$

Proof For brevity, all functions below are appropriately evaluated along the solution curve. We have

$$\begin{aligned} \frac{d}{dt} H &= \frac{\partial H}{\partial q^i} \dot{q}^i + \frac{\partial H}{\partial p_j} \dot{p}_j + \frac{\partial H}{\partial u^a} \dot{u}^a + \frac{\partial H}{\partial \lambda_A} \dot{\lambda}_A \\ &= \frac{\partial H}{\partial q^i} \dot{q}^i + \frac{\partial H}{\partial p_j} \dot{p}_j + \left(p_i \frac{\partial f^i}{\partial u^a} + \lambda_A \frac{\partial \phi^A}{\partial u^a} \right) \dot{u}^a + \phi^A \dot{\lambda}_A \\ &= \frac{\partial H}{\partial q^i} \dot{q}^i + \frac{\partial H}{\partial p_j} \dot{p}_j \\ &= \frac{\partial H_d}{\partial q^i} \dot{q}^i + \frac{\partial H_d}{\partial p_j} \dot{p}_j = \frac{d}{dt} H_d, \end{aligned}$$

where in the third equality, we used (2.14c) and (2.14d). □

Remark 2.13 A more geometric way to view the above proposition is as follows: note that if a partially defined vector field X exists such that $i_X \Omega_0 = dH$, then the change of H in a given direction Y , at any point where X is defined, can be computed as $dH(Y) = \Omega_0(X, Y)$. Observe that the kernel of Ω_0 is locally spanned by $\partial/\partial u, \partial/\partial \lambda$, i.e., it is spanned by the coordinate vectors in the algebraic coordinates. Hence, the change of H in the algebraic coordinate directions is zero. This justifies referring to (u, λ) as “algebraic” variables.

We now show that the adjoint system (2.14a)–(2.14d) formally arises from a variational principle. To do so, let Θ_0 be the pullback of the tautological one-form Θ_d on the cotangent bundle T^*M_d by the maps Π and $\tilde{\pi}_d$, $\Theta_0 = \Pi^* \circ \tilde{\pi}_d^*(\Theta_d)$. Of course, one has $\Omega_0 = -d\Theta_0$. Consider the action S defined by

$$S[\psi] = \int_I [\psi^* \Theta_0 - (H \circ \psi) dt],$$

where $\psi(t) = (q(t), u(t), p(t), \lambda(t))$ is a curve on the generalized phase space bundle over the interval $I = (t_0, t_1)$. We consider the variational principle $\delta S[\psi] = 0$, subject to variations which fix the endpoints $q(t_0), q(t_1)$.

Proposition 2.10 *Let ψ be a curve on the generalized phase space bundle over the interval I . Then, ψ is a stationary point of S with respect to variations which fix $q(t_0)$, $q(t_1)$ if and only if (2.14a)–(2.14d) hold.*

Proof In $\psi = (q, u, p, \lambda)$ coordinates, the action has the expression

$$\begin{aligned} S[q, u, p, \lambda] &= \int_I \left(p_i \dot{q}^i - p_i f^i(q, u) - \lambda_A \phi^A(q, u) \right) dt \\ &= \int_I \left(p_i (\dot{q}^i - f^i(q, u)) - \lambda_A \phi^A(q, u) \right) dt. \end{aligned}$$

The variation of the action reads

$$\begin{aligned} \delta S[q, u, p, \lambda] \cdot (\delta q, \delta u, \delta p, \delta \lambda) &= \int_I \left[\delta p_i (\dot{q}^i - f^i) + p_j \frac{d}{dt} \delta q^j - p_i \frac{\partial f^i}{\partial q^j} \delta q^j - \lambda_A \frac{\partial \phi^A}{\partial q^j} \delta q^j \right. \\ &\quad \left. - \delta \lambda_A \phi^A + \left(-p_i \frac{\partial f^i}{\partial u^a} - \lambda_A \frac{\partial \phi^A}{\partial u^a} \right) \delta u^a \right] dt \\ &= \int_I \left[\delta p_i (\dot{q}^i - f^i) - \left(\dot{p}_j + p_i \frac{\partial f^i}{\partial q^j} + \lambda_A \frac{\partial \phi^A}{\partial q^j} \right) \delta q^j \right. \\ &\quad \left. - \delta \lambda_A \phi^A + \left(-p_i \frac{\partial f^i}{\partial u^a} - \lambda_A \frac{\partial \phi^A}{\partial u^a} \right) \delta u^a \right] dt, \end{aligned}$$

where we used integration by parts and the vanishing of the variations at the endpoints to drop any boundary terms. Clearly, if (2.14a)–(2.14d) hold, then $\delta S = 0$ for all such variations. Conversely, by the fundamental lemma of the calculus of variations, if $\delta S = 0$ for all such variations, then (2.14a)–(2.14d) hold. \square

Remark 2.14 We will use the variational structure associated with the adjoint DAE system to construct numerical integrators in Sect. 3.2.

We now prove a result regarding the conservation of a quadratic invariant, analogous to the case of cotangent lifted adjoint systems in the ODE case. To do this, we define the variational equations as the linearization of the DAE (2.12a)–(2.12b). The coordinate expressions for the variational equations are obtained by taking the variation of equations (2.12a)–(2.12b) with respect to variations $(\delta q, \delta u)$,

$$\dot{q}^i = f^i(q, u), \tag{2.15a}$$

$$0 = \phi^A(q, u), \tag{2.15b}$$

$$\frac{d}{dt} \delta q^i = \frac{\partial f^i(q, u)}{\partial q^j} \delta q^j + \frac{\partial f^i(q, u)}{\partial u^a} \delta u^a, \tag{2.15c}$$

$$0 = \frac{\partial \phi^A(q, u)}{\partial q^j} \delta q^j + \frac{\partial \phi^A(q, u)}{\partial u^a} \delta u^a. \tag{2.15d}$$

Proposition 2.11 For a solution (q, u, p, λ) of the adjoint DAE system (2.14a)–(2.14d) and a solution $(q, u, \delta q, \delta u)$ of the variational equations (2.15a)–(2.15d), covering the same curve (q, u) , one has

$$\frac{d}{dt} \langle p(t), \delta q(t) \rangle = 0.$$

Proof This follows from a direct computation,

$$\begin{aligned} \frac{d}{dt} \langle p, \delta q \rangle &= \frac{d}{dt} \left(p_i \delta q^i \right) = \dot{p}_j \delta q^j + p_i \frac{d}{dt} \delta q^i \\ &= -p_i \frac{\partial f^i}{\partial q^j} \delta q^j - \lambda_A \frac{\partial \phi^A}{\partial q^j} \delta q^j + p_i \frac{\partial f^i}{\partial q^j} \delta q^j + p_i \frac{\partial f^i}{\partial u^a} \delta u^a \\ &= -\lambda_A \frac{\partial \phi^A}{\partial q^j} \delta q^j + p_i \frac{\partial f^i}{\partial u^a} \delta u^a \\ &= \left(\lambda_A \frac{\partial \phi^A}{\partial u^a} + p_i \frac{\partial f^i}{\partial u^a} \right) \delta u^a = 0, \end{aligned}$$

where we used (2.14b), (2.15c), (2.15d), and (2.14d). □

Remark 2.15 Although we proved the previous proposition in coordinates, it can be understood intrinsically through the presymplecticity of the adjoint DAE flow. To see this, assume a partially defined vector field X exists such that $i_X \Omega_0 = dH$. Then, the flow of X preserves Ω_0 , which follows from

$$\mathcal{L}_X \Omega_0 = i_X d\Omega_0 + d(i_X \Omega_0) = d(i_X \Omega_0) = d^2 H = 0.$$

The coordinate expression for the preservation of the presymplectic form $\Omega_0 = dq^i \wedge dp_i$, with the appropriate choice of first variations, gives the previous proposition, analogous to the argument that we made in the symplectic (unconstrained) case.

Additionally, as we will see in Sect. 3.1, Proposition 2.11 will provide a method for computing adjoint sensitivities.

These two observations are interesting when constructing numerical methods to compute adjoint sensitivities, since if we can construct integrators that preserve the presymplectic form, then it will preserve the quadratic invariant and hence, be suitable for computing adjoint sensitivities efficiently.

Remark 2.16 For an index 1 DAE (2.12a)–(2.12b), since $\partial \phi / \partial u$ is (pointwise) invertible for a fixed curve (q, u) , one can solve for δu as a function of δq in the variational equation (2.15d) and substitute this into (2.15c) to obtain an explicit ODE for δq . Hence, in the index 1 case, given a solution (q, u) of the DAE (2.12a)–(2.12b) and an initial condition $\delta q(0)$ in the tangent fiber over $q(0)$, there is a corresponding (at least local) unique solution of the variational equations.

2.3.1 DAE Index and the Presymplectic Constraint Algorithm

In this section, we relate the index of the DAE (2.12a)–(2.12b) to the number of steps for convergence in the presymplectic constraint algorithm associated with the adjoint DAE system (2.13). In particular, we show that for an index 1 DAE, the presymplectic constraint algorithm for the associated adjoint DAE system terminates after $\nu_P = 1$ step. Subsequently, we discuss how one can formally handle the more general index ν DAE case.

We consider again the presymplectic system given by the adjoint DAE system, $P = T^*M_d \oplus \Phi^*$ equipped with the presymplectic form $\Omega_0 = dq \wedge dp$ and Hamiltonian $H(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle$, as discussed in the previous section. Our goal is to bound the number of steps in the presymplectic constraint algorithm ν_P for this presymplectic system in terms of the index ν of the underlying DAE (2.12a)–(2.12b).

Recall the presymplectic constraint algorithm discussed in Sect. 1.2. We first determine the primary constraint manifold P_1 . Observe that since $\Omega_0 = dq \wedge dp$, we have the local expression $\ker(\Omega_0)|_{(q,u,p,\lambda)} = \text{span}\{\partial/\partial u, \partial/\partial \lambda\}$. Thus, we require that

$$\begin{aligned} \frac{\partial H}{\partial u} &= 0, \\ \frac{\partial H}{\partial \lambda} &= 0, \end{aligned}$$

i.e., P_1 consists of the points (q, u, p, λ) such that

$$\begin{aligned} 0 &= \frac{\partial H(q, u, p, \lambda)}{\partial u^a} = p_i \frac{\partial f^i(q, u)}{\partial u^a} + \lambda_A \frac{\partial \phi^A(q, u)}{\partial u^a}, \\ 0 &= \frac{\partial H(q, u, p, \lambda)}{\partial \lambda^A} = \phi^A(q, u). \end{aligned}$$

These are of course the constraint equations (2.14c)–(2.14d) of the adjoint DAE system.

We now consider first the case when the DAE system (2.12a)–(2.12b) has index $\nu = 1$ and subsequently, consider the general case $\nu \geq 1$.

The Presymplectic Constraint Algorithm for $\nu = 1$. For the case $\nu = 1$, we will show that the presymplectic constraint algorithm terminates after 1 step, i.e., $\nu_P = \nu = 1$.

Now, assume that the DAE system (2.12a)–(2.12b) has index $\nu = 1$, i.e., for each $(q, u) \in M_d \times M_a$ such that $\phi(q, u) = 0$, the matrix with A^{th} row and a^{th} column entry

$$\frac{\partial \phi^A(q, u)}{\partial u^a}$$

is invertible. Observe that the definition of the presymplectic constraint algorithm, equation (1.1), is local and hence, we seek a local coordinate expression for $\Omega_1 \equiv \Omega_0|_{P_1}$ and its kernel.

Let $(q, u, p, \lambda) \in P_1$. In particular, $\phi(q, u) = 0$. Since $\partial\phi(q, u)/\partial u$ is invertible, by the implicit function theorem, one can locally solve for u as a function of q , which we denote $u = u(q)$, such that $\phi(q, u(q)) = 0$. Then, one can furthermore locally solve for λ as a function of q and p from the second constraint equation,

$$\lambda_A(q, p) = - \left[\left(\frac{\partial\phi(q, u(q))}{\partial u} \right)^{-1} \right]_A^a p_i \frac{\partial f^i(q, u(q))}{\partial u^a}.$$

Thus, we can coordinatize P_1 via coordinates (q', p') , where the inclusion $i_1 : P_1 \hookrightarrow P$ is given by the coordinate expression

$$i_1 : (q', p') \mapsto (q', u(q'), p', \lambda(q', p')).$$

Then, one obtains the local expression for Ω_1 ,

$$\Omega_1 = i_1^* \Omega_0 = i_1^*(dq) \wedge i_1^*(dp) = dq' \wedge dp'.$$

This is clearly nondegenerate, i.e., $Z_p = 0$ for any $Z \in \ker(\Omega_1)$, $p \in P_1$, so the presymplectic constraint algorithm terminates, $P_2 = P_1$. We conclude that $\nu_P = 1$.

To conclude the discussion of the index 1 case, we obtain coordinate expressions for the resulting nondegenerate Hamiltonian system. The Hamiltonian on P_1 can be expressed as

$$\begin{aligned} H_1(q', p') &= H(i_1(q', p')) = \langle p', f(q', u(q')) \rangle \\ &\quad + \langle \lambda(q', p'), \phi(q', u(q')) \rangle = \langle p', f(q', u(q')) \rangle. \end{aligned}$$

Thus, with the coordinate expression $X = \dot{q}^i \partial/\partial q^i + \dot{p}'_i \partial/\partial p'_i$, Hamilton's equations $i_X \Omega_1 = dH_1$ can be expressed as

$$\begin{aligned} \dot{q}^i &= \frac{\partial H_1}{\partial p'_i} = f^i(q', u(q')), \\ \dot{p}'_i &= -\frac{\partial H_1}{\partial q^i} = -p'_j \frac{\partial f^j(q', u(q'))}{\partial q^i} - p'_j \frac{\partial f^j(q', u(q'))}{\partial u^a} \frac{\partial u^a(q')}{\partial q^i}. \end{aligned}$$

We will now show explicitly that this Hamiltonian system solves (2.14a)–(2.14d) along the submanifold P_1 . Clearly, the latter two equations (2.14c)–(2.14d) are satisfied, by definition of P_1 . So, we want to show that the first two equations (2.14a)–(2.14b) are satisfied. Using the second constraint equation (2.14d), we have

$$-p'_j \frac{\partial f^j(q', u(q'))}{\partial u^a} = \lambda_A(q', p') \frac{\partial \phi^A(q', u(q'))}{\partial u^a}.$$

Substituting this into the equation for \dot{p}'_i above gives

$$\dot{p}'_i = -p'_j \frac{\partial f^j(q', u(q'))}{\partial q^i} + \lambda_A(q', p') \frac{\partial \phi^A(q', u(q'))}{\partial u^a} \frac{\partial u^a(q')}{\partial q^i}.$$

By the implicit function theorem, one has

$$\frac{\partial \phi^A(q', u(q'))}{\partial u^a} \frac{\partial u^a(q')}{\partial q^i} = - \frac{\partial \phi^A(q', u(q'))}{\partial q^i}.$$

Hence, the Hamiltonian system on P_1 can be equivalently expressed as

$$\begin{aligned} \dot{q}^i &= f^i(q', u(q')), \\ \dot{p}'_i &= -p'_j \frac{\partial f^j(q', u(q'))}{\partial q^i} - \lambda_A(q', p') \frac{\partial \phi^A(q', u(q'))}{\partial q^i}. \end{aligned}$$

Thus, we have explicitly verified that (2.14a)–(2.14d) are satisfied along P_1 . Note that since the presymplectic constraint algorithm terminates at $\nu_P = 1$, X is guaranteed to be tangent to P_1 . One can also verify this explicitly by computing the pushforward $Ti_1(X)$ and verifying that it annihilates the constraint functions whose zero level set defines P_1 ,

$$\begin{aligned} (q, u, p, \lambda) &\mapsto \phi^A(q, u), \\ (q, u, p, \lambda) &\mapsto p_i \frac{\partial f^i(q, u)}{\partial u^a} + \lambda_A \frac{\partial \phi^A(q, u)}{\partial u^a}. \end{aligned}$$

Remark 2.17 It is interesting to note that the Hamiltonian system $i_X \Omega_1 = dH_1$, which we obtained by forming the adjoint system of the underlying index 1 DAE and subsequently, reducing the index of the adjoint DAE system through the presymplectic constraint algorithm, can be equivalently obtained (at least locally) by first reducing the index of the underlying DAE and then forming the adjoint system.

More precisely, if one locally solves $\phi(q, u) = 0$ for $u = u(q)$, then the index 1 DAE can be reduced to an ODE,

$$\dot{q} = f(q, u(q)).$$

Subsequently, we can form the adjoint system to this ODE, as discussed in Sect. 2.2. The corresponding Hamiltonian is $H(q, p) = \langle p, f(q, u(q)) \rangle$, which is the same as H_1 .

Thus, for the index 1 case, the process of forming the adjoint system and reducing the index commute.

Remark 2.18 In the language of the presymplectic constraint algorithm, Proposition 2.9 can be restated as the statement that the Hamiltonian H and its first derivatives, restricted to the primary constraint manifold, agrees with the dynamical Hamiltonian H_1 and its first derivatives.

Remark 2.19 An alternative view of the solution theory of the presymplectic adjoint DAE system (2.14a)–(2.14d) is through singular perturbation theory (see, for example,

Berglund 2007 and Chen and Trenn 2021). We proceed by writing (2.14a)–(2.14d) as

$$\begin{aligned}\dot{q} &= \frac{\partial H}{\partial p} = f(q, u), \\ \dot{p} &= -\frac{\partial H}{\partial q} = -[D_q f(q, u)]^* p - [D_q \phi(q, u)]^* \lambda, \\ 0 &= \frac{\partial H}{\partial \lambda} = \phi(q, u), \\ 0 &= -\frac{\partial H}{\partial u} = -[D_u f(q, u)]^* p - [D_u \phi(q, u)]^* \lambda.\end{aligned}$$

Applying a singular perturbation to the constraint equations yields the system

$$\begin{aligned}\dot{q} &= \frac{\partial H}{\partial p}, \\ \dot{p} &= -\frac{\partial H}{\partial q}, \\ \epsilon \dot{u} &= \frac{\partial H}{\partial \lambda}, \\ \epsilon \dot{\lambda} &= -\frac{\partial H}{\partial u},\end{aligned}$$

where $\epsilon > 0$. Observe that this is a nondegenerate Hamiltonian system with $H(q, u, p, \lambda)$ as previously defined but with the modified symplectic form $\Omega_\epsilon = dq \wedge dp + \epsilon du \wedge d\lambda$. Then, the above system can be expressed $i_{X_H} \Omega_\epsilon = dH$. In the language of perturbation theory, the primary constraint manifold for the presymplectic system is precisely the slow manifold of the singularly perturbed system. One can utilize techniques from singular perturbation theory to develop a solution theory for this system, using Tihonov's theorem, whose assumptions for this particular system depend on the eigenvalues of the algebraic Hessian $D_{u,\lambda}^2 H$ (see Berglund 2007). Although we will not elaborate on this here, this could be an interesting approach for the existence, stability, and approximation theory of such systems. In particular, the slow manifold integrators introduced in Burby and Klotz (2020) may be relevant to their discretization. It is also interesting to note that for a solution $(q_\epsilon, p_\epsilon, u_\epsilon, \lambda_\epsilon)$ of the singularly perturbed system and a solution $(\delta q_\epsilon, \delta u_\epsilon)$ of the variational equations,

$$\begin{aligned}\frac{d}{dt} \delta q_\epsilon &= D_q f(q_\epsilon, u_\epsilon) \delta q_\epsilon + D_u f(q_\epsilon, u_\epsilon) \delta u_\epsilon, \\ \epsilon \frac{d}{dt} \delta u_\epsilon &= D_q \phi(q_\epsilon, u_\epsilon) \delta q_\epsilon + D_u \phi(q_\epsilon, u_\epsilon) \delta u_\epsilon,\end{aligned}$$

one has the perturbed adjoint variational quadratic conservation law

$$\frac{d}{dt} \left(\langle p_\epsilon, \delta q_\epsilon \rangle + \epsilon \langle \lambda_\epsilon, \delta u_\epsilon \rangle \right) = 0,$$

which follows immediately from the preservation of Ω_ϵ under the symplectic flow.

