Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnline CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 54-5 (2021) 181–186

Smooth Approximations for Hybrid Optimal Control Problems with Application to Robotic Walking

Tyler Westenbroek * **Xiaobin Xiong** ** **S Shankar Sastry** *
Aaron D. Ames **
* *University of California, Berkeley*
** *California Institute of Technology*

Abstract: This paper investigates optimal control problems formulated over a class of hybrid dynamical systems which display event-triggered discrete jumps. Due to the discontinuous nature of the underlying dynamics, previous approaches to solving optimal control problems over this class of systems generally rely on fixing the number and sequence of discrete jumps *a priori*, or search over possible mode sequences in a combinatorial manner. Employing contributions from the geometric theory of hybrid systems, we instead formulate a family of smooth approximate problems formulated over a family of smooth control systems which faithfully approximate the dynamics of the original hybrid system, in an appropriate metric. Efficient gradient-based methods can be used to solve the smooth approximations, without specifying the sequence of discrete transitions ahead of time. Under appropriate hypothesis, the gradients of the smooth problem are shown to be well-conditioned and closely approximate the gradients of the non-smooth problem (when they exist). Two cases studies demonstrate the utility of the approach, including an in-depth application to generating a stable walking motion for a bipedal robot.

Copyright © 2021 The Authors. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0)

1. INTRODUCTION

Hybrid dynamical systems represent a powerful, expressive framework for modeling physical systems which undergo sudden changes in their dynamics due to impulsive state-triggered events. Indeed, hybrid models have been used extensively in application domains ranging from bipedal robotic walking (Grizzle et al. (2014)) to power systems (Hiskens and Pai (2000)). Given the broad applicability of the framework, the generation of reliable numerical algorithms solving optimal control problems over hybrid models is of great practical interest.

Necessary and sufficient conditions for optimality in hybrid optimal control problems have been given by considering extensions to classic Dynamic Programming techniques (Branicky et al. (1998), Lygeros et al. (1999), Caines et al. (2007), Pakniyat and Caines (2017), as well as Pontryagin's Maximum Principle (Shaikh and Caines (2007), Sussmann (1999)). However, due to the non-smooth nature of hybrid systems, traditional approaches for systematically translating optimality conditions into reliable numerical optimal control algorithms (Polak (2012)) cannot be directly applied to this setting. Consequently, implementable derivative-based algorithms for solving hybrid optimal control problems to local optimality generally fix the number and sequence of discrete transitions *a priori* (Hereid et al. (2018), Shaikh and Caines (2007)), restricting the behaviors the optimization can produce. Rather than attempting to solve this difficult problem directly, this paper instead introduces a family of smooth approximations to optimal control problems formulated over a restricted, but still widely studied, class of hybrid systems with state-triggered jumps. These relaxations are formulated over the smooth approximations to the hybrid dynamics introduced in Westenbroek et al. (2018) and can be solved with standard derivative-based numerical optimal control algorithms without pre-specifying a desired sequence of discrete jumps.

Our approach is rooted in the long-building geometric theory for hybrid dynamical systems first proposed in Simic et al. (2005). Under a number of regularity assumptions about the hybrid system – most notably that the reset dynamics are invertible – this work demonstrated that the trajectories of a hybrid system with state-triggered jumps can be represented using the integral curves of a piecewise smooth vector field defined on a smooth manifold or *hybridfold*. The underlying manifold structure is obtained by applying a topological gluing construction which identifies each point in a guard set with its image under the associated reset map, effectively 'removing' the discrete jumps from the dynamics of the system. Homotopically meaningful generalizations of this were considered in Ames and Sastry (2005), which adds relaxed 'strips' which preserve topological properties of the hybrid system under the aforementioned gluing constructions. This has parallels to the practical regularizations in Johansson et al. (1999), in which added strips are used to effectively delay the instantaneous nature of hybrid transitions. The recent contribution Westenbroek et al. (2018) smooths the hybrid dynamics along the added strip using techniques from singular perturbation theory (Sotomayor (1996)) to produce a family of smooth control systems whose dynamics uniformly approximate the original hybrid dynamics in the state-space metric from Burden et al. (2015).

This paper advances this theory by demonstrating that these smooth approximations retain important variational

2405-8963 Copyright © 2021 The Authors. This is an open access article under the CC BY-NC-ND license.
Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2021.08.495

properties of the original discontinuous dynamics. In the context of optimal control, these results show that the gradients of our smooth approximate problems converge to the gradients of the original hybrid optimal control problem (when they exist) in the appropriate limit. This result is foundational for applying well-established numerical optimal control frameworks (Polak (2012)) to provably find approximate minimizers of the original hybrid optimal control problem using our relaxations. We apply the approach in two in-depth numerical case studies including the generation of stable walking motion for a bipedal robot, by far the most complicated system to which these folding constructions have been applied. Due to space constraints, in this paper we focus primarily on relating our main theoretical developments, leaving a number of practical implementation details to the technical report Westenbroek et al. (2021), where the interested reader can also find formal proofs for our claims and a lengthy discussion of avenues for future work.