The Presymplectic Constraint Algorithm for General $\nu \geq 1$. Note that for the general case, we assume that the index of the DAE is finite, $1 \leq \nu < \infty$.

In this case, there are two possible approaches to reduce the adjoint system: either form the adjoint system associated with the index ν DAE and then successively apply the presymplectic constraint algorithm or, alternatively, reduce the index of the DAE, form the adjoint system, and then apply the presymplectic constraint algorithm as necessary.

Since we have already worked out the presymplectic constraint algorithm for the index 1 case, we will take the latter approach. Namely, we reduce an index ν DAE to an index 1 DAE, and subsequently, apply the presymplectic constraint algorithm to the reduced index 1 DAE. Given an index ν DAE, it is generally possible to reduce the DAE to an index 1 DAE using the algorithm introduced in Mattsson and Söderlind (1993). The process of index reduction is given by differentiating the equations of the DAE to reveal hidden constraints. Geometrically, the process of index reduction can be understood as the successive jet prolongation of the DAE and subsequent projection back onto the first jet (see, Reid et al. 2001).

Thus, given an index ν DAE $\dot{x} = \tilde{f}(x, y)$, $\tilde{\phi}(x, y) = 0$, we can, after $\nu - 1$ reduction steps, transform it into an index 1 DAE of the form $\dot{q} = f(q, u)$, $\phi(q, u) = 0$. Subsequently, we can form the adjoint DAE system and apply one iteration of the presymplectic constraint algorithm to obtain the underlying nondegenerate dynamical system. If we let the $\nu_{R,P}$ denote the minimum number of DAE index reduction steps plus presymplectic constraint algorithm iterations necessary to take an index ν DAE and obtain the underlying nondegenerate Hamiltonian system associated with the adjoint, we have $\nu_{R,P} \leq \nu$.

Remark 2.20 Note that we could have reduced the index ν DAE to an explicit ODE after ν reduction steps, and subsequently, formed the adjoint. While this is formally equivalent to the above procedure by Remark 2.17, we prefer to keep the DAE in index 1 form. This is especially preferable from the viewpoint of numerics: if one reduces an index 1 DAE to an ODE and attempts to apply a numerical integrator, it is generically the case that the discrete flow drifts off the constraint manifold. For this reason, it is preferable to develop numerical integrators for the index 1 adjoint DAE system directly to prevent constraint violation.

Example 2.1 (Hessenberg Index 2 DAE) Consider a Hessenberg index 2 DAE, i.e., a DAE of the form

$$\begin{aligned}\dot{q} &= f(q, u), \\ 0 &= g(q),\end{aligned}$$

where $(q, u) \in \mathbb{R}^n \times \mathbb{R}^m$, $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and $\frac{\partial g}{\partial q} \frac{\partial f}{\partial u}$ is pointwise invertible. We reduce this to an index 1 DAE (2.12a)–(2.12b) as follows. Let $M_d = g^{-1}(\{0\})$ be the dynamical configuration space which we will assume is a submanifold of \mathbb{R}^n . For example, this is true if g is a constant rank map. Furthermore,

let $M_a = \mathbb{R}^m$ be the algebraic configuration space. To reduce the index, we differentiate the constraint $g(q) = 0$ with respect to time. This is equivalent to enforcing that the dynamics are tangent to M_d . This gives

$$0 = \frac{\partial g^A(q)}{\partial q^i} \dot{q}^i = \frac{\partial g^A(q)}{\partial q^i} f^i(q, u) \equiv \phi^A(q, u).$$

Hence, we can form the semi-explicit index 1 system on $M_d \times M_a$ given by

$$\begin{aligned} \dot{q} &= f(q, u), \\ 0 &= \phi(q, u). \end{aligned}$$

The above system is an index 1 DAE since $\frac{\partial \phi}{\partial u} = \frac{\partial g}{\partial q} \frac{\partial f}{\partial u}$ is pointwise invertible.

We now form the adjoint DAE system associated with this index 1 DAE, (2.14a)–(2.14d). Expressing the constraint in terms of g and f , instead of ϕ , gives

$$\begin{aligned} \dot{q}^i &= f^i(q, u), \\ \dot{p}_j &= -p_i \frac{\partial f^i(q, u)}{\partial q^j} - \lambda_A \left(\frac{\partial^2 g^A(q)}{\partial q^j \partial q^i} f^i(q, u) + \frac{\partial g^A(q)}{\partial q^i} \frac{\partial f^i(q, u)}{\partial q^j} \right), \\ 0 &= \frac{\partial g^A(q)}{\partial q^i} f^i(q, u), \\ 0 &= p_i \frac{\partial f^i(q, u)}{\partial u^a} + \lambda_A \left(\frac{\partial g^A(q)}{\partial q^i} \frac{\partial f^i(q, u)}{\partial u^a} \right). \end{aligned}$$

We can then apply one iteration of the presymplectic constraint algorithm, as discussed above in the index $\nu = 1$ case, to obtain the underlying nondegenerate Hamiltonian dynamics. Restricting to the primary constraint manifold, using the first constraint equation to solve for $u = u(q)$ by the implicit function theorem and subsequently, using the second constraint equation to solve for $\lambda = \lambda(q, p)$ by inverting $\left(\frac{\partial g}{\partial q} \frac{\partial f}{\partial u} \right)^T$, gives the Hamiltonian system

$$\begin{aligned} \dot{q}'^i &= f^i(q', u(q')), \\ \dot{p}'_j &= -p'_i \frac{\partial f^i(q', u(q'))}{\partial q^j} \\ &\quad - \lambda_A(q', p') \left(\frac{\partial^2 g^A(q')}{\partial q^j \partial q^i} f^i(q', u(q')) + \frac{\partial g^A(q')}{\partial q^i} \frac{\partial f^i(q', u(q'))}{\partial q^j} \right). \end{aligned}$$

2.3.2 Adjoint Systems for DAEs with Augmented Hamiltonians

In Sect. 2.2.1, we augmented the adjoint ODE Hamiltonian by some function L . In this section, we do analogously for the adjoint DAE system.

To begin, let $H(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle$ be the Hamiltonian on the generalized phase space bundle corresponding to the DAE $\dot{q} = f(q, u), 0 = \phi(q, u)$,

and let $L : M_d \times M_a \rightarrow \mathbb{R}$ be the function that we would like to augment. We identify L with its pullback through $T^*M_d \oplus \Phi^* \rightarrow M_d \times M_a$. Then, we define the augmented Hamiltonian

$$H_L \equiv H + L : T^*M_d \oplus \Phi^* \rightarrow \mathbb{R}$$

$$(q, u, p, \lambda) \mapsto H(q, u, p, \lambda) + L(q, u).$$

We define the augmented adjoint DAE system as the presymplectic system

$$i_{X_{H_L}} \Omega_0 = dH_L. \tag{2.16}$$

A direct calculation yields the coordinate expression, along an integral curve of such a (generally, partially defined) vector field X_{H_L} ,

$$\dot{q}^i = f^i(q, u), \tag{2.17a}$$

$$\dot{p}_j = -p_i \frac{\partial f^i}{\partial q^j} - \lambda_A \frac{\partial \phi^A}{\partial q^j} - \frac{\partial L}{\partial q^j}, \tag{2.17b}$$

$$0 = \phi^A(q, u), \tag{2.17c}$$

$$0 = p_i \frac{\partial f^i}{\partial u^a} + \lambda_A \frac{\partial \phi^A}{\partial u^a} + \frac{\partial L}{\partial u^a}. \tag{2.17d}$$

Remark 2.21 Observe that if the base DAE (2.12a)–(2.12b) has index 1, then the above system has index 1 by the exact same argument given in the nonaugmented case. After reduction by applying the presymplectic constraint algorithm and solving for u as a function of q and λ as a function of (q, p) , the underlying nondegenerate Hamiltonian system on the primary (final) constraint manifold corresponds to the Hamiltonian

$$(H_L)_1(q', p') = \langle p', f(q', u(q')) \rangle + L(q', u(q')),$$

which is the adjoint Hamiltonian for the ODE $\dot{q}' = f(q', u(q'))$, augmented by $L(q', u(q'))$.

However, as we will discuss in Sect. 3.3, it is not uncommon in optimal control problems for $\partial\phi/\partial u$ to be singular, but the presence of $\int L dt$ in the minimization objective may uniquely specify the singular degrees of freedom.

We now prove an analogous proposition to Proposition 2.11, modified by the presence of L in the Hamiltonian. We again consider the variational equations (2.15a)–(2.15d) associated with the base DAE (2.12a)–(2.12b), which for simplicity we express in matrix derivative notation as

$$\dot{q} = f(q, u), \tag{2.18a}$$

$$0 = \phi(q, u), \tag{2.18b}$$

$$\frac{d}{dt} \delta q = D_q f(q, u) \delta q + D_u f(q, u) \delta u, \tag{2.18c}$$

$$0 = D_q\phi(q, u)\delta q + D_u\phi(q, u)\delta u. \quad (2.18d)$$

Proposition 2.12 For a solution (q, u, p, λ) of the augmented adjoint DAE system (2.17a)–(2.17d) and a solution $(q, u, \delta q, \delta u)$ of the variational equations (2.18a)–(2.18d), covering the same solution (q, u) of the base DAE (2.12a)–(2.12b),

$$\frac{d}{dt}\langle p, \delta q \rangle = -\langle \nabla_q L, \delta q \rangle - \langle \nabla_u L, \delta u \rangle. \quad (2.19)$$

Proof This follows from a direct computation:

$$\begin{aligned} \frac{d}{dt}\langle p, \delta q \rangle &= \langle \dot{p}, \delta q \rangle + \langle p, \frac{d}{dt}\delta q \rangle \\ &= -\langle [D_q f]^* p, \delta q \rangle - \langle [D_q \phi]^* \lambda, \delta q \rangle - \langle \nabla_q L, \delta q \rangle \\ &\quad + \langle p, D_q f \delta q \rangle + \langle p, D_u f \delta u \rangle \\ &= -\langle \lambda, D_q \phi \delta q \rangle - \langle \nabla_q L, \delta q \rangle + \langle p, D_u f \delta u \rangle \\ &= \langle \lambda, D_u \phi \delta u \rangle - \langle \nabla_q L, \delta q \rangle + \langle p, D_u f \delta u \rangle \\ &= -\langle \nabla_q L, \delta q \rangle + \langle [D_u \phi]^* \lambda + [D_u f]^* p, \delta u \rangle \\ &= -\langle \nabla_q L, \delta q \rangle - \langle \nabla_u L, \delta u \rangle, \end{aligned}$$

where in the fourth equality above we used (2.18d) and in the sixth equality above we used (2.17d). \square

Remark 2.22 Analogous to the ODE case discussed in Remark 2.9, we remark that for the nonaugmented adjoint DAE system (2.14a)–(2.14d), we have preservation of $\langle p, \delta q \rangle$ by virtue of presymplecticity. On the other hand, for the augmented adjoint DAE system, despite preserving the same presymplectic form, the change of $\langle p, \delta q \rangle$ now measures the change in L with respect to variations in q and u . This can be understood from the fact that the adjoint equations for (p, λ) in the nonaugmented case, (2.14b) and (2.14d), are linear in (p, λ) , so that one can identify first variations in (p, λ) with (p, λ) , whereas, in the augmented case, equations (2.17b) and (2.17d) are affine in (p, λ) , so such an identification cannot be made. Furthermore, the failure of (2.17b) and (2.17d) to be linear in (p, λ) are given precisely by $\nabla_q L$ and $\nabla_u L$, respectively. Thus, in the augmented case, this leads to the additional terms $-\langle \nabla_u L, \delta q \rangle - \langle \nabla_q L, \delta u \rangle$ in equation (2.19).

3 Applications

3.1 Adjoint Sensitivity Analysis for Semi-Explicit Index 1 DAEs

In this section, we discuss how one can utilize adjoint systems to compute sensitivities. We will split this into four cases; namely, we want to compute sensitivities for ODEs or DAEs (we will focus on index 1 DAEs), and whether we are computing the sensitivity of a terminal cost or the sensitivity of a running cost.

The relevant adjoint system used to compute sensitivities in all four cases are summarized below.

	Terminal Cost	Running Cost
ODE	Adjoint ODE System (2.6a)–(2.6b)	Augmented Adjoint ODE System (2.11a)–(2.11b)
DAE	Adjoint DAE System (2.14a)–(2.14d)	Augmented Adjoint DAE System (2.17a)–(2.17d)

Note that in our calculations below, the top row (the ODE case) can be formally obtained from the bottom row (the DAE case) simply by ignoring the algebraic variables (u, λ) and letting the constraint function ϕ be identically zero. Thus, we will focus on the bottom row, i.e., computing sensitivities of a terminal cost function and of a running cost function, subject to a DAE constraint. In both cases, we will first show how the adjoint sensitivity can be derived using a traditional variational argument. Subsequently, we will show how the adjoint sensitivity can be derived more simply by using Propositions 2.11 and 2.12.

Adjoint Sensitivity of a Terminal Cost Consider the DAE $\dot{q} = f(q, u), 0 = \phi(q, u)$ as in Sect. 2.3. We will assume that M_d is a vector space and additionally, that the DAE has index 1. We would like to extract the gradient of a terminal cost function $C(q(t_f))$ with respect to the initial condition $q(0) = \alpha$, i.e., we want to extract the sensitivity of $C(q(t_f))$ with respect to an infinitesimal perturbation in the initial condition, given by $\nabla_\alpha C(q(t_f))$. Consider the functional J defined by

$$J = C(q(t_f)) - \langle p_0, q(0) - \alpha \rangle - \int_0^{t_f} [\langle p, \dot{q} - f(q, u) \rangle - \langle \lambda, \phi(q, u) \rangle] dt.$$

Observe that for (q, u) satisfying the given DAE with initial condition $q(0) = \alpha$, J coincides with $C(q(t_f))$. We think of p_0 as a free parameter. For simplicity, we will use matrix derivative notation instead of indices. Computing the variation of J yields

$$\begin{aligned} \delta J &= \langle \nabla_q C(q(t_f)), \delta q(t_f) \rangle - \langle p_0, \delta q(0) - \delta \alpha \rangle \\ &\quad - \int_0^{t_f} \left[\langle p, \frac{d}{dt} \delta q - D_q f(q, u) \delta q \rangle - \langle p, D_u f(q, u) \delta u \rangle \right. \\ &\quad \left. - \langle \lambda, D_q \phi(q, u) \delta q + D_u \phi(q, u) \delta u \rangle \right] dt. \end{aligned}$$

Integrating by parts in the term containing $\frac{d}{dt} \delta q$ and restricting to a solution (q, u, p, λ) of the adjoint DAE system (2.14a)–(2.14d) yields

$$\delta J = \langle \nabla_q C(q(t_f)) - p(t_f), \delta q(t_f) \rangle - \langle p_0, \delta \alpha \rangle + \langle p(0) - p_0, \delta q(0) \rangle.$$

We enforce the endpoint condition $p(t_f) = \nabla_q C(q(t_f))$ and choose $p_0 = p(0)$, which yields

$$\delta J = \langle p(0), \delta \alpha \rangle.$$

Hence, the sensitivity of $C(q(t_f))$ is given by

$$p(0) = \nabla_\alpha J = \nabla_\alpha C(q(t_f)),$$

with initial condition $q(0) = \alpha$ and terminal condition $p(t_f) = \nabla_q C(q(t_f))$. Thus, the adjoint sensitivity can be computed by setting the terminal condition on $p(t_f)$ above and subsequently, solving for the momenta p at time 0. In order for this to be well-defined, we have to verify that the given initial and terminal conditions lie on the primary constraint manifold P_1 . However, as discussed in Sect. 2.3.1, since the DAE has index 1, we can always solve for the algebraic variables $u = u(q)$ and $\lambda = \lambda(q, p)$ and thus, we are free to choose the initial and terminal values of q and p , respectively. For higher index DAEs, one has to ensure that these conditions are compatible with the final constraint manifold. For example, this is done in Cao et al. (2003) in the case of Hessenberg index 2 DAEs. Alternatively, at least theoretically, for higher index DAEs, one can reduce the DAE to an index 1 DAE and then the above discussion applies; however, this reduction may fail in practice due to numerical cancellation.

Note that the above adjoint sensitivity result is also a consequence of the preservation of the quadratic invariant $\langle p, v \rangle$ as in Proposition 2.11. From this proposition, one has that

$$\langle p(t_f), \delta q(t_f) \rangle = \langle p(0), \delta q(0) \rangle,$$

where δq satisfies the variational equations. Setting $p(t_f) = \nabla_q C(q(t_f))$ and $\delta q(0) = \delta \alpha$ gives the same result. As mentioned in Remark 2.15, this quadratic invariant arises from the presymplecticity of the adjoint DAE system. Thus, a numerical integrator which preserves the presymplectic structure is desirable for computing adjoint sensitivities, as it exactly preserves the quadratic invariant that allows the adjoint sensitivities to be accurately and efficiently computed. We will discuss this in more detail in Sect. 3.2.

Adjoint Sensitivity of a Running Cost Again, consider an index 1 DAE $\dot{q} = f(q, u)$, $0 = \phi(q, u)$. We would like to extract the sensitivity of a running cost function

$$\int_0^{t_f} L(q, u) dt,$$

where $L : M_d \times M_a \rightarrow \mathbb{R}$, with respect to an infinitesimal perturbation in the initial condition $q(0) = \alpha$. Consider the functional J defined by

$$J = -\langle p_0, q(0) - \alpha \rangle + \int_0^{t_f} [L(q, u) + \langle p, f(q, u) - \dot{q} \rangle + \langle \lambda, \phi(q, u) \rangle] dt.$$

Observe that when the DAE is satisfied with initial condition $q(0) = \alpha$, $J = \int_0^{t_f} L dt$. Now, we would to compute the implicit change in $\int_0^{t_f} L dt$ with respect to a perturbation $\delta \alpha$ in the initial condition. Taking the variation in J yields

$$\delta J = -\langle p_0, \delta q(0) - \delta \alpha \rangle$$

$$\begin{aligned}
 & + \int_0^{t_f} \left[\langle \nabla_q L, \delta q \rangle + \langle \nabla_u L, \delta u \rangle + \left\langle p, D_q f \delta q - \frac{d}{dt} \delta q \right\rangle \right. \\
 & \left. + \langle p, D_u f \delta u \rangle + \langle \lambda, D_q \phi \delta q + D_u \phi \delta u \rangle \right] dt \\
 = & - \langle p_0, \delta q(0) - \delta \alpha \rangle - \langle p(t_f), \delta q(t_f) \rangle + \langle p(0), \delta q(0) \rangle \\
 & + \int_0^{t_f} \left[\langle \nabla_q L + [D_q f]^* p + [D_q \phi]^* \lambda + \dot{p}, \delta q \rangle \right. \\
 & \left. + \langle \nabla_u L + [D_u f]^* p + [D_u \phi]^* \lambda, \delta u \rangle \right] dt.
 \end{aligned}$$

Restricting to a solution (q, u, p, λ) of the augmented adjoint DAE system (2.17a)–(2.17d), setting the terminal condition $p(t_f) = 0$, and choosing $p_0 = p(0)$ gives $\delta J = \langle p(0), \delta \alpha \rangle$. Hence, the implicit sensitivity of $\int_0^{t_f} L dt$ with respect to a change $\delta \alpha$ in the initial condition is given by

$$p(0) = \delta_\alpha J = \delta_\alpha \int_0^{t_f} L(q, u) dt.$$

Thus, the adjoint sensitivity of a running cost functional with respect to a perturbation in the initial condition can be computed by using the augmented adjoint DAE system (2.17a)–(2.17d) with terminal condition $p(t_f) = 0$ to solve for the momenta p at time 0.

Note that the above adjoint sensitivity result can be obtained from Proposition 2.12 as follows. We write equation (2.19) as

$$\frac{d}{dt} \langle p, \delta q \rangle = - \langle dL, (\delta q, \delta u) \rangle,$$

to highlight that the right-hand side measures the total induced variation of L . Now, we integrate this equation from 0 to t_f , which gives

$$\langle p(t_f), \delta q(t_f) \rangle - \langle p(0), \delta q(0) \rangle = - \int_0^{t_f} \langle dL, (\delta q, \delta u) \rangle dt.$$

Since we want to determine the change in the running cost functional with respect to a perturbation in the initial condition, we set $p(t_f) = 0$ which yields

$$\langle p(0), \delta q(0) \rangle = \int_0^{t_f} \langle dL, (\delta q, \delta u) \rangle dt.$$

The right-hand side is the total change induced on the running cost functional, whereas the left-hand side tells us how this change is implicitly induced from a perturbation $\delta q(0)$ in the initial condition. Note that a perturbation in the initial condition $\delta q(0)$ will generally induce perturbations in both q and u , according to the variational equations. Such a curve $(\delta q, \delta u)$ satisfying the variational equations exists in the index 1 case

as noted in Remark 2.16. Thus, we arrive at the same conclusion as the variational argument: $p(0)$ is the desired adjoint sensitivity.

To summarize, adjoint sensitivities for terminal and running costs can be computed using the properties of adjoint systems, such as the various aforementioned propositions regarding $\frac{d}{dt} \langle p, \delta q \rangle$, which is zero in the nonaugmented case and measures the variation of L in the augmented case. In the case of a terminal cost, one sets an inhomogeneous terminal condition $p(t_f) = \nabla_q C(q(t_f))$ and backpropagates the momenta through the nonaugmented adjoint DAE system (2.14a)–(2.14d) to obtain the sensitivity $p(0)$. On the other hand, in the case of a running cost, one sets a homogeneous terminal condition $p(t_f) = 0$ and backpropagates the momenta through the augmented adjoint DAE system (2.17a)–(2.17d) to obtain the sensitivity $p(0)$.

The various propositions used to derive the above adjoint sensitivity results are summarized below. We also include the ODE case, since it follows similarly. In

	Terminal Cost	Running Cost
ODE	Proposition 2.3, $\frac{d}{dt} \langle p, \delta q \rangle = 0$	Proposition 2.7, $\frac{d}{dt} \langle p, \delta q \rangle = -\langle dL, \delta q \rangle$
DAE	Proposition 2.11, $\frac{d}{dt} \langle p, \delta q \rangle = 0$	Proposition 2.12, $\frac{d}{dt} \langle p, \delta q \rangle = -\langle dL, (\delta q, \delta u) \rangle$

Sect. 3.2, we will construct integrators that admit discrete analogues of the above propositions, and hence, are suitable for computing discrete adjoint sensitivities.

3.2 Structure-Preserving Discretizations of Adjoint Systems

In this section, we utilize the Galerkin Hamiltonian variational integrators of Leok and Zhang (2011) to construct structure-preserving integrators which admit discrete analogues of Propositions 2.3, 2.7, 2.11, and 2.12, and are therefore suitable for numerical adjoint sensitivity analysis. For brevity, the proofs of these discrete analogues can be found in Appendix A.

We start by recalling the construction of Galerkin Hamiltonian variational integrators as introduced in Leok and Zhang (2011). We assume that the base manifold Q is a vector space, and thus, we have the identification $T^*Q \cong Q \times Q^*$. To construct a variational integrator for a Hamiltonian system on T^*Q , one starts with the exact Type II generating function

$$H_{d,\text{exact}}^+(q_0, p_1) = \text{ext} \left[\langle p_1, q_1 \rangle - \int_0^{\Delta t} [\langle p, \dot{q} \rangle - H(q, p)] dt \right],$$

where one extremizes over C^2 curves on the cotangent bundle satisfying $q(0) = q_0, p(\Delta t) = p_1$. This is a Type II generating function in the sense that it defines a symplectic map $(q_0, p_1) \mapsto (q_1, p_0)$ by $q_1 = D_2 H_{d,\text{exact}}^+(q_0, p_1), p_0 = D_1 H_{d,\text{exact}}^+(q_0, p_1)$.

To approximate this generating function, one approximates the integral above using a quadrature rule and extremizes the resulting expression over a finite-dimensional subspace satisfying the prescribed boundary conditions. This yields the Galerkin discrete Hamiltonian

$$H_d^+(q_0, p_1) = \text{ext} \left[\langle p_1, q_1 \rangle - \Delta t \sum_i b_i \left(\langle P^i, V^i \rangle - H(Q^i, P^i) \right) \right],$$

where $\Delta t > 0$ is the timestep, q_0, q_1, p_0, p_1 are numerical approximations to $q(0), q(\Delta t), p(0), p(\Delta t)$, respectively, $b_i > 0$ are quadrature weights corresponding to quadrature nodes $c_i \in [0, 1]$, Q^i and P^i are internal stages representing $q(c_i \Delta t), p(c_i \Delta t)$, respectively, and V is related to Q by $Q^i = q_0 + \Delta t \sum_j a_{ij} V^j$, where the coefficients a_{ij} arise from the choice of function space. The expression above is extremized over the internal stages Q^i, P^i and subsequently, one applies the discrete right Hamilton’s equations

$$\begin{aligned} q_1 &= D_2 H_d^+(q_0, p_1), \\ p_0 &= D_1 H_d^+(q_0, p_1), \end{aligned}$$

to obtain a Galerkin Hamiltonian variational integrator. The extremization conditions and the discrete right Hamilton’s equations can be expressed as

$$q_1 = q_0 + \Delta t \sum_i b_i D_p H(Q^i, P^i), \tag{3.1a}$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} D_p H(Q^j, P^j), \tag{3.1b}$$

$$p_1 = p_0 - \Delta t \sum_i b_i D_q H(Q^i, P^i), \tag{3.1c}$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} D_q H(Q^j, P^j), \tag{3.1d}$$

where we interpret a_{ij} as Runge–Kutta coefficients and $\tilde{a}_{ij} = (b_i b_j - b_j a_{ji})/b_i$ as the symplectic adjoint of the a_{ij} coefficients. Thus, (3.1a)–(3.1d) can be viewed as a symplectic partitioned Runge–Kutta method.

We will consider such methods in four cases: adjoint systems corresponding to a base ODE or DAE, and whether or not the corresponding system is augmented. Note that in the DAE case, we will have to modify the above construction because the system is presymplectic. Furthermore, we will assume that all of the relevant configuration spaces are vector spaces.

Nonaugmented Adjoint ODE System The simplest case to consider is the nonaugmented adjoint ODE system (2.6a)–(2.6b). Since the quadratic conservation law in Proposition 2.3,

$$\frac{d}{dt} \langle p, \delta q \rangle = 0,$$

arises from symplecticity, a structure-preserving discretization can be obtained by applying a symplectic integrator. This case is already discussed in Sanz-Serna (2016), so we will only outline it briefly.

Applying the Galerkin Hamiltonian variational integrator (3.1a)–(3.1d) to the Hamiltonian for the adjoint ODE system, $H(q, p) = \langle p, f(q) \rangle$, yields

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i), \quad (3.2a)$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j), \quad (3.2b)$$

$$p_1 = p_0 - \Delta t \sum_i b_i [Df(Q^i)]^* P^i, \quad (3.2c)$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} [Df(Q^j)]^* P^j. \quad (3.2d)$$

In the setting of adjoint sensitivity analysis of a terminal cost function, the appropriate boundary condition to prescribe on the momenta is $p_1 = \nabla_q C(q(t_f))$, as discussed in Sect. 3.1.

Since the above integrator is symplectic, we have the symplectic conservation law,

$$dq_1 \wedge dp_1 = dq_0 \wedge dp_0,$$

when evaluated on discrete first variations of (3.2a)–(3.2d). In this setting, a discrete first variation can be identified with solutions of the linearization of (3.2a)–(3.2d). For the linearization of the equations in the position variables, (3.2a)–(3.2b), we have

$$\delta q_1 = \delta q_0 + \Delta t \sum_i b_i Df(Q^i) \delta Q^i, \quad (3.3a)$$

$$\delta Q^i = \delta q_0 + \Delta t \sum_j a_{ij} Df(Q^j) \delta Q^j. \quad (3.3b)$$

As observed in Sanz-Serna (2016), while we obtained this by linearizing the discrete equations, one could also obtain this by first linearizing (2.1) and subsequently, applying the Runge–Kutta scheme to the linearization. For the linearization of the equations for the adjoint variables, (3.2c)–(3.2d), observe that they are already linear in the adjoint variables, so we can identify the linearization with itself. Thus, we can choose for first variations vector fields V as the first variation corresponding to the solution of the linearized position equation and W as the first variation corresponding to the solution of the adjoint equation itself. With these choices, the above symplectic conservation law yields

$$0 = dq_1 \wedge dp_1(V, W)|_{(q_1, p_1)} - dq_0 \wedge dp_0(V, W)|_{(q_0, p_0)} = \langle p_1, \delta q_1 \rangle - \langle p_0, \delta q_0 \rangle.$$

This is of course a discrete analogue of Proposition 2.3. Note that one can derive the conservation law $\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle$ directly by starting with the expression

$\langle p_1, \delta q_1 \rangle$ and substituting the discrete equations where appropriate. We will do this in the more general augmented case below.

Augmented Adjoint ODE System We now consider the case of the augmented adjoint ODE system (2.11a)–(2.11b). In the continuous setting, we have from Proposition 2.7,

$$\frac{d}{dt} \langle p, \delta q \rangle = -\langle dL, \delta q \rangle.$$

We would like to construct an integrator which admits a discrete analogue of this equation. To do this, we apply the Galerkin Hamiltonian variational integrator, equations (3.1a)–(3.1d), to the augmented Hamiltonian $H_L(q, p) = \langle p, f(q) \rangle + L(q)$. This gives

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i), \tag{3.4a}$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j), \tag{3.4b}$$

$$p_1 = p_0 - \Delta t \sum_i b_i ([Df(Q^i)]^* P^i + dL(Q^i)), \tag{3.4c}$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} ([Df(Q^j)]^* P^j + dL(Q^j)). \tag{3.4d}$$

We now prove a discrete analogue of Proposition 2.7. To do this, we again consider the discrete variational equations for the position variables, (3.3a)–(3.3b).

Proposition 3.1 *With the above notation, the above integrator satisfies*

$$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta Q^i \rangle. \tag{3.5}$$

Proof See Appendix A. □

Remark 3.1 To see that this is a discrete analogue of $\frac{d}{dt} \langle p, \delta q \rangle = -\langle dL, \delta q \rangle$, we write it in integral form as

$$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle - \int_0^{\Delta t} \langle dL(q), \delta q \rangle dt.$$

Then, applying the quadrature rule on $[0, \Delta t]$ given by quadrature weights $b_i \Delta t$ and quadrature nodes $c_i \Delta t$, the above integral is approximated by

$$\int_0^{\Delta t} \langle dL(q), \delta q \rangle dt \approx \Delta t \sum_i b_i \langle dL(q(c_i \Delta t)), \delta q(c_i \Delta t) \rangle = \Delta t \sum_i b_i \langle dL(Q^i), \delta Q^i \rangle,$$

which yields equation (3.5). The discrete analogue is natural in the sense that the quadrature rule for which the discrete equation (3.5) approximates the continuous equation is the same as the quadrature rule used to approximate the exact discrete generating function. This occurs more generally for such Hamiltonian variational integrators, as noted in Tran and Leok (2022) for the more general setting of multisymplectic Hamiltonian variational integrators.

For adjoint sensitivity analysis of a running cost $\int L dt$, the appropriate boundary condition to prescribe on the momenta is $p_1 = 0$, as discussed in Sect. 3.1. With such a boundary condition, equation (3.5) reduces to

$$\langle p_0, \delta q_0 \rangle = \Delta t \sum_i b_i \langle dL(Q^i), \delta Q^i \rangle.$$

Thus, p_0 gives the discrete sensitivity, i.e., the change in the quadrature approximation of $\int L dt$ induced by a change in the initial condition along a discrete solution trajectory. One can compute this quantity directly via the direct method, where one needs to integrate the discrete variational equations for every desired search direction δq_0 . On the other hand, by the above proposition, one can compute this quantity using the adjoint method: one integrates the adjoint equation with $p_1 = 0$ once to compute p_0 and subsequently, pair p_0 with any search direction δq_0 to obtain the sensitivity in that direction. By the above proposition, both methods give the same sensitivities. However, assuming the search space has dimension $n > 1$, the adjoint method is more efficient since it only requires $\mathcal{O}(1)$ integrations and $\mathcal{O}(n)$ vector–vector products, whereas the direct method requires $\mathcal{O}(n)$ integrations and $\mathcal{O}(ns)$ vector–vector products where $s \geq 1$ is the number of Runge–Kutta stages, since, in the direct method, one has to compute $\langle dL(Q^i), \delta Q^i \rangle$ for each i and for each choice of δq_0 .

Nonaugmented Adjoint DAE System We will now construct discrete Hamiltonian variational integrators for the adjoint DAE system (2.14a)–(2.14d), where we assume that the base DAE has index 1. To construct such a method, we have to modify the Galerkin Hamiltonian variational integrator (3.1a)–(3.1d), so that it is applicable to the presymplectic adjoint DAE system.

First, consider a general presymplectic system $i_X \Omega' = dH$. Note that, locally, any presymplectic system can be transformed to the canonical form (see, Cariñena et al. 1987),

$$\begin{aligned} \dot{q} &= D_p H(q, p, r), \\ \dot{p} &= -D_q H(q, p, r), \\ 0 &= D_r H(q, p, r), \end{aligned}$$

where, in these coordinates, $\Omega' = dq \wedge dp$, so that $\ker(\Omega') = \text{span}\{\partial/\partial r\}$. The action for this system is given by $\int_0^{\Delta t} (\langle p, \dot{q} \rangle - H(q, p, r)) dt$. We approximate this integral by quadrature, introduce internal stages for q, p as before, and additionally introduce

internal stages $R^i = r(c_i h)$. This gives the discrete generating function

$$H_d^+(q_0, p_1) = \text{ext} \left[\langle p_1, q_1 \rangle - \Delta t \sum_i b_i \left(\langle P^i, V^i \rangle - H(Q^i, P^i, R^i) \right) \right],$$

where again V is related to the internal stages of Q by $Q^i = q_0 + \Delta t \sum_j a_{ij} V^j$ and the above expression is extremized over the internal stages Q^i, P^i, R^i . The discrete right Hamilton's equations are again given by

$$q_1 = H_d^+(q_0, p_1), \quad p_0 = H_d^+(q_0, p_1),$$

which we interpret as the evolution equations of the system. There are no evolution equations for r due to the presymplectic structure and the absence of derivatives of r in the action. This gives the integrator

$$q_1 = q_0 + \Delta t \sum_i b_i D_p H(Q^i, P^i, R^i), \quad (3.6a)$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} D_p H(Q^i, P^i, R^i), \quad (3.6b)$$

$$p_1 = p_0 - \Delta t \sum_i b_i D_q H(Q^i, P^i, R^i), \quad (3.6c)$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} D_q H(Q^i, P^i, R^i), \quad (3.6d)$$

$$0 = D_r H(Q^i, P^i, R^i), \quad (3.6e)$$

where (3.6b), (3.6d), (3.6e) arise from extremizing with respect to P^i, Q^i, R^i , respectively, while (3.6a) and (3.6c) arise from the discrete right Hamilton's equations. This integrator is presymplectic, in the sense that

$$dq_1 \wedge dp_1 = dq_0 \wedge dp_0,$$

when evaluated on discrete first variations. The proof is formally identical to the symplectic case. For this reason, we refer to (3.6a)–(3.6e) as a presymplectic Galerkin Hamiltonian variational integrator.

Remark 3.2 In general, the system (3.6a)–(3.6e) evolves on the primary constraint manifold given implicitly by the zero level set of $D_r H$, however, it may not evolve on the final constraint manifold. This is not an issue for us since we are dealing with adjoint DAE systems for index 1 DAEs, for which we know the primary constraint manifold and the final constraint manifold coincide. For the general case, one may need to additionally differentiate the constraint equation $D_r H = 0$ to obtain hidden constraints.

Thus, the method (3.6a)–(3.6e) is generally only applicable to index 1 presymplectic systems, unless we add in further hidden constraints. In order for the continuous

presymplectic system to have index 1, it is sufficient that the Hessian of H with respect to the algebraic variables, $D_r^2 H$, is (pointwise) invertible on the primary constraint manifold. This is the case for the adjoint DAE system corresponding to an index 1 DAE.

We now specialize to the adjoint DAE system (2.14a)–(2.14d), corresponding to an index 1 DAE, which is already in the above canonical form with $r = (u, \lambda)$ and $H(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle$. Note that we reordered the argument of H , $(q, p, r) = (q, p, u, \lambda) \rightarrow (q, u, p, \lambda)$, in order to be consistent with the previous notation used throughout. We label the internal stages for the algebraic variables as $R^i = (U^i, \Lambda^i)$. Applying the presymplectic Galerkin Hamiltonian variational integrator to this particular system yields

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i, U^i), \tag{3.7a}$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j, U^j), \tag{3.7b}$$

$$p_1 = p_0 - \Delta t \sum_i b_i \left([D_q f(Q^i, U^i)]^* P^i + [D_q \phi(Q^i, U^i)]^* \Lambda^i \right), \tag{3.7c}$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} \left([D_q f(Q^j, U^j)]^* P^j + [D_q \phi(Q^j, U^j)]^* \Lambda^j \right), \tag{3.7d}$$

$$0 = \phi(Q^i, U^i), \tag{3.7e}$$

$$0 = [D_u f(Q^i, U^i)]^* P^i + [D_u \phi(Q^i, U^i)]^* \Lambda^i, \tag{3.7f}$$

where (3.7b), (3.7d), (3.7e), (3.7f) arise from extremizing over P^i, Q^i, Λ^i, U^i , respectively, while (3.7a), (3.7c) arise from the discrete right Hamilton’s equations.

Remark 3.3 In order for q_1 to appropriately satisfy the constraint, we should take the final quadrature point to be $c_s = 1$ (for an s -stage method), so that $\phi(q_1, U^s) = \phi(Q^s, U^s) = 0$. In this case, equation (3.7a) and equation (3.7b) with $i = s$ are redundant. Note that with the choice $c_s = 1$, they are still consistent (i.e., are the same equation), since in the Galerkin construction, the coefficients a_{ij} and b_i are defined as

$$a_{ij} = \int_0^{c_i} \phi_j(\tau) d\tau, \quad b_i = \int_0^1 \phi_j(\tau) d\tau,$$

where ϕ_j are functions on $[0, 1]$ which interpolate the nodes c_j (see, Leok and Zhang 2011). Hence, $a_{sj} = b_j$, so that the two equations are consistent. However, we will write the system as above for conceptual clarity. Furthermore, even in the case where one does not take $c_s = 1$, the proposition that we prove below still holds, despite the possibility of constraint violations.