Preliminaries: Given a set $U \subset \mathbb{R}^n$, we denote the set of square integrable functions into U with $L^2([0,T],U)$. Given a map $\mathcal{L}\colon V \to W$ between Hilbert spaces V, W and $u, \delta u \in V$, the directional derivative of $\mathcal{L}$ at the point u in the direction δu is $D\mathcal{L}(u;\delta u) = \lim_{\lambda \to 0^+} \frac{\mathcal{L}(u+\lambda\delta u)-\mathcal{L}(\delta u)}{\lambda}$, if the preceding limit exists. Given a smooth manifold with boundary M, ∂M will denote the boundary and $int(M)$ will denote the interior. Given two smooth manifolds M, N and a smooth map $G\colon M \to N$, we denote the associated *pushforward* of G by $DG\colon TM \to TN$.

This paper will make extensive use of the notion of 'gluing constructions' or 'quotient topologies'. A comprehensive introduction to the topic can be forund in (Kelley, 2017, Chapter 3), while examples of how the topic will be applied here can be found in Simic et al. (2005) and Westenbroek et al. (2018). Given a topological space D and an equivalence relation $\sim \subset D \times D$, we let $D/\sim$ denote the corresponding set of equivalence classes. There exists a natural *quotient projection* $\pi\colon D \to D/\sim$ taking each point $x \in D$ to its equivalence class $[x] \in D/\sim$. Each map $R\colon G \to D$ with $G \subset D$ induces an equivalence relation $\overset{R}{\sim} = \{(x,y) \in D \times D\colon x = y,\ x \in R^{-1}(y) \text{ or } y \in R^{-1}(x)\}$. To emphasize the dependence on R, we will denote the corresponding quotient space by $D/\sim = D/(G \overset{R}{\sim} R(G))$.

2. (RELAXED) HYBRID DYNAMICAL SYSTEMS

2.1 Hybrid Dynamical Systems

To simplify exposition, throughout the paper we will consider hybrid systems with a single continuous domain, guard and reset map. As discussed in Westenbroek et al. (2021), the method naturally extends to hybrid systems with multiple discrete modes so long as the guard sets do not intersect and each reset map is invertible.

Definition 1. $\mathcal{H} = (D, U, f, G, R)$ is a *hybrid dynamical system* where

(1) $D \subset \mathbb{R}^n$ is a bounded n-dimensional manifold;
(2) $U \subset \mathbb{R}^m$ is a compact set of allowable inputs;
(3) $f\colon D \times U \to T\mathbb{R}^n$ is a smooth controlled vector field;
(4) $G \subset \partial D$ is an embedded $(n-1)$-dimensional submanifold referred to as the guard;

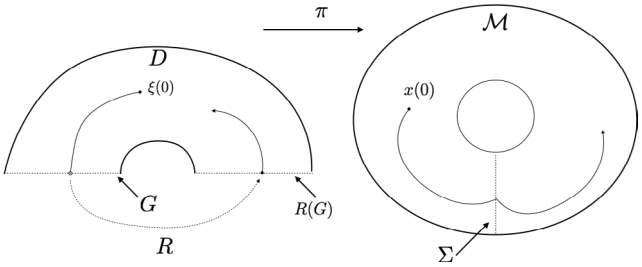

Fig. 1. (left) The continuous domain D. A hybrid executions ξ undergoes a single discrete jump. (right) The hybrid manifold and the curve $x = \pi \circ \xi$, which is an integral curve of F.

(5) $R\colon G \to \partial D$ is a smooth reset map.

Before defining the executions of $\mathcal{H}$, we first discuss several additional assumptions. The first assumption ensures that discrete jumps are isolated, precluding pathologies such as Zeno executions and, together with Assumption 1, ensures that the gluing constructions we utilize lead to a smooth manifold structure, which ensures that the underlying topology of the hybrid optimal control problems we consider are well-posed.

Assumption 1. The guard and its image under the reset do not overlap, i.e., $G \cap R(G) = \emptyset$.

Assumption 2. The reset map R is a diffeomorphism onto its image.

Although Assumption 2 is somewhat restrictive, a number of important hybrid models such as the bipedal robot investigated in Section 4 satisfy this condition. The following Assumption ensures that trajectories do not skim the guard set, a standard condition which supports the uniqueness of hybrid executions:

Assumption 3. The normal component of the vector field f along ∂D is non-zero. Moreover, f points out of D along G and into the interior of D along $R(G)$.