A similar remark holds for the adjoint variable p and the associated constraint (3.7f), except we think of p_0 as the unknown, instead of p_1 .

Note that (3.7a), (3.7b), (3.7e) is a standard Runge–Kutta discretization of an index 1 DAE $\dot{q} = f(q, u)$, $0 = \phi(q, u)$, where again, usually $c_s = 1$. Associated with these equations are the variational equations given by their linearization,

$$\delta q_1 = \delta q_0 + \Delta t \sum_i b_i (D_q f(Q^i, U^i) \delta Q^i + D_u f(Q^i, U^i) \delta U^i), \quad (3.8a)$$

$$\delta Q^i = \delta q_0 + \Delta t \sum_j a_{ij} (D_q f(Q^j, U^j) \delta Q^j + D_u f(Q^j, U^j) \delta U^j), \quad (3.8b)$$

$$0 = D_q \phi(Q^i, U^i) \delta Q^i + D_u \phi(Q^i, U^i) \delta U^i, \quad (3.8c)$$

which is the Runge–Kutta discretization of the continuous variational equations (2.15c)–(2.15d).

Proposition 3.2 *With the above notation, the above integrator satisfies*

$$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle.$$

Proof See Appendix A. □

Thus, the above integrator admits a discrete analogue of Proposition 2.11 for the nonaugmented adjoint DAE system. By setting $p_1 = \nabla_q C(q(t_f))$, one can use this integrator to compute the sensitivity p_0 of a terminal cost function with respect to a perturbation in the initial condition. As discussed before, this only requires $\mathcal{O}(1)$ integrations instead of $\mathcal{O}(n)$ integrations via the direct method (for a dimension n search space). Furthermore, the adjoint method requires only $\mathcal{O}(1)$ numerical solves of the constraints, while the direct method requires $\mathcal{O}(n)$ numerical solves.

Remark 3.4 Since we are assuming the DAE has index 1, it is always possible to prescribe an arbitrary initial condition q_0 (and δq_0) and terminal condition p_1 , since the corresponding algebraic variables can always formally be solved for using the corresponding constraints. In practice, one generally has to solve the constraints to some tolerance, e.g., through an iterative scheme. If the constraints are only satisfied to a tolerance $\mathcal{O}(\epsilon)$, then the above proposition holds to $\mathcal{O}(s\epsilon)$, where s is the number of Runge–Kutta stages.

Remark 3.5 The above method (3.7a)–(3.7f) is presymplectic, since it is a special case of the more general presymplectic Galerkin Hamiltonian variational integrator (3.6a)–(3.6e). Although we proved it directly, the above proposition could also have been proven from presymplecticity, with the appropriate choices of first variations.

Augmented Adjoint DAE System Finally, we construct a discrete Hamiltonian variational integrator for the augmented adjoint DAE system (2.17a)–(2.17d) associated with an index 1 DAE. To do this, we apply the presymplectic Galerkin Hamiltonian variational integrator (3.6a)–(3.6e) with $r = (u, \lambda)$ and with Hamiltonian given by the augmented adjoint DAE Hamiltonian,

$$H_L(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle + L(q, u).$$

The presymplectic integrator is then

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i, U^i), \tag{3.9a}$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j, U^j), \tag{3.9b}$$

$$p_1 = p_0 - \Delta t \sum_i b_i \left([D_q f(Q^i, U^i)]^* P^i + [D_q \phi(Q^i, U^i)]^* \Lambda^i + D_q L(Q^i, U^i) \right), \tag{3.9c}$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} \left([D_q f(Q^j, U^j)]^* P^j + [D_q \phi(Q^j, U^j)]^* \Lambda^j + D_q L(Q^j, U^j) \right), \tag{3.9d}$$

$$0 = \phi(Q^i, U^i), \tag{3.9e}$$

$$0 = [D_u f(Q^i, U^i)]^* P^i + [D_u \phi(Q^i, U^i)]^* \Lambda^i + D_u L(Q^i, U^i). \tag{3.9f}$$

The associated variational equations are again (3.8a)–(3.8c). Remarks analogous to the nonaugmented case regarding setting the quadrature node $c_s = 1$ and solvability of these systems under the index 1 assumption can be made.

Proposition 3.3 *With the above notation, the above integrator satisfies*

$$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i, U^i), (\delta Q^i, \delta U^i) \rangle.$$

Proof See Appendix A. □

Remark 3.6 Analogous to the remark in the augmented adjoint ODE case, the above proposition is a discrete analogue of Proposition 2.12, in integral form,

$$\langle p_1, \delta q_1 \rangle - \langle p_0, \delta q_0 \rangle = - \int_0^{\Delta t} \langle dL(q, u), (\delta q, \delta u) \rangle dt.$$

The discrete analogue is natural in the sense that it is just quadrature applied to the right-hand side of this equation, with the same quadrature rule used to discretize the generating function.

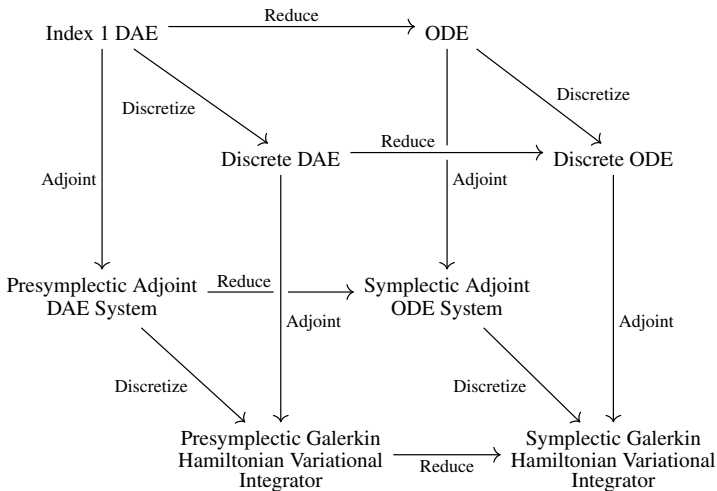
Remark 3.7 As with the augmented adjoint ODE case, the above proposition allows one to compute numerical sensitivities of a running cost function by solving for p_0 with $p_1 = 0$, which is more efficient than the direct method.

To summarize, we have utilized Galerkin Hamiltonian variational integrators to construct methods which admit natural discrete analogues of the various propositions used for sensitivity analysis. We summarize the results below.

	Terminal Cost	Running Cost
ODE	$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle$	$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta Q^i \rangle$
DAE	$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle$	$\langle p_1, \delta q_1 \rangle = \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i, U^i), (\delta Q^i, \delta U^i) \rangle$

3.2.1 Naturality of the Adjoint DAE System Discretization

To conclude our discussion of discretizing adjoint systems, we prove a discrete extension of the fact that, for an index 1 DAE, the process of index reduction and forming the adjoint system commute, as discussed in Sect. 2.3.1. Namely, we will show that, starting from an index 1 DAE (2.12a)–(2.12b), the processes of reduction, forming the adjoint system, and discretization all commute, for particular choices of these processes which we will define and choose below. This can be summarized in the following commutative diagram.



In the above diagram, we will use the convention that the “Discretize” arrows point forward, the “Adjoint” arrows point downward, and the “Reduce” arrows point to the right. For the “Discretize” arrows on the top face, we take the discretization to be a Runge–Kutta discretization (of a DAE on the left and of an ODE on the right, with the same Runge–Kutta coefficients in both cases). For the “Discretize” arrows on the bottom face, we take the discretization to be the symplectic partitioned Runge–Kutta discretization induced by the discretization of the base DAE or ODE, i.e., the momenta expansion coefficients \tilde{a}_{ij} are the symplectic adjoint of the coefficients a_{ij} used on the top face. We have already defined the “Adjoint” arrows on the back face, as discussed in Sect. 2. For the “Adjoint” arrows on the front face, we define them as forming the discrete adjoint system corresponding to a discrete (and generally nonlinear) system of equations and we will review this notion where needed in the proof. We have already

defined the “Reduce” arrows on the back face, as discussed in Sect. 2.3.1. For the “Reduce” arrows on the front face, we define this as solving for the discrete algebraic variables in terms of the discrete kinematic variables through the discrete constraint equations. With these choices, the above diagram commutes, as we will show. To prove this, it suffices to prove that the diagrams on each of the six faces commutes. To keep the exposition concise, we provide the proof in Appendix B and move on to discuss the implications of this result.

The previous discussion shows that the presymplectic Galerkin Hamiltonian variational integrator construction is natural for discretizing adjoint (index 1) DAE systems, in the sense that the integrator is equivalent to the integrator produced from applying a symplectic Galerkin Hamiltonian variational integrator to the underlying nondegenerate Hamiltonian system. Of course, in practice, one cannot generally determine the function $u = u(q)$ needed to reduce the DAE to an ODE. Therefore, one generally works with the presymplectic Galerkin Hamiltonian variational integrator instead, where one iteratively solves the constraint equations. However, although reduction then symplectic integration is often impractical, one can utilize this naturality to derive properties of the presymplectic integrator. For example, we will use this naturality to prove a variational error analysis result.

The basic idea for the variational error analysis result goes as follows: one utilizes the naturality to relate the presymplectic variational integrator to a symplectic variational integrator of the underlying nondegenerate Hamiltonian system and subsequently, applies the variational error analysis result in the symplectic case (Schmitt and Leok 2017). Recall the discrete generating function for the previously constructed presymplectic variational integrator,

$$H_d^+(q_0, p_1; \Delta t) = \text{ext} \left[\langle p_1, q_1 \rangle - \Delta t \sum_i b_i \left(\langle P^i, V^i \rangle - H(Q^i, U^i, P^i, \Lambda^i) \right) \right],$$

where we have now explicitly included the timestep dependence in H_d^+ and H is the Hamiltonian for the adjoint DAE system (augmented or nonaugmented), corresponding to an index 1 DAE.

Proposition 3.4 *Suppose the discrete generating function $H_d^+(q_0, p_1; \Delta t)$ for the presymplectic variational integrator approximates the exact discrete generating function $H_d^{+,E}(q_0, p_1; \Delta t)$ to order r , i.e.,*

$$H_d^+(q_0, p_1; \Delta t) = H_d^{+,E}(q_0, p_1; \Delta t) + \mathcal{O}(\Delta t^{r+1}),$$

and the Hamiltonian H is continuously differentiable, then the Type II map $(q_0, p_1) \mapsto (q_1, p_0)$ and the evolution map $(q_0, p_0) \mapsto (q_1, p_1)$ are order- r accurate.

Proof The proof follows from two simple steps. First, observe that the discrete generating function $H_d^+(q_0, p_1; \Delta t)$ for the presymplectic integrator is also the discrete generating function for the symplectic integrator for the underlying nondegenerate Hamiltonian system. This follows since in the definition of H_d^+ , one extremizes over the algebraic variables U^i, Λ^i which enforces the constraints and hence, determines

U^i, Λ^i as functions of the kinematic variables Q^i, P^i . Thus, the discrete (or continuous) Type II map determined by H_d^+ (or $H_d^{+,E}$, respectively), $(q_0, p_1) \mapsto (q_1, p_0)$, is the same as the Type II map for the underlying nondegenerate Hamiltonian system, which is just another consequence of the aforementioned naturality. One then applies the variational error analysis result in Schmitt and Leok (2017). \square

Remark 3.8 Another way to view this result is that the order of an implicit (partitioned) Runge–Kutta scheme for index 1 DAEs is the same as the order of an implicit (partitioned) Runge–Kutta scheme for ODEs (Roche 1989), since the aforementioned discretization generates a partitioned Runge–Kutta scheme. To be complete, we should determine the order for the full presymplectic flow, i.e., including also the algebraic variables. As discussed in Roche (1989), as long as $a_{si} = b_i$ for each i , which, as we have discussed, is a natural choice and holds as long as $c_s = 1$, there is no order reduction arising from the algebraic variables. Thus, with this assumption, the presymplectic variational integrator in the previous proposition approximates the presymplectic flow, in both the kinematic and algebraic variables, to order r .

Remark 3.9 In the above proposition, we considered both the Type II map $(q_0, p_1) \mapsto (q_1, p_0)$ and the evolution map $(q_0, p_0) \mapsto (q_1, p_1)$. The latter is of course the traditional way to view the map corresponding to a numerical method, but the former is the form of the map used in adjoint sensitivity analysis.

Furthermore, in light of this naturality, we can view Propositions 3.2 and 3.3 as following from the analogous propositions for symplectic Galerkin Hamiltonian variational integrators, applied to the underlying nondegenerate Hamiltonian system.

3.2.2 Numerical Example

For our numerical example, we consider the planar pendulum. Although one can formulate this system as an ODE in the angular variable θ , we instead work with this system in Cartesian coordinates xy where this system is formulated as a DAE, as an academic example of the theory presented in this paper. We will derive the adjoint DAE system associated to the planar pendulum DAE and, subsequently, perform a numerical test demonstrating the presymplecticity of a presymplectic Galerkin Hamiltonian variational integrator applied to this system.

Consider a pendulum of mass $m > 0$ and length $L > 0$ confined to the xy plane, where gravity acts in the vertical y direction, with acceleration $-g < 0$. This is described by the system

$$\begin{aligned} m\ddot{x} &= \rho x, \\ m\ddot{y} &= \rho y - mg, \\ x^2 + y^2 &= L^2. \end{aligned}$$

This system can be derived from the Lagrangian

$$L = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2) - mg(y - L) + \frac{1}{2}\rho(x^2 + y^2 - L^2),$$

where the first term is the kinetic energy, the second term is (minus) the potential energy, and the third term enforces the constraint $x^2 + y^2 = L^2$ where ρ is interpreted as a Lagrange multiplier.

If we restrict to the region $y < 0$, the above system can be expressed as a semi-explicit index 1 DAE of the form

$$\dot{x} = v_x, \tag{3.10a}$$

$$\dot{v}_x = \rho x/m, \tag{3.10b}$$

$$0 = x^2 + y^2 - L^2, \tag{3.10c}$$

$$0 = v_x x + v_y y, \tag{3.10d}$$

$$0 = m(v_x^2 + v_y^2) - mgy + L^2 \rho. \tag{3.10e}$$

In terms of the notation of Sect. 2.3, we have $(x, v_x) \in M_d = (-1, 1) \times \mathbb{R}$ and $(y, v_y, \rho) \in M_a = \mathbb{R}_- \times \mathbb{R} \times \mathbb{R}$. Letting $q = (x, v_x)$ denote the coordinates for the dynamical variables and $u = (y, v_y, \rho)$ denote the coordinates for the algebraic variables, this system can be expressed in the form (2.12a)–(2.12b), where

$$f(q, u) = \begin{pmatrix} v_x \\ \rho x/m \end{pmatrix},$$

$$\phi(q, u) = \begin{pmatrix} x^2 + y^2 - L^2 \\ v_x x + v_y y \\ m(v_x^2 + v_y^2) - mgy + L^2 \rho \end{pmatrix}.$$

We regard ϕ as a section of the constraint bundle Φ given by the trivial vector bundle $(M_d \times M_a) \times \mathbb{R}^3 \rightarrow M_d \times M_a$. Coordinatize $\overline{T^*M_d}$ by (q, u, p) where $p = (p_x, p_{v_x})$ are the momenta dual to $q = (x, v_x)$ and coordinatize Φ^* by (q, u, λ) where $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ are the coordinates of the fibers dual to the constraint bundle fibers. The Hamiltonian $H : \overline{T^*M_d} \oplus \Phi^* \rightarrow \mathbb{R}$ is then given by

$$H(q, u, p, \lambda) = \langle p, f(q, u) \rangle + \langle \lambda, \phi(q, u) \rangle$$

$$= (p_x \ p_{v_x}) \begin{pmatrix} v_x \\ \rho x \end{pmatrix} + (\lambda_1 \ \lambda_2 \ \lambda_3) \begin{pmatrix} x^2 + y^2 - L^2 \\ v_x x + v_y y \\ m(v_x^2 + v_y^2) - mgy + L^2 \rho \end{pmatrix}.$$

The presymplectic form Ω_0 on $\overline{T^*M_d} \oplus \Phi^*$ is given by

$$\Omega_0 = dq \wedge dp = dx \wedge dp_x + dv_x \wedge dp_{v_x}.$$

To obtain an expression for the adjoint DAE system (2.14a)–(2.14d), we compute the derivative matrices of f and ϕ .

$$D_q f(q, u) = \begin{pmatrix} 0 & 1 \\ \rho/m & 0 \end{pmatrix},$$

$$\begin{aligned}
 D_u f(q, u) &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & x/m \end{pmatrix}, \\
 D_q \phi(q, u) &= \begin{pmatrix} 2x & 0 \\ v_x & x \\ 0 & 2mv_x \end{pmatrix}, \\
 D_u \phi(q, u) &= \begin{pmatrix} 2y & 0 & 0 \\ v_y & y & 0 \\ -mg & 2mv_y & L^2 \end{pmatrix}.
 \end{aligned}$$

Note that $\det(D_u \phi(q, u)) = 2L^2 y^2 \neq 0$ for $(q, u) \in M_d \times M_a$, and hence, the system is an index 1 DAE as previously claimed.

The adjoint DAE system (2.14a)–(2.14d) for the planar pendulum is then given by

$$\frac{d}{dt} \begin{pmatrix} x \\ v_x \end{pmatrix} = \begin{pmatrix} v_x \\ \rho x/m \end{pmatrix}, \tag{3.11a}$$

$$\frac{d}{dt} \begin{pmatrix} p_x \\ p_{v_x} \end{pmatrix} = - \begin{pmatrix} 0 & 1 \\ \rho/m & 0 \end{pmatrix}^T \begin{pmatrix} p_x \\ p_{v_x} \end{pmatrix} - \begin{pmatrix} 2x & 0 \\ v_x & x \\ 0 & 2mv_x \end{pmatrix}^T \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix}, \tag{3.11b}$$

$$0 = \begin{pmatrix} x^2 + y^2 - L^2 \\ v_x x + v_y y \\ m(v_x^2 + v_y^2) - mgy + L^2 \rho \end{pmatrix}, \tag{3.11c}$$

$$0 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & x/m \end{pmatrix}^T \begin{pmatrix} p_x \\ p_{v_x} \end{pmatrix} + \begin{pmatrix} 2y & 0 & 0 \\ v_y & y & 0 \\ -mg & 2mv_y & L^2 \end{pmatrix}^T \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{pmatrix}. \tag{3.11d}$$

We will apply a presymplectic Galerkin Hamiltonian variational integrator (3.7a)–(3.7f) to the above system. We choose a first-order Runge–Kutta method, with Runge–Kutta coefficients $a = 1, b = 1, c = 1$ and hence, $\tilde{a} = 0$. Thus, the internal stages for the position and momenta are given by $Q = q_1$ and $P = p_0$. With these choices, the presymplectic Galerkin Hamiltonian variational integrator can be expressed as

$$\begin{aligned}
 q_1 &= q_0 + \Delta t f(q_1, U), \\
 p_1 &= p_0 - \Delta t ([D_q f(q_1, U)]^* p_0 + [D_q \phi(q_1, U)]^* \Lambda), \\
 0 &= \phi(q_1, U), \\
 0 &= [D_u f(q_1, U)]^* p_0 + [D_u \phi(q_1, U)]^* \Lambda.
 \end{aligned}$$

For our example, we set $m = g = L = 1$. Letting $U = (Y, V_y, \mathcal{P})$ and $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$ denote the internal stages corresponding to $u = (y, v_y, \rho)$ and $\lambda = (\lambda_1, \lambda_2, \lambda_3)$, respectively, the above integrator applied to the adjoint DAE system for the planar pendulum (3.11a)–(3.11d), with $m = g = L = 1$, can be expressed as

$$\begin{pmatrix} x_1 \\ (v_x)_1 \end{pmatrix} = \begin{pmatrix} x_0 \\ (v_x)_0 \end{pmatrix} + \Delta t \begin{pmatrix} (v_x)_1 \\ \mathcal{P}x_1 \end{pmatrix},$$

$$\begin{aligned} \begin{pmatrix} (p_x)_1 \\ (p_{v_x})_1 \end{pmatrix} &= \begin{pmatrix} (p_x)_0 \\ (p_{v_x})_0 \end{pmatrix} - \Delta t \left(\begin{pmatrix} 0 & 1 \\ \mathcal{P} & 0 \end{pmatrix}^T \begin{pmatrix} (p_x)_0 \\ (p_{v_x})_0 \end{pmatrix} + \begin{pmatrix} 2x_1 & 0 \\ (v_x)_1 & x_1 \\ 0 & 2(v_x)_1 \end{pmatrix}^T \begin{pmatrix} \Lambda_1 \\ \Lambda_2 \\ \Lambda_3 \end{pmatrix} \right), \\ 0 &= \begin{pmatrix} x_1^2 + Y^2 - 1 \\ (v_x)_1 x_1 + V_y Y \\ (v_x)_1^2 + V_y^2 - Y + \mathcal{P} \end{pmatrix}, \\ 0 &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & x_1 \end{pmatrix}^T \begin{pmatrix} (p_x)_0 \\ (p_{v_x})_0 \end{pmatrix} + \begin{pmatrix} 2Y & 0 & 0 \\ V_y & Y & 0 \\ -1 & 2V_y & 1 \end{pmatrix}^T \begin{pmatrix} \Lambda_1 \\ \Lambda_2 \\ \Lambda_3 \end{pmatrix}. \end{aligned}$$

We refer to this method as PGHVI-1. We will compare this to the first-order method where the Runge–Kutta coefficients are the same for both q and p , i.e., $a = 1 = \tilde{a}$. This method, which we refer to as BE-1, is given by applying the backward Euler method in both the q and p variables, i.e.,

$$\begin{aligned} \begin{pmatrix} x_1 \\ (v_x)_1 \end{pmatrix} &= \begin{pmatrix} x_0 \\ (v_x)_0 \end{pmatrix} + \Delta t \begin{pmatrix} (v_x)_1 \\ \mathcal{P} x_1 \end{pmatrix}, \\ \begin{pmatrix} (p_x)_1 \\ (p_{v_x})_1 \end{pmatrix} &= \begin{pmatrix} (p_x)_0 \\ (p_{v_x})_0 \end{pmatrix} - \Delta t \left(\begin{pmatrix} 0 & 1 \\ \mathcal{P} & 0 \end{pmatrix}^T \begin{pmatrix} (p_x)_1 \\ (p_{v_x})_1 \end{pmatrix} + \begin{pmatrix} 2x_1 & 0 \\ (v_x)_1 & x_1 \\ 0 & 2(v_x)_1 \end{pmatrix}^T \begin{pmatrix} \Lambda_1 \\ \Lambda_2 \\ \Lambda_3 \end{pmatrix} \right), \\ 0 &= \begin{pmatrix} x_1^2 + Y^2 - 1 \\ (v_x)_1 x_1 + V_y Y \\ (v_x)_1^2 + V_y^2 - Y + \mathcal{P} \end{pmatrix}, \\ 0 &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & x_1 \end{pmatrix}^T \begin{pmatrix} (p_x)_1 \\ (p_{v_x})_1 \end{pmatrix} + \begin{pmatrix} 2Y & 0 & 0 \\ V_y & Y & 0 \\ -1 & 2V_y & 1 \end{pmatrix}^T \begin{pmatrix} \Lambda_1 \\ \Lambda_2 \\ \Lambda_3 \end{pmatrix}. \end{aligned}$$

For our numerical test, we will qualitatively compare the preservation of the presymplectic form $\Omega_0 = dx \wedge dp_x + dv_x \wedge dp_{v_x}$ between the two methods. Since Type II boundary conditions arise in adjoint sensitivity analysis, we place Type II boundary conditions, i.e., by specifying $q_0 = (x_0, (v_x)_0)$ and $p_1 = ((p_x)_1, (p_{v_x})_1)$, and subsequently, numerically solve the resulting system for q_1, p_0, U, Λ . We use various nearby values for the initial position $q_0 = (x_0, (v_x)_0)$ and various nearby values for the final momenta $p_1 = ((p_x)_1, (p_{v_x})_1)$. For a presymplectic integrator applied to a presymplectic system with presymplectic form $dx \wedge dp_x + dv_x \wedge dp_{v_x}$, we expect that the area occupied by the distribution of points $(x_0, (p_x)_0)$ is the same as the area occupied by the distribution of points $(x_1, (p_x)_1)$; similarly, we expect that the area occupied by the distribution of points $((v_x)_0, (p_{v_x})_0)$ is the same as the area occupied by the distribution of points $((v_x)_1, (p_{v_x})_1)$. Since we choose to only solve the system for one timestep, we take a large timestep to highlight the difference between the two methods, $\Delta t = 2$, which corresponds to roughly one-third of the period of the pendulum.

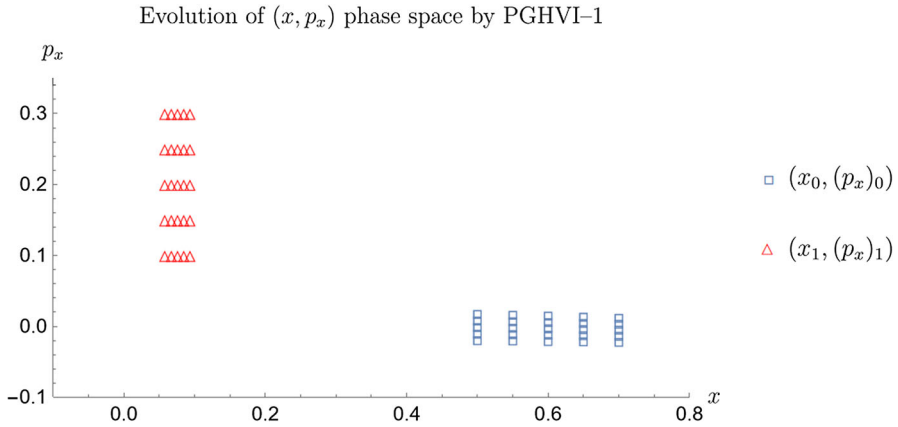


Fig. 1 (x, p_x) phase space cross section of PGHVI-1 applied to a distribution of initial conditions q_0 and final momenta p_1

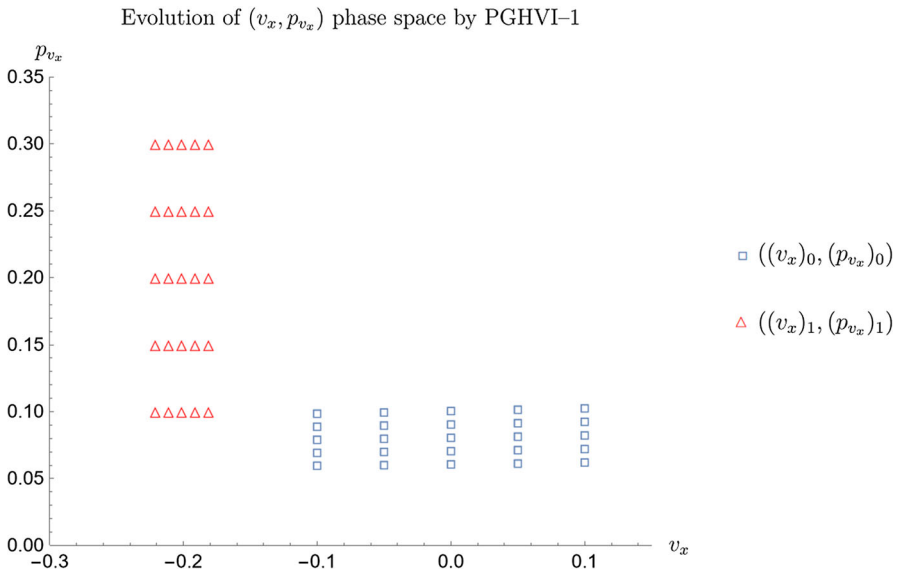


Fig. 2 (v_x, p_{v_x}) phase space cross section of PGHVI-1 applied to a distribution of initial conditions q_0 and final momenta p_1

Note that, with Type II boundary conditions, both methods give a map $(q_0, p_1) \mapsto (q_1, p_0)$ which implicitly determines an evolution map $(q_0, p_0) \mapsto (q_1, p_1)$; below, we plot the phase space cross sections of these implicit evolution maps. The evolution of the (x, p_x) and (v_x, p_{v_x}) distributions by PGHVI-1 is shown in Figs. 1 and 2, respectively. The evolution of the (x, p_x) and (v_x, p_{v_x}) distributions by BE-1 is shown in Figs. 3 and 4, respectively. As can be qualitatively seen from Figs. 1, 2, 3, 4, the PGHVI-1 method preserves the phase space area in both the (x, p_x) and (v_x, p_{v_x}) cross sections, whereas the B-1 method does not.

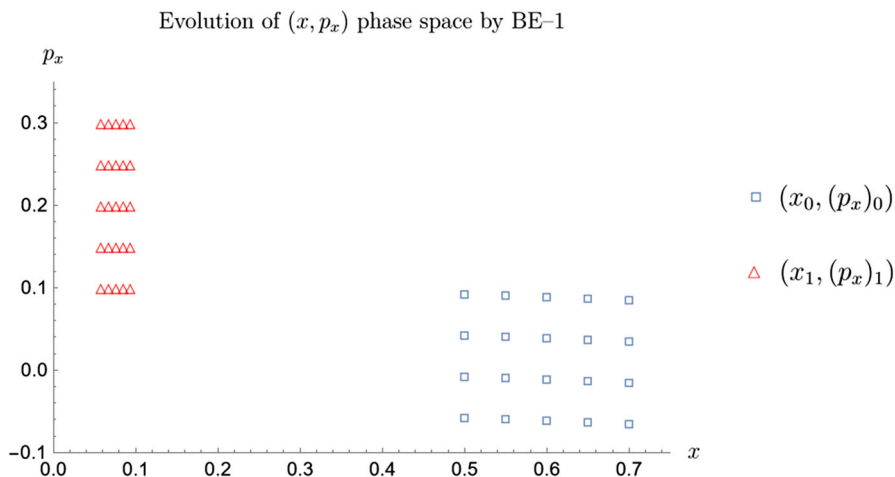


Fig. 3 (x, p_x) phase space cross section of BE-1 applied to a distribution of initial conditions q_0 and final momenta p_1

3.3 Optimal Control of DAE Systems

In this section, we derive the optimality conditions for an optimal control problem (OCP) subject to a semi-explicit DAE constraint. It is known that the optimality conditions can be described as a presymplectic system on the generalized phase space bundle (Delgado-Télllez and Ibrort 2003; Echeverría-Enríquez et al. 2003). For a discussion of the presymplectic geometry of optimal control systems and in particular, symmetries of such systems, see de León et al. (2004). We will subsequently consider a variational discretization of such OCPs and discuss the naturality of such discretizations.

Consider the following optimal control problem in Bolza form, subject to a DAE constraint, which we refer to as (OCP-DAE),

$$\min C(q(t_f)) + \int_0^{t_f} L(q, u)dt$$

subject to

$$\begin{aligned} \dot{q} &= f(q, u), \\ 0 &= \phi(q, u), \\ q_0 &= q(0), \\ 0 &= \phi_f(q(t_f)), \end{aligned}$$

where the DAE system $\dot{q} = f(q, u), 0 = \phi(q, u)$ is over $M_d \times M_a$ as described in Sect. 2.3, $C : M_d \rightarrow \mathbb{R}$ is the terminal cost, $L : M_d \times M_a \rightarrow \mathbb{R}$ is the running cost, the initial condition $q(0) = q_0$ is prescribed, and for generality, a terminal constraint $\phi_f(q(t_f)) = 0$ is also imposed, where ϕ_f is a map from M_d into some vector space V .

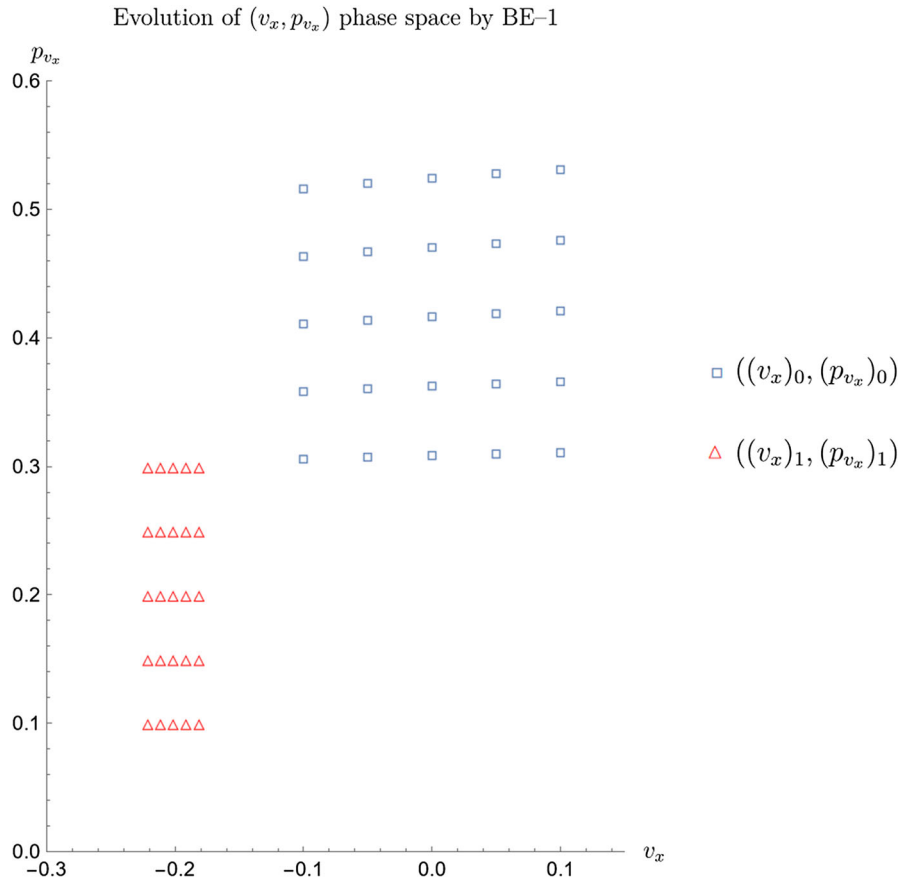


Fig. 4 (v_x, p_{v_x}) phase space cross section of BE-1 applied to a distribution of initial conditions q_0 and final momenta p_1

We assume a local optimum to (OCP-DAE). We then adjoin the constraints to J using adjoint variables, which gives the adjoined functional

$$\mathcal{J} = C(q(t_f)) + \langle \lambda_f, \phi_f(q(t_f)) \rangle + \int_0^{t_f} [L(q, u) + \langle p, f(q, u) - \dot{q} \rangle + \langle \lambda, \phi(q, u) \rangle] dt.$$

The optimality conditions are given by the condition that \mathcal{J} is stationary about the local optimum, $\delta \mathcal{J} = 0$ (Biegler 2010). For simplicity in the notation, we will use matrix derivatives instead of indices. Note also that we will implicitly leave out the variation of the adjoint variables, since those terms pair with the DAE constraints, which vanish at the local optimum. The optimality condition $\delta \mathcal{J} = 0$ is then

$$0 = \delta \mathcal{J} = \langle \nabla_q C(q(t_f)), \delta q(t_f) \rangle + \langle \lambda_f, D_q \phi_f(q(t_f)) \delta q(t_f) \rangle + \int_0^{t_f} \left[\langle \nabla_q L(q, u), \delta q \rangle + \langle \nabla_u L(q, u), \delta u \rangle + \langle p, D_q f(q, u) \delta q \rangle + \langle p, D_u f(q, u) \delta u \rangle \right] dt$$

$$\begin{aligned}
 & - \left\langle p, \frac{d}{dt} \delta q \right\rangle + \langle \lambda, D_q \phi(q, u) \delta q \rangle + \langle \lambda, D_u \phi(q, u) \delta u \rangle \Big] dt \\
 = & \langle \nabla_q C(q(t_f)) + [D_q \phi_f(q(t_f))]^* \lambda_f - p(t_f), \delta q(t_f) \rangle \\
 & + \int_0^{t_f} \left[\langle \nabla_q L(q, u) + [D_q f(q, u)]^* p + \dot{p} + [D_q \phi(q, u)]^* \lambda, \delta q \rangle \right. \\
 & \left. + \langle \nabla_u L(q, u) + [D_u f(q, u)]^* p + [D_u \phi(q, u)]^* \lambda, \delta u \rangle \right] dt,
 \end{aligned}$$

where we integrated by parts on the term $\langle p, \frac{d}{dt} \delta q \rangle$ and used $\delta q(0) = 0$ since the initial condition is fixed. Enforcing stationarity for all such variations gives the optimality conditions,

$$\dot{q} = f(q, u), \tag{3.12a}$$

$$\dot{p} = -[D_q f(q, u)]^* p - [D_q \phi(q, u)]^* \lambda - \nabla_q L(q, u), \tag{3.12b}$$

$$0 = \phi(q, u), \tag{3.12c}$$

$$0 = \nabla_u L(q, u) + [D_u f(q, u)]^* p + [D_u \phi(q, u)]^* \lambda, \tag{3.12d}$$

$$0 = \phi_f(q(t_f)), \tag{3.12e}$$

$$p(t_f) = \nabla_q C(q(t_f)) + [D_q \phi_f(q(t_f))]^* \lambda_f. \tag{3.12f}$$

The first four optimality conditions (3.12a)–(3.12d) are precisely the augmented adjoint DAE equations, (2.17a)–(2.17d). The last two optimality conditions (3.12e), (3.12f) are the terminal constraint and the associated transversality condition, respectively. Note that these conditions are only sufficient for a trajectory (q, u, p, λ) to be an extremum of the optimal control problem; whether or not the trajectory is optimal depends on the properties of the DAE constraint and cost function, e.g., convexity of L .

Regular Index 1 Optimal Control In the literature, the problem (OCP-DAE) is usually formulated by making a distinction between algebraic variables and control variables, (q, y, u) , instead of (q, u) (see, for example, Biegler 2010 and Aguiar et al. 2021). This does not change any of the previous discussion of the optimality conditions, except that (3.12d) splits into two equations for y and u . That is, the distinction is not formally important for the previous discussion. It is of course important when actually solving such an optimal control problem. For example, the constraint function $\phi(q, y, u)$ may have a singular matrix derivative with respect to (y, u) but may have a nonsingular matrix derivative with respect to y . In such a case, one interprets y as the algebraic variable, in that it can locally be solved in terms of (q, u) via the constraint, and the control variable u as “free” to optimize over. We now briefly elaborate on this case.