Our final Assumption is standard even for optimal control problems defined over classical control systems:

Assumption 4. The vector field f is Lipschitz continuous.

The above assumptions ensure the uniqueness of the following solution concept (Lygeros et al. (2003)):

Definition 2. Given $(\xi_0, u) \in D \times L^2([0,T],U)$, the right-continuous function $\xi\colon [0,T] \to D$ is an *execution* of $\mathcal{H}$ of length $T > 0$ corresponding to (ξ, u) if $\xi(0) = \xi_0$ and

$$\begin{cases} \dot{\xi} = f(\xi(t), u(t)) & \text{if } \xi(t^-) \notin G \\ \xi^+(t) = R(\xi(t^-)) & \text{if } \xi(t^-) \in G. \end{cases} \tag{1}$$

2.2 The Hybrid Manifold

The discontinuities in the executions of $\mathcal{H}$ make it difficult to directly apply classical optimal control techniques to this class of systems. Thus, we will instead work with an equivalent representation of the hybrid dynamics wherein the executions are represented using the continuous integral curves of a piecewise smooth controlled vector field

defined on a smooth manifold. In Section 2.3 we then construct smooth approximations to these dynamics. Our description of these constructions is brief, as they are covered extensively elsewhere (Ames and Sastry (2005); Westenbroek et al. (2018); Simic et al. (2005)).

The main idea behind the re-topologization of the hybrid system $\mathcal{H}$ is to 'glue' each point $x \in G$ to its image under the reset map, $R(x) \in R(G)$. This process is depicted in Figure 1. More formally, we define the *Hybrid Manifold*

$$\mathcal{M} := \frac{D}{G \overset{R}{\sim} R(G)}$$

which is defined by collapsing $x \in G$ and $R(x) \in R(G)$ to a single point. Letting $\pi \colon D \to \mathcal{M}$ denote the quotient projection taking each point $x \in D$ to its equivalence class in $x \in \mathcal{M}$, this process amounts to identifying the surface G with the surface $R(G)$ to produce the single surface

$$\Sigma := \pi(G) = \pi(R(G)) \subset \mathcal{M}.$$

As depicted in Figure 1, each execution $\xi \colon [0, T] \to D$ of $\mathcal{H}$ is represented by a continuous curve on $\mathcal{M}$, namely, $x = \pi \circ \xi \colon [0, T] \to \mathcal{M}$. In effect, by identifying G with $R(G)$ the gluing construction 'removes' the discrete jumps in the executions of $\mathcal{H}$.

The term 'hybrid manifold' or 'hybridfold' has been justified by a number of previous works (Simic et al. (2005); Westenbroek et al. (2018)) which have demonstrated that it is possible to endow $\mathcal{M}$ with the structure of a smooth topological manifold by constructing a smoothly compatible atlas of coordinate charts $\{(U_\alpha, \varphi_\alpha)\}_{\alpha \in \mathcal{A}}$ for the space. The specific form of these coordinates is not needed for the statement of the theoretical results presented in this manuscript, but is critical for successful numerical implementations of the approach and is discussed further in Westenbroek et al. (2021).

As demonstrated by Simic et al. (2005), Westenbroek et al. (2018), the dynamics of $\mathcal{H}$ can be generated using the integral curves of the piece-wise smooth vector field $F \colon \mathcal{M} \times \mathbb{R}^m \to T\mathcal{M}$ which is defined by

$$F(\pi(x), u) = D\pi \circ f(x, u),$$

for each $x \in int(D)$ and $u \in U$. Intuitively, F is simply the representation of f on the new space $\mathcal{M}$. In general, F will be discontinuous and undefined along Σ. Formally, the relationship between $\mathcal{H}$ and F are encapsulated by the following result from (Westenbroek et al. (2018)):

Lemma 1. Let $\xi_0 \in D$ and $u \in L^2([0, T], U)$, and let $\xi \colon [0, T] \to D$ be the corresponding hybrid execution. Further let $x \colon [0, T] \to \mathcal{M}$ be given by $x = \pi \circ \xi$. Then x satisfies $\frac{d}{dt} = F(x(t), u(t))$ for almost every $t \in [0, T]$.

Specifically, given an execution ξ of $\mathcal{H}$, the instances of time when ξ undergoes a discrete jump correspond to time instances when the curve $x = \pi \circ \xi$ crosses the surface Σ and experiences a discontinuity in F.