We take the configuration manifold for the algebraic variables to be $M_a = Y_a \times U \ni (y, u)$, where y is interpreted as the algebraic constraint variable and u is interpreted as the control variable. We will assume that the control space U is compact. The constraint has the form $\phi(q, y, u) = 0$, and we assume that $\partial \phi / \partial y$ is pointwise invertible. We

consider the following optimal control problem,

$$\begin{aligned} & \min \int_0^{t_f} L(q, y, u) dt \\ & \text{subject to} \\ & \quad \dot{q} = f(q, y, u), \\ & \quad 0 = \phi(q, y, u), \\ & \quad q_0 = q(0). \end{aligned}$$

We perform an analogous argument to before, except that, in this case, since U may have a boundary, the optimality for the control variable u will either require u to lie on ∂U or will require the stationarity of the adjoined functional with respect to variations in u . In any case, the necessary conditions for optimality can be expressed as

$$\dot{q} = f(q, y, u), \quad (3.13a)$$

$$\dot{p} = -[D_q f(q, y, u)]^* p - [D_q \phi(q, y, u)]^* \lambda - \nabla_q L(q, y, u), \quad (3.13b)$$

$$0 = \phi(q, y, u), \quad (3.13c)$$

$$0 = \nabla_y L(q, y, u) + [D_y f(q, y, u)]^* p + [D_y \phi(q, y, u)]^* \lambda, \quad (3.13d)$$

$$u = \arg \min_{u' \in U} H_L(q, y, u'), \quad (3.13e)$$

$$0 = p(t_f), \quad (3.13f)$$

where H_L is the augmented Hamiltonian $H_L(q, y, u) = L(q, y, u) + \langle p, f(q, y, u) \rangle + \langle \lambda, \phi(q, y, u) \rangle$. Assuming that u lies in the interior of U , (3.13e) can be expressed as

$$0 = \nabla_u L(q, y, u) + [D_u f(q, y, u)]^* p + [D_u \phi(q, y, u)]^* \lambda,$$

or $D_u H_L(q, y, u) = 0$. We say that an optimal control problem with a DAE constraint forms a regular index 1 system if both $\partial \phi / \partial y$ and the Hessian $D_u^2 H_L$ are pointwise invertible. In this case, whenever u lies in the interior of U , (y, u, λ) can be locally solved as functions of (q, p) . Thus, in principle, the resulting Hamiltonian ODE for (q, p) can be integrated to yield extremal trajectories for the optimal control problem. As mentioned before, without additional assumptions on the DAE and cost function, such a trajectory will only generally be an extremum but not necessarily optimal.

Of course, in practice, one cannot generally analytically integrate the resulting ODE nor determine the functions which give (y, u, λ) in terms of (q, p) . Thus, the only practical option is to discretize the presymplectic system above to compute approximate extremal trajectories. To integrate such a presymplectic system, one can again use the presymplectic Galerkin Hamiltonian variational integrator construction discussed in Sect. 3.2. Such an integrator would be natural in the following sense. First, as discussed in Sect. 3.2, a presymplectic Galerkin Hamiltonian variational integrator applied to the augmented adjoint DAE system is equivalent to applying a symplectic Galerkin Hamiltonian variational integrator to the underlying Hamiltonian ODE, with the same Runge–Kutta expansions for q_1, Q^j in both methods. Furthermore, as shown

in Sanz-Serna (2016), utilizing a symplectic integrator to discretize the extremality conditions is equivalent to first discretizing the ODE constraint by a Runge–Kutta method and then enforcing the associated discrete extremality conditions. This also holds in the DAE case.

More precisely, beginning with a regular index 1 optimal control problem, the processes of reduction, extremization, and discretization commute, for suitable choices of these processes, analogous to those used in the naturality result discussed in Sect. 3.2.1. The proof is similar to the naturality result discussed in Sect. 3.2.1, where the arrow given by forming the adjoint is replaced by extremization. In essence, these are the same, since the extremization condition is given by the adjoint system, so we will just elaborate briefly. We already know how to extremize the continuous optimal control problem, with either a DAE constraint or an ODE constraint after reduction, which results in an adjoint system. We also already know how to discretize the resulting adjoint system after discretization, using a (pre)symplectic partitioned Runge–Kutta method. Furthermore, at any step, reduction is just defined to be solving the continuous or discrete constraints for y in terms of (q, u) . Thus, the only major difference compared to the previous naturality result is defining the discretization of the optimal control problem and subsequently, how to extremize the discrete optimal control problem. For the regular index 1 optimal control problem,

$$\begin{aligned} & \min \int_0^{t_f} L(q, y, u) dt \\ & \text{subject to} \\ & \quad \dot{q} = f(q, y, u), \\ & \quad 0 = \phi(q, y, u), \\ & \quad q_0 = q(0), \end{aligned}$$

its discretization is obtained by replacing the constraints with a Runge–Kutta discretization and replacing the cost function with its quadrature approximation, using the same quadrature weights as those in the Runge–Kutta discretization. This can be written as

$$\begin{aligned} & \min \Delta t \sum_i b_i L(Q^i, Y^i, U^i) \\ & \text{subject to} \\ & \quad V^i = f(Q^i, Y^i, U^i), \\ & \quad 0 = \phi(Q^i, Y^i, U^i), \end{aligned}$$

where $Q^i = q_0 + \Delta t \sum_j a_{ij} V^j$, which implicitly encodes $q(0) = q_0$. One can then extremize this discrete system, which is given by the discrete Euler–Lagrange equations for the discrete action

$$\mathbb{S} = \Delta t \sum_i b_i \left(\langle P^i, V^i - f(Q^i, Y^i, U^i) \rangle - \langle \Lambda^i, \phi(Q^i, Y^i, U^i) \rangle - L(Q^i, Y^i, U^i) \right).$$

That is, we enforce the discrete constraints by adding to the discrete Lagrangian the appropriate Lagrange multiplier terms paired with the constraints, where we weighted the Lagrange multipliers P^i , Λ^i by $\Delta t b_i$ just as convention, in order to interpret them as the appropriate variables, as discussed in Appendix B. Enforcing extremality of this action recovers a partitioned Runge–Kutta method applied to the adjoint system corresponding to extremizing the continuous optimal control problem, as discussed in Appendix B, where the Runge–Kutta coefficients for the momenta are the symplectic adjoint of the original Runge–Kutta coefficients. Alternatively, starting from the original continuous optimal control problem, one could first reduce the DAE constraint to an ODE constraint using the invertibility of $D_y\phi$ to give

$$\begin{aligned} & \min \int_0^{t_f} L(q, y(q, u), u) dt \\ & \text{subject to} \\ & \quad \dot{q} = f(q, y(q, u), u), \\ & \quad q_0 = q(0). \end{aligned}$$

One can then discretize this using the same Runge–Kutta method as before, where the cost function is replaced with a quadrature approximation, and then extremize using Lagrange multipliers. Alternatively, one can extremize the continuous problem to yield an adjoint system and then apply a partitioned Runge–Kutta method to that system, where the momenta Runge–Kutta coefficients are again the symplectic adjoint of the original Runge–Kutta coefficients. Having defined all of these processes, a direct computation yields that all of the processes commute, analogous to the computation in Appendix B.

4 Conclusion and Future Research Directions

In this paper, we utilized symplectic and presymplectic geometry to study the properties of adjoint systems associated with ODEs and DAEs, respectively. The (pre)symplectic structure of these adjoint systems led us to a geometric characterization of the adjoint variational quadratic conservation law used in adjoint sensitivity analysis. As an application of this geometric characterization, we constructed structure-preserving discretizations of adjoint systems by utilizing (pre)symplectic integrators, which led to natural discrete analogues of the quadratic conservation laws.

A natural research direction is to extend the current framework to adjoint systems for differential equations with nonholonomic constraints, in order to more generally allow for constraints between configuration variables and their derivatives. In this setting, it is reasonable to expect that the geometry of the associated adjoint systems can be described using Dirac structures (see, for example, Yoshimura and Marsden 2006a, b), which generalize the symplectic and presymplectic structures of adjoint ODE and DAE systems, respectively. Structure-preserving discretizations of such systems could then be studied through the lens of discrete Dirac structures (Leok and Ohsawa 2011). These discrete Dirac structures make use of the notion of a retraction (Absil et al.

2008). The tangent and cotangent lifts of a retraction also provide a useful framework for constructing geometric integrators (Barbero-Liñán and Martín de Diego 2021). It would be interesting to synthesize the notion of tangent and cotangent lifts of retraction maps with discrete Dirac structures in order to construct discrete Dirac integrators for adjoint systems with nonholonomic constraints which generalize the presymplectic integrators constructed in Barbero-Liñán and Martín de Diego (2022).

Another natural research direction is to extend the current framework to evolutionary partial differential equations (PDEs). There are two possible approaches in this direction. The first is to consider evolutionary PDEs as ODEs evolving on infinite-dimensional spaces, such as Banach or Hilbert manifolds. One can then investigate the geometry of the infinite-dimensional symplectic structure associated with the corresponding adjoint system. In practice, adjoint systems for evolutionary PDEs are often formed after semi-discretization, leading to an ODE on a finite-dimensional space. Understanding the reduction of the infinite-dimensional symplectic structure of the adjoint system to a finite-dimensional symplectic structure under semi-discretization could provide useful insights into structure preservation. The second approach would be to explore the multisymplectic structure of the adjoint system associated with a PDE. This approach would be insightful for several reasons. First, an adjoint variational quadratic conservation law arising from multisymplecticity would be adapted to spacetime instead of just time. With appropriate spacetime splitting and boundary conditions, such a quadratic conservation law would induce either a temporal or spatial conservation law. As such, one could use the multisymplectic conservation law to determine adjoint sensitivities for a PDE with respect to spatial or temporal directions, which could be useful in practice (Li and Petzold 2004). Furthermore, the multisymplectic framework would apply equally as well to nonevolutionary (elliptic) PDEs, where there is no interpretation of a PDE as an infinite-dimensional evolutionary ODE. Additionally, adjoint systems for PDEs with constraints could be investigated with multi-Dirac structures (Vankerschaver et al. 2012). In future work, we aim to explore both approaches, relate them once a spacetime splitting has been chosen, and investigate structure-preserving discretizations of such systems by utilizing the multisymplectic variational integrators constructed in Tran and Leok (2022).

Acknowledgements BT was supported by the NSF Graduate Research Fellowship DGE-2038238, and by NSF under grants DMS-1813635. ML was supported by NSF under grants DMS-1345013, DMS-1813635, DMS-2307801, and by AFOSR under grants FA9550-18-1-0288, FA9550-23-1-0279.

Data Availability The data generated in this paper are available from the corresponding author upon reasonable request.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Proofs of Discrete Adjoint Variational Quadratic Conservation Laws

Proof of Proposition 3.1 We begin by substituting (3.4c) and (3.3a) into the left-hand side of (3.5),

$$\begin{aligned}
 \langle p_1, \delta q_1 \rangle &= \langle p_0, \delta q_0 \rangle + \Delta t \sum_i b_i \langle p_0, Df(Q^i) \delta Q^i \rangle \\
 &\quad - \Delta t \sum_i b_i \langle [Df(Q^i)]^* P^i, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta q_0 \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle [Df(Q^i)]^* P^i, Df(Q^j) \delta Q^j \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle dL(Q^i), Df(Q^j) \delta Q^j \rangle \\
 &= \langle p_0, \delta q_0 \rangle + \Delta t \sum_i b_i \left\langle P^i + \Delta t \sum_j \tilde{a}_{ij} ([Df(Q^j)]^* P^j + dL(Q^j)), Df(Q^i) \delta Q^i \right\rangle \\
 &\quad - \Delta t \sum_i b_i \left\langle [Df(Q^i)]^* P^i, \delta Q^i - \Delta t \sum_j a_{ij} Df(Q^j) \delta Q^j \right\rangle \\
 &\quad - \Delta t \sum_i b_i \langle dL(Q^i), \delta q_0 \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle [Df(Q^i)]^* P^i, Df(Q^j) \delta Q^j \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle dL(Q^i), Df(Q^j) \delta Q^j \rangle,
 \end{aligned}$$

where, in the last equality, we substituted (3.4d) and (3.3b). We now group and simplify the above expression,

$$\begin{aligned}
 \langle p_1, \delta q_1 \rangle &= \langle p_0, \delta q_0 \rangle + \Delta t \sum_i b_i \left\langle P^i + \Delta t \sum_j \tilde{a}_{ij} [Df(Q^j)]^* P^j, Df(Q^i) \delta Q^i \right\rangle \\
 &\quad - \Delta t \sum_i b_i \left\langle [Df(Q^i)]^* P^i, \delta Q^i - \Delta t \sum_j a_{ij} Df(Q^j) \delta Q^j \right\rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle [Df(Q^i)]^* P^i, Df(Q^j) \delta Q^j \rangle \\
 &\quad + \Delta t \sum_i b_i \left\langle \Delta t \sum_j \tilde{a}_{ij} dL(Q^j), Df(Q^i) \delta Q^i \right\rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta q_0 \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle dL(Q^i), Df(Q^j) \delta Q^j \rangle
 \end{aligned}$$

$$\begin{aligned}
 &= \langle p_0, \delta q_0 \rangle + \Delta t^2 \sum_{ij} \underbrace{(b_j \tilde{a}_{ji} + b_i a_{ij} - b_i b_j)}_{=0} \langle [Df(Q^i)]^* P^i, Df(Q^j) \delta Q^j \rangle \\
 &\quad + \Delta t \sum_i b_i \left\langle \Delta t \sum_j \tilde{a}_{ij} dL(Q^j), Df(Q^i) \delta Q^i \right\rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta q_0 \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle dL(Q^i), Df(Q^j) \delta Q^j \rangle \\
 &= \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta q_0 \rangle \\
 &\quad - \Delta t^2 \sum_{ij} \underbrace{(b_i b_j - b_j \tilde{a}_{ji})}_{=b_i a_{ij}} \langle dL(Q^i), Df(Q^j) \delta Q^j \rangle \\
 &= \langle p_0, \delta q_0 \rangle - \Delta t^2 \sum_i b_i \left\langle dL(Q^i), \frac{\delta q_0}{\Delta t} + \sum_j a_{ij} Df(Q^j) \delta Q^j \right\rangle \\
 &= \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle dL(Q^i), \delta Q^i \rangle,
 \end{aligned}$$

where, in the last equality, we used (3.3b). □

Proof of Proposition 3.2 For brevity, we denote

$$\begin{aligned}
 D_q f_i &\equiv D_q f(Q^i, U^i), \\
 D_u f_i &\equiv D_u f(Q^i, U^i), \\
 D_q \phi_i &\equiv D_q \phi(Q^i, U^i), \\
 D_u \phi_i &\equiv D_u \phi(Q^i, U^i).
 \end{aligned}$$

Starting from $\langle p_1, \delta q_1 \rangle$, we substitute the evolution equations (3.7c), (3.7d), (3.8a), (3.8b),

$$\begin{aligned}
 \langle p_1, \delta q_1 \rangle &= \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle [D_q f_i]^* P^i + [D_q \phi_i]^* \Lambda^i, \delta q_0 \rangle \\
 &\quad + \Delta t \sum_i b_i \langle p_0, D_q f_i \delta Q^i + D_u f_i \delta U^i \rangle \\
 &\quad - \Delta t^2 \sum_{ij} b_i b_j \langle [D_q f_i]^* P^i + [D_q \phi_i]^* \Lambda^i, D_q f_j \delta Q^j + D_u f_j \delta U^j \rangle \\
 &= \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \left\langle [D_q f_i]^* P^i + [D_q \phi_i]^* \Lambda^i, \delta Q^i \right. \\
 &\quad \left. - \Delta t \sum_j a_{ij} (D_q f_j \delta Q^j + D_u f_j \delta U^j) \right\rangle \\
 &\quad + \Delta t \sum_i b_i \left\langle P^i + \Delta t \sum_j \tilde{a}_{ij} ([D_q f_j]^* P^j + [D_q \phi_j]^* \Lambda^j), D_q f_i \delta Q^i + D_u f_i \delta U^i \right\rangle
 \end{aligned}$$

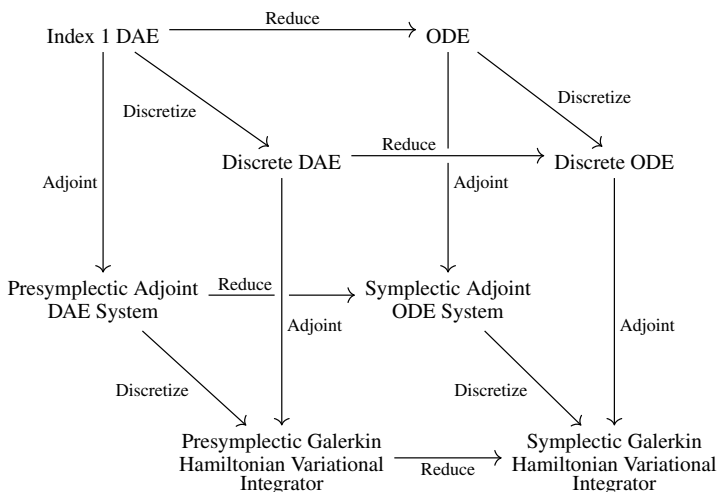
$$\begin{aligned}
 & - \Delta t^2 \sum_{ij} b_i b_j \langle [D_q f_i]^* P^i + [D_q \phi_i]^* \Lambda^i, D_q f_j \delta Q^j + D_u f_j \delta U^j \rangle \\
 = & \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle [D_q f_i]^* P^i + [D_q \phi_i]^* \Lambda^i, \delta Q^i \rangle \\
 & + \Delta t \sum_i b_i \langle P^i, D_q f_i \delta Q^i + D_u f_i \delta U^i \rangle \\
 & + \Delta t^2 \sum_{ij} \underbrace{(b_j \tilde{a}_{ji} + b_i a_{ij} - b_i b_j)}_{=0} \langle [D_q f_i]^* P^i \\
 & + [D_q \phi_i]^* \Lambda^i, D_q f_j \delta Q^j + D_u f_j \delta U^j \rangle \\
 = & \langle p_0, \delta q_0 \rangle + \Delta t \sum_i b_i \left(\langle [D_u f_i]^* P^i, \delta U^i \rangle - \langle [D_q \phi_i]^* \Lambda^i, \delta Q^i \rangle \right) \\
 = & \langle p_0, \delta q_0 \rangle + \Delta t \sum_i b_i \left(- \langle [D_u \phi_i]^* \Lambda^i, \delta U^i \rangle - \langle [D_q \phi_i]^* \Lambda^i, \delta Q^i \rangle \right) \\
 = & \langle p_0, \delta q_0 \rangle - \Delta t \sum_i b_i \langle \Lambda^i, D_u \phi_i \delta U^i + D_q \phi_i \delta Q^i \rangle = \langle p_0, \delta q_0 \rangle,
 \end{aligned}$$

where in the third to last equality, we used the constraint equation (3.7f) and in the last equality, we used the constraint equation (3.8c). □

Proof of Proposition 3.3 The proof uses computations analogous to those used in the proofs of Propositions 3.1 and 3.2. In particular, starting from the simplest case of the nonaugmented adjoint ODE system, Proposition 3.1 considers the case of augmenting the Hamiltonian, whereas Proposition 3.2 considers the case of replacing the ODE with a DAE. The case at hand combines both and the proof involves a combination of both computations. □

Appendix B: Proof of Naturality of Adjoint System Discretization

In this appendix, we prove the statement in Sect. 3.2.1 that (for suitable choices of) discretization, reduction, and forming the adjoint all commute when applied to an index 1 DAE. The definitions and choices of these processes were made in Sect. 3.2.1. To prove that the diagram commutes, we prove that each face of the diagram commutes. We again include the relevant diagram which we wish to show commutes below.



Back Face We have already proved that the back face commutes (i.e., that reduction and forming the adjoint commute when starting with an index 1 DAE), as discussed in Sect. 2.3.1. One can then interpret the above diagram as an extension of this result with an extra dimension corresponding to discretization.

Right Face This was proven in Sanz-Serna (2016). One can then interpret the above diagram as an extension of the result in Sanz-Serna (2016) by adding the reduction operation.

Bottom Face Consider the augmented adjoint DAE system corresponding to the DAE (2.12a)–(2.12b), which we take to have index 1, i.e., $\partial\phi/\partial u$ is pointwise invertible. We consider the augmented case because the nonaugmented case can be obtained by taking $L \equiv 0$. We show that reducing the system first and then applying a symplectic Galerkin Hamiltonian variational integrator is equivalent to applying a presymplectic Galerkin Hamiltonian variational integrator, with the same partitioned Runge–Kutta coefficients, and then reducing.

We start with the former approach. The symplectic adjoint ODE system given by reduction, as discussed in Sect. 2.3.1, is the Hamiltonian system corresponding to the Hamiltonian

$$H(q', p') = \langle p', f'(q') \rangle + L'(q'),$$

where we have solved $u = u(q')$ and defined $f'(q') \equiv f(q', u(q'))$, $L'(q') \equiv L(q', u(q'))$. Applying the symplectic Galerkin Hamiltonian variational integrator construction yields the integrator

$$q_1 = q_0 + \Delta t \sum_i b_i f'(Q^i) \tag{B.1a}$$

$$\begin{aligned}
 &= q_0 + \Delta t \sum_i b_i f(Q^i, u(Q^i)), \\
 Q^i &= q_0 + \Delta t \sum_j a_{ij} f'(Q^j) \tag{B.1b}
 \end{aligned}$$

$$\begin{aligned}
 &= q_0 + \Delta t \sum_j a_{ij} f(Q^j, u(Q^j)), \\
 p_1 &= p_0 - \Delta t \sum_i b_i ([Df'(Q^i)]^* P^i + dL'(Q^i)), \tag{B.1c}
 \end{aligned}$$

$$P^i = p_0 - \Delta t \sum_j \tilde{a}_{ij} ([Df'(Q^j)]^* P^j + dL'(Q^j)). \tag{B.1d}$$

Note that the derivative Df' can be equivalently expressed as

$$Df'(Q^i) = D_1 f(Q^i, u(Q^i)) + D_2 f(Q^i, u(Q^i)) Du(Q^i),$$

where D_i denotes differentiation with respect to the i^{th} argument. We switch to indexing the derivative operator here, so we do not have to make the distinction between total derivatives D_q and partial derivatives ∂_q . Similarly, we can express dL' as follows. First, note that we have been implicitly identifying the row vector dL' with the column vector given by its transpose $\nabla L'$. Thus, dL' in equations (B.1c)–(B.1d) should really be written as $\nabla L'$. Thus,

$$dL'(Q^i) \cong \nabla L'(Q^i) = \nabla_1 L(Q^i, u(Q^i)) + [Du(Q^i)]^* \nabla_2 L(Q^i, u(Q^i)).$$

Now, we show that the second approach is equivalent to the above system. The starting point is the presymplectic Galerkin Hamiltonian variational integrator, equations (3.9a)–(3.9f). From (3.9e), we can solve for U^i in terms of Q^i as $U^i = u(Q^i)$. Plugging this into (3.9a)–(3.9b) gives precisely (B.1a)–(B.1b). Thus, we just need to see that, after solving the constraint (3.9f) for Λ^i , the two momenta equations (3.9c)–(3.9d) are equivalent to (B.1c)–(B.1d). Solving (3.9f) for Λ^i gives

$$\begin{aligned}
 \Lambda^i &= -([D_2 \phi(Q^i, u(Q^i))]^*)^{-1} [D_2 f(Q^i, u(Q^i))]^* P^i \\
 &\quad - ([D_2 \phi(Q^i, u(Q^i))]^*)^{-1} \nabla_2 L(Q^i, u(Q^i)).
 \end{aligned}$$

Multiplying both sides by $[D_1 \phi(Q^i, u(Q^i))]^*$ yields

$$\begin{aligned}
 &[D_1 \phi(Q^i, u(Q^i))]^* \Lambda^i \\
 &= -[D_1 \phi(Q^i, u(Q^i))]^* ([D_u \phi(Q^i, u(Q^i))]^*)^{-1} [D_2 f(Q^i, u(Q^i))]^* P^i \\
 &\quad - [D_1 \phi(Q^i, u(Q^i))]^* ([\nabla_u \phi(Q^i, u(Q^i))]^*)^{-1} \nabla_2 L(Q^i, u(Q^i)) \\
 &= [Du(Q^i)]^* [D_2 f(Q^i, u(Q^i))]^* P^i + [Du(Q^i)]^* \nabla_2 L(Q^i, u(Q^i)) \\
 &= [D_2 f(Q^i, u(Q^i)) Du(Q^i)]^* P^i + [Du(Q^i)]^* \nabla_2 L(Q^i, u(Q^i)),
 \end{aligned}$$

where in the second equality, we used $D_1\phi(Q^i, u(Q^i)) = -D_2\phi(Q^i, u(Q^i))Du(Q^i)$ from the implicit function theorem. Plugging this expression and $U^i = u(Q^i)$ into (3.9c)–(3.9d) yields (B.1c)–(B.1d), noting the above expressions for Df' , dL' .

Remark B.1 Note that, in the above, we used the implicit function theorem to obtain the local function $u = u(q)$. This is sufficient to prove that the two processes are the same for a single integration step, assuming that the timestep Δt is sufficiently small and the vector field f and constraint ϕ are sufficiently regular, so that q_0, q_1 , and all of the internal stages Q^i are in the neighborhood where the local function is defined. For each subsequent time step, one generally needs a different local function. This does not matter in practice since one works directly with the presymplectic integrator and solves the constraints iteratively.

Top Face We want to prove that, starting from an index 1 DAE, the processes of discretization and reduction commute, where the discretization of the ODE and DAE have the same Runge–Kutta coefficients.

We start first with reduction then discretization. Starting from the index 1 DAE $\dot{q} = f(q, u)$, $\phi(q, u) = 0$, we apply the reduction operation, which gives the ODE $\dot{q} = f(q, u(q))$. Applying a Runge–Kutta discretization gives

$$\begin{aligned} q_1 &= q_0 + \Delta t \sum_i b_i f(Q^i, u(Q^i)), \\ Q^i &= q_0 + \Delta t \sum_j a_{ij} f(Q^j, u(Q^j)). \end{aligned}$$

On the other hand, we can discretize the DAE and then reduce. We discretize the DAE $\dot{q} = f(q, u)$, $\phi(q, u) = 0$ by applying a Runge–Kutta discretization with the same coefficients as before,

$$\begin{aligned} q_1 &= q_0 + \Delta t \sum_i b_i f(Q^i, U^i), \\ Q^i &= q_0 + \Delta t \sum_j a_{ij} f(Q^j, U^j), \\ 0 &= \phi(Q^i, U^i). \end{aligned}$$

To reduce this system, we solve the constraint equations $U^i = u(Q^i)$ and substitute these into the two evolution equations, which yields the same system obtained from first reducing and then discretizing.

Front Face The starting point for this loop is a discrete DAE system, which arises as a Runge–Kutta discretization of an index 1 DAE, i.e., it is given by the discrete system

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i, U^i), \quad (\text{B.2a})$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j, U^j), \quad (\text{B.2b})$$

$$0 = \phi(Q^i, U^i). \quad (\text{B.2c})$$

From here, we wish to show that reducing and forming the discrete adjoint system commute.

First, we recall the notion of a discrete adjoint system. Suppose we are given a generally nonlinear system of equations, $F(x_1) = x_0$, where $x_1 \in V$ is unknown, $x_0 \in W$ is given, and $F : V \rightarrow W$ (where V and W are vector spaces). To define the adjoint system, we first consider the variational equations associated with this nonlinear system given by its linearization,

$$DF(x_1)\delta x_1 = \delta x_0,$$

where $DF(x_1)$ is a linear map $V \rightarrow W$ and $\delta x_0 \in W$ is given. Suppose that we are interested in computing the quantity $\langle s_1, \delta x_1 \rangle$ for a given vector $s_1 \in V^*$. In the setting of adjoint sensitivity analysis, the quantity $\langle s_1, \delta x_1 \rangle$ is the sensitivity of the terminal cost function. We define the associated adjoint equation as

$$[DF(x_1)]^* s_0 = s_1.$$

For a solution $s_0 \in W^*$ of this system, one has

$$\langle s_1, \delta x_1 \rangle = \langle [DF(x_1)]^* s_0, \delta x_1 \rangle = \langle s_0, DF(x_1)\delta x_1 \rangle = \langle s_0, \delta x_0 \rangle.$$

Thus, to compute $\langle s_1, \delta x_1 \rangle$, one could solve the variational equation for δx_1 and pair it with s_1 which is given, or, alternatively, solve the adjoint equation for s_0 and pair it with δx_0 which is given, since these linear systems are solvable by assumption. We define the adjoint system associated with the equation $F(x_1) = x_0$ as this equation combined with the associated adjoint equation, i.e., as the combined system

$$\begin{aligned} F(x_1) &= x_0, \\ [Df(x_1)]^* s_0 &= s_1. \end{aligned}$$

Following Ibragimov (2006), we will utilize an alternative characterization of the adjoint system. We define the discrete adjoint action

$$\mathbb{S}(x_1, s_0) \equiv \langle s_0, F(x_1) \rangle.$$

Then, observe that \mathbb{S} is a generating function for the adjoint system $(x_1, s_0) \mapsto (x_0, s_1)$, in the sense that

$$\begin{aligned} x_0 &= \frac{\delta}{\delta s_0} \mathbb{S}(x_1, s_0) = F(x_1), \\ s_1 &= \frac{\delta}{\delta x_1} \mathbb{S}(x_1, s_0) = [Df(x_1)]^* s_0. \end{aligned}$$

This characterization serves two purposes. First, it will simplify the calculation of the adjoint system for the case at hand. Furthermore, it resembles the process of forming the adjoint at the continuous level: starting from the (discrete or continuous) differential(-algebraic) equation at hand, one forms the (discrete or continuous) adjoint action and applies the variational principle to obtain the adjoint system. To obtain the augmented adjoint system, we add a discrete Lagrangian $\mathbb{L} : V \rightarrow \mathbb{R}$ to the action (as a convention, we subtract the discrete Lagrangian). We define the augmented discrete adjoint action to be

$$\mathbb{S}_{\mathbb{L}}(x_1, s_0) \equiv \langle s_0, F(x_1) \rangle - \mathbb{L}(x_1).$$

The map that this generates defines the augmented discrete adjoint system,

$$\begin{aligned} x_0 &= \frac{\delta}{\delta s_0} \mathbb{S}_{\mathbb{L}}(x_1, s_0) = F(x_1), \\ s_1 &= \frac{\delta}{\delta x_1} \mathbb{S}_{\mathbb{L}}(x_1, s_0) = [Df(x_1)]^* s_0 - d\mathbb{L}(x_1). \end{aligned}$$

Observe that this definition of an augmented discrete adjoint system is natural in the sense that,

$$\langle s_1, \delta x_1 \rangle = \langle [Df(x_1)]^* s_0 + d\mathbb{L}(x_1), \delta x_1 \rangle = \langle s_0, x_0 \rangle - \langle dL(x_1), \delta x_1 \rangle,$$

which resembles the continuous analogue of the adjoint sensitivity result for a running cost function.

Now, we use this notion of a discrete adjoint system for the problem at hand. We begin first with reduction and then forming the adjoint system. Applying the reduction operation to the discrete DAE system (B.2a)–(B.2c), given by solving $\phi(Q^i, U^i) = 0$ for $U^i = u(Q^i)$, we obtain

$$q_1 = q_0 + \Delta t \sum_i b_i f(Q^i, u(Q^i)), \tag{B.3a}$$

$$Q^i = q_0 + \Delta t \sum_j a_{ij} f(Q^j, u(Q^j)). \tag{B.3b}$$

Let us define $Q^i = q_0 + \Delta t \sum_j a_{ij} V^j$. We think of the internal stages Q as functions of the internal stages V , which are the internal stage proxies for \dot{q} . Our discrete system (B.3a)–(B.3b) can then be defined by $x_1 = \{V^i\}_{i=1}^s$, $x_0 = \{0\}_{i=1}^s$, where s is the number of internal stages, and

$$x_0 = F(x_1) \equiv \begin{pmatrix} V^1 - f(Q^1(V), u(Q^1(V))) \\ \vdots \\ V^s - f(Q^s(V), u(Q^s(V))) \end{pmatrix}.$$

Observe that $F = 0$ only gives the internal stage equations (B.3b). We do this for simplicity, since we will assume $c_s = 1$ as is typical for a Runge–Kutta discretization of a DAE as previously discussed and hence, equation (B.3a) is redundant, since $a_{sj} = b_j$.

We define F and x_1 in terms of V instead of Q because when we form the adjoint action, we pair the components of F with the dual variable s_0 . In order to interpret s_0 as representing the momenta internal stages P^i , it should be paired with the proxy for the tangent vector V , instead of Q . We now form the discrete adjoint action. We define the dual variable for the adjoint system to be $s_0 = \{\Delta t b_i P^i\}_{i=1}^s$. The normalization factor $\Delta t b_i$ is used so that the discrete action is the quadrature approximation of the continuous action. This is just a convention, but we would have to reinterpret the components of s_0 if we did not choose this convention. Finally, we define the discrete Lagrangian to be the quadrature approximation of the continuous Lagrangian $L'(q) \equiv L(q, u(q))$, i.e., $\mathbb{L}(x_1) = \Delta t \sum_i b_i L'(Q^i(V))$. This is the natural choice because the discrete sensitivity of a running cost function is $\Delta t \sum_i b_i \langle dL'(Q^i(V)), \delta Q^i(V) \rangle$, which equals $\langle d\mathbb{L}(x_1), \delta x_1 \rangle$ with the above choice of \mathbb{L} . The augmented discrete adjoint action is then

$$\begin{aligned} \mathbb{S}_{\mathbb{L}}(\{V^i\}, \{b_i P^i\}) &= \mathbb{S}_{\mathbb{L}}(x_1, s_0) = \langle s_0, F(x_1) \rangle - \mathbb{L}(x_1) \\ &= \Delta t \sum_i b_i \left(\langle P^i, V^i - f(Q^i(V), u(Q^i(V))) \rangle - L'(Q^i(V)) \right). \end{aligned}$$

To define the discrete adjoint system, we have to give s_1 , which we take to be $s_1 = \{\Delta t b_i p_1\}_{i=1}^s$, where p_1 is given. Thus, the augmented discrete adjoint system is given by

$$\begin{aligned} 0 &= \frac{\delta}{\delta P^k} \mathbb{S}_{\mathbb{L}} = V^k - f(Q^k(V), u(Q^k(V))), \\ \Delta t b_k p_1 &= \frac{\delta}{\delta V^k} \mathbb{S}_{\mathbb{L}} \\ &= \Delta t b_k P^k - \Delta t^2 \sum_i b_i a_{ik} \left([D_1 f(Q^i(V), u(Q^i(V)))]^* P^i \right. \\ &\quad \left. + [D_2 f(Q^i(V), u(Q^i(V))) Du(Q^i(V))]^* P^i + dL'(Q^i(V)) \right). \end{aligned}$$

The first set of equations above, combined with the definition of Q in terms of V , gives (B.3b). For the second set of equations, we first divide through by $\Delta t b_k$ and rearrange to obtain

$$\begin{aligned} P^k &= p_1 + \Delta t \sum_i \frac{b_i a_{ik}}{b_k} \left([D_1 f(Q^i(V), u(Q^i(V)))]^* P^i \right. \\ &\quad \left. + [D_2 f(Q^i(V), u(Q^i(V))) Du(Q^i(V))]^* P^i + dL'(Q^i(V)) \right). \\ &= p_1 + \sum_i (b_i - \tilde{a}_{ki}) \left([D_1 f(Q^i(V), u(Q^i(V)))]^* P^i \right. \end{aligned}$$

$$+ [D_2 f(Q^i(V), u(Q^i(V))) Du(Q^i(V))]^* P^i + dL'(Q^i(V)).$$

Note that this is the usual symplectic partitioned Runge–Kutta expansion for the internal stages P^i , expressed in terms of p_1 instead of p_0 . Thus, the full adjoint system, combined with the redundant $k = s$ stages, yields a symplectic partitioned Runge–Kutta method.

Now, in the other direction, we first form the adjoint system corresponding to the discrete DAE system and subsequently reduce. We begin by forming the adjoint system. We form the discrete action analogously to before, but now the discrete system (B.2a)–(B.2c) also has constraints which we must incorporate into F , since we have not yet reduced the system. We take $x_1 = \{\{V^i\}, \{U^i\}\}_{i=1}^s$ and $s_0 = \{\{\Delta t b_i P^i\}, \{\Delta t b_i \Lambda^i\}\}_{i=1}^s$. We define F as

$$x_0 = F(x_1) \equiv \begin{pmatrix} V^1 - f(Q^1(V), U^1) \\ \vdots \\ V^s - f(Q^s(V), U^s) \\ -\phi(Q^1(V), U^1) \\ \vdots \\ -\phi(Q^s(V), U^s) \end{pmatrix}.$$

Note again that Q is a function of V as $Q^i = q_0 + \Delta t \sum_j a_{ij} V^j$. It is not a priori a function of U because the condition $V^i = f(Q^i(V), U^i)$ has not yet been enforced. Rather, it is a consequence of the variational principle, which formally matters when one computes the variation of the discrete action. Define the discrete Lagrangian $\mathbb{L}(x_1) = \sum_i b_i L(Q^i(V), U^i)$. We form the augmented discrete adjoint action

$$\begin{aligned} \mathbb{S}_{\mathbb{L}}(\{V^i\}, \{b_i P^i\}) &= \mathbb{S}_{\mathbb{L}}(x_1, s_0) = \langle s_0, F(x_1) \rangle - \mathbb{L}(x_1) \\ &= \Delta t \sum_i b_i \left(\langle P^i, V^i - f(Q^i(V), U^i) \rangle \right. \\ &\quad \left. - \langle \Lambda^i, \phi(Q^i, U^i) \rangle - L(Q^i(V), U^i) \right). \end{aligned}$$

We use this as a generating function to compute the adjoint system as before. The computation is analogous so we will just state the result,

$$\begin{aligned} V^k &= f(Q^k(V), U^k), \\ P^k &= p_1 + \Delta t \sum_i (b_i - \tilde{a}_{ki}) \left([D_1 f(Q^i(V), U^i)]^* P^i \right. \\ &\quad \left. + [D_1 \phi(Q^i(V), U^i)]^* \Lambda^i + D_q L(Q^i(V), U^i) \right), \\ 0 &= \phi(Q^i(V), U^i), \\ 0 &= [D_2 f(Q^i(V), U^i)]^* P^i + [D_2 \phi(Q^i(V), U^i)]^* \Lambda^i + D_2 L(Q^i(V), U^i). \end{aligned}$$

Finally, we reduce by solving the last two equations for U^i, Λ^i as functions of $Q^i(V^i), P^i$. Finally, an implicit function theorem computation analogous to the proof of the bottom face shows that this is the same as the system obtained by first reducing and then forming the discrete adjoint.

Left Face The proof for the left face is formally similar to the right face, but since we have already computed both directions, we will include it for completeness. Starting from an index 1 DAE, forming the adjoint and then discretizing just give the presymplectic Galerkin Hamiltonian variational integrator (3.9a)–(3.9f). In the other direction, we first discretize the DAE and then take the adjoint which we did in the proof of the front face. Expressed in terms of Q , instead of V , this is

$$\begin{aligned} Q^k &= q_0 + \Delta t \sum_j a_{ij} f(Q^j, U^j), \\ P^k &= p_1 + \Delta t \sum_i (b_i - \tilde{a}_{ki}) \left([D_1 f(Q^i, U^i)]^* P^i \right. \\ &\quad \left. + [D_1 \phi(Q^i, U^i)]^* \Lambda^i + D_q L(Q^i, U^i) \right), \\ 0 &= \phi(Q^i, U^i), \\ 0 &= [D_2 f(Q^i, U^i)]^* P^i + [D_2 \phi(Q^i, U^i)]^* \Lambda^i + D_2 L(Q^i, U^i). \end{aligned}$$

Returning to the system given by first forming the adjoint and then discretizing, (3.9a)–(3.9f), one substitutes (3.9c) into (3.9d) to write the internal stages for P^i in terms of p_1 , and this gives the above system.

Appendix C: An Intrinsic Type II Variational Principle for Adjoint Systems

We show that the adjoint system (2.10) arises from an intrinsic Type II variational principle. In coordinates, the type II variational principle corresponds to fixed initial position $q(t_0) = q_0$ and fixed final momenta $p(t_1) = p_1$, which are the boundary conditions used in adjoint sensitivity analysis, as discussed in Sect. 3.1.

Consider the augmented adjoint system


$$\begin{aligned} \dot{q} &= \partial H_L / \partial p = f(q), \\ \dot{p} &= -\partial H_L / \partial q = -[Df(q)]^* p - dL(q), \end{aligned}$$

where H_L is the augmented Hamiltonian. Recall that H_L is intrinsically defined by $H_L = i_{\hat{f}} \Theta + \pi_{T^*M}^* L$, where Θ is the tautological one-form on $T^*M, \pi_{T^*M} : T^*M \rightarrow M$ is the cotangent bundle projection, and $L : M \rightarrow \mathbb{R}$.

We would like to show that the above system arises from a Type II variational principle. We consider the action

$$S[\psi] = \int_{t_0}^{t_1} \psi^* (\Theta - H_L dt),$$

Fig. 5 Type II boundary conditions for adjoint systems



where $\psi : (t_0, t_1) \rightarrow T^*M$ is a curve on T^*M .

We would to place Type II boundary conditions, $q(t_0) = q_0$ and $p(t_1) = p_1$, on the variational principle. However, Type II boundary conditions for Hamiltonian systems, in general, suffer the drawback that they do not make intrinsic sense on a manifold, since one cannot specify a covector $p(t_1) = p_1$ without specifying the basepoint $q(t_1)$. Fortunately, for Hamiltonian systems which are adjoint systems, Type II boundary conditions do make intrinsic sense, due to the fact they cover an ODE on the base manifold M . To see this, if we fix the boundary condition $q(t_0) = q_0$, the time $t_1 - t_0$ flow of f , assuming it exists for this time, fixes the basepoint $q(t_1) = \Phi_{t_1-t_0}(q(t_0))$. In terms of the curve ψ , this means that once we fix $\pi_{T^*M}(\psi(t_0)) = q_0$, we have $\psi(t_1) \in T_{q(t_1)}^*M$, where $q(t_1) = \Phi_{t_1-t_0}(q(t_0))$. Thus, it then makes sense to specify a boundary condition on $\psi(t_1) \in T_{q(t_1)}^*M$ of the form $\psi(t_1) = p_1$, for any $p_1 \in T_{q(t_1)}^*M$. Fig. 5 illustrates Type II boundary conditions for an adjoint system; the flow of f on the base manifold evolves the initial condition q_0 forward to q_1 and subsequently, the vertical component of the lifted vector field X_{H_L} evolves the final momenta p_1 , based at q_1 , backwards to the initial momenta p_0 . As discussed in Sect. 3.1, p_1 can be chosen by taking $p_1 = dC|_{q_1}$ to compute the sensitivity of a terminal cost function $C : M \rightarrow \mathbb{R}$ with a nonaugmented Hamiltonian $H_L = H$ or by taking $p_1 = 0$ to compute the sensitivity of a running cost function L with an augmented Hamiltonian $H_L = H + L$.

Remark C.1 It is interesting to note that the reason for which Type I boundary conditions for adjoint systems are generally inconsistent (namely, that they cover an ODE on the base manifold) is precisely the reason that one can make intrinsic sense of Type II boundary conditions for adjoint systems. That is, Type II boundary conditions are consistent while Type I boundary conditions are generally inconsistent precisely because an adjoint system is a Hamiltonian system which covers an ODE on the base manifold. Conversely, every Hamiltonian system on T^*M which covers an ODE on the base manifold M is locally an adjoint system. To see this, if a Hamiltonian system covers an ODE on the base manifold, then Hamilton’s equation in the position variable $\dot{q} = \partial H / \partial p$ must equal $f(q)$ for some vector field f on M . Thus, we have $\partial H / \partial p = f(q)$. Integrating this equation yields a coordinate expression for the

Hamiltonian

$$H(q, p) = \langle p, f(q) \rangle + L(q),$$

where the “constant of integration” (constant with respect to the p variable) $L(q)$ is some arbitrary function of q . This is precisely the form of the Hamiltonian for an augmented adjoint system.

To state an intrinsic Type II variational principle for adjoint systems, we regard the integrand of the above action (before pulling back by ψ) as a contact form on the extended phase space $I \times T^*M$. Namely, given an interval $I = (t_0, t_1) \subset \mathbb{R}, t_0 \neq t_1$, let $\pi_{I \times T^*M} : I \times T^*M \rightarrow T^*M$ denote the projection onto the second factor. Then, define the contact form

$$\Theta_H = \pi_{I \times T^*M}^* \Theta - H dt,$$

where we have identified $H : T^*M \rightarrow \mathbb{R}$ with its pullback through $\pi_{I \times T^*M}$. In coordinates, $\Theta_H(q, p) = pdq - H dt$. Additionally, we define the presymplectic form $\Omega_H = -d\Theta_H$. Furthermore, we identify curves on T^*M , of the form $\psi : I \rightarrow T^*M$, with curves on $I \times T^*M$ which cover the identity on I ; in coordinates, this identification reads $\psi(t) = (t, q(t), p(t))$. The above action can then be expressed as

$$S[\psi] = \int_I \psi^* \Theta_H.$$

To enforce Type II boundary conditions $\pi_{T^*M}(\psi(t_0)) = q_0 \in M$ and $\psi(t_1) = p_1 \in T_{q_1}^*M$ where $q_1 = \Phi_{t_1-t_0}(q_0)$, we define the space of admissible variations with respect to these boundary conditions as the space of vector fields X on T^*M (identified with vertical vector fields on $I \times T^*M \rightarrow T^*M$) such that $(T\pi_{T^*M}X)(q_0) = 0$ and $X(\psi_1) = 0$, where $\psi_1 = (q_1, p_1) \in T_{q_1}^*M$. Intuitively, the first condition states that an admissible variation does not vary the initial position $q(t_0) = q_0$, whereas the second condition states that an admissible variation does not vary the final momenta ψ_1 .

Proposition C.1 *Fix an interval $I = (t_0, t_1) \subset \mathbb{R}, t_0 \neq t_1$. Consider the above augmented Hamiltonian, where we assume that the time $t_1 - t_0$ flow of the vector field f exists. Let $q_0 \in M$ and let $p_1 \in T_{q_1}^*M$ where $q_1 = \Phi_{t_1-t_0}(q_0)$. Then, the augmented adjoint system with Type II boundary conditions*

$$\begin{aligned} \dot{q} &= f(q), \\ \dot{p} &= -[Df(q)]^* p - dL(q), \\ q(t_0) &= q_0, \\ p(t_1) &= p_1, \end{aligned}$$

is intrinsically given by the variational principle: enforce the stationarity of the action

$$S[\psi] = \int_I \psi^* \Theta_H$$

with respect to admissible variations.

Proof Let φ_ϵ denote the time- ϵ flow of an admissible variation X . Then, the variational principle for the action with respect to admissible variations is given by

$$\begin{aligned} 0 &= dS[\psi] \cdot X = \left. \frac{d}{d\epsilon} \right|_{\epsilon=0} S[\varphi_\epsilon \circ \psi] = \int_I \psi^* \left. \frac{d}{d\epsilon} \right|_0 \varphi_\epsilon^* \Theta_H = \int_I \psi^* \mathcal{L}_X \Theta_H \\ &= - \int_I \psi^* (i_X \Omega_H) + \int_I \psi^* d(i_X \Theta_H) = - \int_I \psi^* (i_X \Omega_H) + \int_I d(\psi^* i_X \Theta_H). \end{aligned}$$

Observe that the boundary term $\int_I d(\psi^* i_X \Theta_H) = (\psi^* i_X \Theta_H)(t_1) - (\psi^* i_X \Theta_H)(t_0)$ vanishes by the fact that X is an admissible variation since $(\psi^* i_X \Theta_H)(t) = \langle p(t), (T\pi_{T^*M}X)(q(t)) \rangle$. Hence, the stationarity condition is given by

$$\int_I \psi^* (i_X \Omega_H) = 0.$$

By the fundamental lemma of the calculus of variations, we have $\psi^* (i_X \Omega_H) = 0$, whose coordinate expression is precisely the adjoint system. \square

Remark C.2 In our definition of the space of admissible variations, we set the conditions that the variation at q_0 is purely vertical, $(T\pi_{T^*M}X)(q_0) = 0$, whereas at q_1 , we enforced that the variation is zero, $X(q_1, p_1) = 0$. In coordinates where

$$X = \delta q \frac{\partial}{\partial q} + \delta p \frac{\partial}{\partial p},$$

the first condition reads $\delta q_0 = 0$ and the second condition reads $\delta q_1 = 0, \delta p_1 = 0$. It would thus seem that we are enforcing an overdetermined set of three boundary conditions $q(t_0) = q_0, q(t_1) = q_1, p(t_1) = p_1$. However, the resolution is that the variations δq_0 and δq_1 are not independent; fixing one to zero sets the other one to zero, by virtue of the fact that the adjoint system covers an ODE on M . Thus, with the chosen variational principle, we are only setting two independent boundary conditions, $q(t_0) = q_0, p(t_1) = p_1$.

Furthermore, in the above proof, by looking at the coordinate expression of the boundary term,

$$\psi^* i_X \Theta_H \Big|_{t_0}^{t_1} = \langle p(t_1), \delta q(t_1) \rangle - \langle p(t_0), \delta q(t_0) \rangle,$$

we see that we only used $\delta q_0 = 0, \delta q_1 = 0$. We did not need that $\delta p_1 = 0$ for the boundary terms to vanish. However, without setting $\delta p_1 = 0$, we only have the system

$$\dot{q} = f(q),$$

$$\begin{aligned}\dot{p} &= -[Df(q)]^* p - dL(q), \\ q(t_0) &= q_0.\end{aligned}$$

Hence, this system is underdetermined; any curve $p(t)$ in the fibers of T^*M satisfying

$$\dot{p} = -[Df(q)]^* p - dL(q)$$

would suffice. Thus, to uniquely fix the system, we must also supply a boundary condition of the form $p(t_1) = p_1$. Thus, even though the condition $\delta p_1 = 0$ is not strictly necessary in the variational principle to derive the equations of motion, it is necessary to fix the curve $p(t)$ in the fibers that define the adjoint system with Type II boundary conditions.

Analogously the adjoint DAE system (2.17a)–(2.17d), for index 1 DAEs, can be derived by an intrinsic Type II variational principle, by considering variations V of the action

$$S[\psi] = \int_{t_0}^{t_1} p dq - (H(q, u, p, \lambda) + L(q, u)) dt$$

such that $T\Pi_q V|_{t_0} = 0$ and $T\Pi_{(q,p)} V|_{t_1} = 0$ where $\Pi_q : (q, u, p, \lambda) \mapsto q$ and $\Pi_{(q,p)} : (q, u, p, \lambda) \mapsto (q, p)$ are the canonical bundle projections on $\overline{T^*M}_d \oplus \Phi^*$.

References

- Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton (2008)
- Aguiar, M.A., Camponogara, E., Foss, B.: An augmented Lagrangian for optimal control of DAE systems: algorithm and properties. *IEEE Trans. Autom. Control* **66**(1), 261–266 (2021)
- Barbero-Liñán, M., Martín de Diego, D.: Retraction maps: a seed of geometric integrators. [arXiv:2106.00607](https://arxiv.org/abs/2106.00607) (2021)
- Barbero-Liñán, M., Martín de Diego, D.: Presymplectic integrators for optimal control problems via retraction maps. [arXiv:2203.00790](https://arxiv.org/abs/2203.00790) (2022)
- Benning, M., Celledoni, E., Ehrhardt, M.J., Owren, B., Schönlieb, C.-B.: Deep learning as optimal control problems: models and numerical methods. *J. Comput. Dyn.* **6**(2), 171–198 (2019)
- Berglund, N.: Perturbation Theory of Dynamical Systems. DEA (2007)
- Biegler, L.T.: Nonlinear Programming: Concepts, Algorithms, and Applications to Chemical Processes. Society for Industrial and Applied Mathematics, Philadelphia (2010)
- Bullo, F., Lewis, A. D.: Supplementary chapters for Geometric Control of Mechanical Systems (2014). <http://motion.mee.ucsb.edu/book-gcms/>
- Burby, J.W., Klotz, T.J.: Slow manifold reduction for plasma science. *Commun. Nonlinear Sci. Numer. Simul.* **89**, 105289 (2020)
- Cacuci, D.G.: Sensitivity theory for nonlinear systems. I. Nonlinear functional analysis approach. *J. Math. Phys.* **22**(12), 2794–2802 (1981)
- Cao, Y., Li, S., Petzold, L., Serban, R.: Adjoint sensitivity analysis for differential-algebraic equations: the adjoint DAE system and its numerical solution. *SIAM J. Sci. Comput.* **24**(3), 1076–1089 (2003)
- Cariñena, J.F., Ibort, L.A., Gomis, J., Román-Roy, N.: Applications of the canonical-transformation theory for presymplectic systems. *Il Nuovo Cimento B* (1971–1996) **98**(2), 172–196 (1987)
- Chen, Y., Trenn, S.: An approximation for nonlinear differential-algebraic equations via singular perturbation theory. [arXiv:2103.12146](https://arxiv.org/abs/2103.12146) (2021)

- de León, Manuel, Cortés, Jorge, de Diego, Martín, Martínez, Sonia: General symmetries in optimal control. *Rep. Math. Phys.* **53**(1), 55–78 (2004)
- Delgado-Téllez, M., Ibrort, A.: A panorama of geometrical optimal control theory. *Extracta Math.* **18**(2), 129–151 (2003)
- Echeverría-Enríquez, A., Marín-Solano, J., Muñoz-Lecanda, M.C., Román-Roy, N.: Geometric reduction in optimal control theory with symmetries. *Rep. Math. Phys.* **52**(1), 89–113 (2003)
- Giles, M.B., Pierce, N.A.: An introduction to the adjoint approach to design. *Flow Turbul. Combust.* **65**(3), 393–415 (2000)
- Gotay, M.J., Nester, J.M.: Presymplectic Lagrangian systems. I: the constraint algorithm and the equivalence theorem. *Ann. de l'I.H.P. Phys. Théorique* **30**(2), 129–142 (1979)
- Gotay, M.J., Nester, J.M., Hinds, G.: Presymplectic manifolds and the Dirac–Bergmann theory of constraints. *J. Math. Phys.* **19**(11), 2388–2399 (1978)
- Griewank, A.: A mathematical view of automatic differentiation. In: *Acta Numer.*, vol. 12, pp. 321–398. Cambridge University Press (2003)
- Ibragimov, N.H.: Integrating factors, adjoint equations and Lagrangians. *J. Math. Anal. Appl.* **318**(2), 742–757 (2006)
- Ibragimov, N.H.: A new conservation theorem. *J. Math. Anal. Appl.* **333**(1), 311–328 (2007)
- Leok, M., Ohsawa, T.: Variational and geometric structures of discrete Dirac mechanics. *Found. Comput. Math.* **11**(5), 529–562 (2011)
- Leok, M., Zhang, J.: Discrete Hamiltonian variational integrators. *IMA J. Numer. Anal.* **31**(4), 1497–1532 (2011)
- Li, S., Petzold, L.: Adjoint sensitivity analysis for time-dependent partial differential equations with adaptive mesh refinement. *J. Comput. Phys.* **198**(1), 310–325 (2004)
- Li, S., Petzold, L. R.: Solution adapted mesh refinement and sensitivity analysis for parabolic partial differential equation systems. In: Biegler, L. T., Heinkenschloss, M., Ghattas, O., van Bloemen Waanders, B., (eds) *Large-Scale PDE-Constrained Optimization*, pp. 117–132. Springer, Berlin (2003)
- Mattsson, S.E., Söderlind, G.: Index reduction in differential-algebraic equations using dummy derivatives. *SIAM J. Sci. Comput.* **14**(3), 677–692 (1993)
- Nguyen, V.T., Georges, D., Besançon, G.: State and parameter estimation in 1-D hyperbolic PDEs based on an adjoint method. *Automatica* **67**(C), 185–191 (2016)
- Pierce, N.A., Giles, M.B.: Adjoint recovery of superconvergent functionals from PDE approximations. *SIAM Rev.* **42**(2), 247–264 (2000)
- Reid, G.J., Lin, P., Wittkopf, A.D.: Differential elimination-completion algorithms for DAE and PDAE. *Stud. Appl. Math.* **106**(1), 1–45 (2001)
- Roche, M.: Implicit Runge–Kutta methods for differential algebraic equations. *SIAM J. Numer. Anal.* **26**(4), 963–975 (1989)
- Ross, I.M.: A roadmap for optimal control: the right way to commute. *Ann. NY Acad. Sci.* **1065**(1), 210–231 (2005)
- Ross, M., Fahroo, F.: A pseudospectral transformation of the convectors of optimal control systems. *IFAC Proc. Ser.* **34**(13), 543–548 (2001)
- Sanz-Serna, J.M.: Symplectic Runge–Kutta schemes for adjoint equations, automatic differentiation, optimal control, and more. *SIAM Rev.* **58**(1), 3–33 (2016)
- Schmitt, J.M., Leok, M.: Properties of Hamiltonian variational integrators. *IMA J. Numer. Anal.* **38**(1), 377–398 (2017)
- Sirkes, Z., Tziperman, E.: Finite difference of adjoint or adjoint of finite difference? *Mon. Weather Rev.* **125**(12), 3373–3378 (1997)
- Tran, B., Leok, M.: Multisymplectic Hamiltonian variational integrators. *Int. J. Comput. Math.* (Special Issue on Geometric Numerical Integration, Twenty-Five Years Later) **99**(1), 113–157 (2022)
- Vankerschaver, J., Yoshimura, H., Leok, M.: The Hamilton–Pontryagin principle and multi-Dirac structures for classical field theories. *J. Math. Phys.* **53**(7), 072903 (25 pages) (2012)
- Wang, Q., Duraisamy, K., Alonso, J.J., Iaccarino, G.: Risk assessment of scramjet unstart using adjoint-based sampling. *AIAA J.* **50**(3), 581–592 (2012)
- Yano, K., Ishihara, S.: *Tangent and Cotangent Bundles: Differential Geometry*. Pure Appl. Math., No. 16. Marcel Dekker, Inc., New York (1973)
- Yoshimura, H., Marsden, J.E.: Dirac structures in Lagrangian mechanics Part I: implicit Lagrangian systems. *J. Geom. Phys.* **57**(1), 133–156 (2006)

Yoshimura, H., Marsden, J.E.: Dirac structures in Lagrangian mechanics Part II: variational structures. *J. Geom. Phys.* **57**(1), 209–250 (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.