2.3 Relaxed Hybrid Dynamics

While $\mathcal{M}$ and F provide a compact representation of the original hybrid dynamics, the piecewise continuous nature of F still prevents the direct application of classical derivative-based search algorithms to optimal control

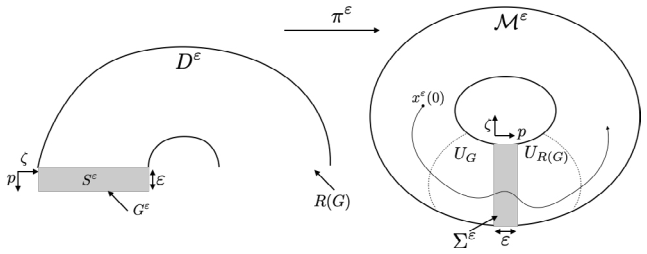

Fig. 2. (left) The relaxed domain D and its components. (right) The relaxed hybrid manifold. An integral curve x^ε of F^ε flows through the relaxed strip.

problems formulated over this space. Thus, we now introduce a family of relaxations to $\mathcal{M}$ and F which result in a smooth control system which accurately approximates the original discontinuous dynamics. This construction, which is depicted in Figure 2.

Formally, we construct this new space by defining for each positive regularization parameter $\varepsilon > 0$ the ε-thick strip

$$S^\varepsilon = G \times [0, \varepsilon].$$

and then attach the strip to G to obtain

$$D^\varepsilon = \frac{D \coprod S^\varepsilon}{G \sim S^\varepsilon},$$

which is depicted in Figure 2. We will use the coordinates $(\zeta, p) \in G \times [0, \varepsilon]$ to refer to points in the relaxed strip. Finally, we attach the relaxed strip to $R(G)$ defining the *relaxed hybrid manifold* by

$$\mathcal{M}^\varepsilon = \frac{D^\varepsilon}{G \times \{\varepsilon\} \overset{R^\varepsilon}{\sim} R(G)},$$

where $R^\varepsilon \colon G \times \{\varepsilon\} \to R(G)$ is defined by $R^\varepsilon(\zeta, \varepsilon) = R(\zeta)$ for each $(\zeta, \varepsilon) \in G \times \{\varepsilon\}$. We will let $\pi^\varepsilon \colon D^\varepsilon \to \mathcal{M}^\varepsilon$ denote the quotient projection taking points from D^ε to $\mathcal{M}^\varepsilon$. It was shown in Westenbroek et al. (2018) that $\mathcal{M}^\varepsilon$ can also be endowed with a smooth manifold structure.

While $\mathcal{M}$ is constructed by 'gluing' G to $R(G)$ to form the co-dimension-1 sub-manifold $\Sigma \subset \mathcal{M}$, intuitively, $\mathcal{M}^\varepsilon$ is then formed by 'pulling' G and $R(G)$ apart to create the ε-thick strip

$$\Sigma^\varepsilon = \pi^\varepsilon(S^\varepsilon) \subset \mathcal{M}^\varepsilon$$

now separating the two surfaces, which are represented by the sets $\pi^\varepsilon(G), \pi^\varepsilon(R(G)) \subset \mathcal{M}^\varepsilon$. Note that the set $\mathcal{M} \setminus \Sigma$ is a proper subset of $\mathcal{M}^\varepsilon$. Specifically, we can decompose $\mathcal{M}^\varepsilon = (\mathcal{M} \setminus \Sigma) \cup \Sigma^\varepsilon$. Indeed, the smooth vector field $F^\varepsilon \colon \mathcal{M}^\varepsilon \times U \to T\mathcal{M}^\varepsilon$ we construct is of the form

$$F^\varepsilon(x, u) = \begin{cases} F(x, u) & \text{if } x \in \mathcal{M} \setminus \Sigma \\ \tilde{F}^\varepsilon(x, u) & \text{if } x \in \Sigma^\varepsilon \end{cases}$$

where $\tilde{F}^\varepsilon \colon \Sigma^\varepsilon \times U \to T\Sigma^\varepsilon$ interpolates between the discontinuities in F. To construct $\tilde{F}^\varepsilon$, we first let $U_G \subset \mathcal{M}^\varepsilon$ be a neighborhood containing $\pi^\varepsilon(G)$ and let $U_{R(G)} \subset \mathcal{M}^\varepsilon$ be a neighborhood containing $\pi^\varepsilon(R(G))$ as in Figure 2. Next, let $\tilde{F}_1 \colon U_G \cup \Sigma^\varepsilon \times U \to T\mathcal{M}^\varepsilon$ smoothly extend F_1 onto Σ^ε and let $\tilde{F}_2 \colon U_G \cup \Sigma^\varepsilon \times U \to T\mathcal{M}^\varepsilon$ smoothly extend $\tilde{F}_2$ onto Σ^ε. We then define the vector field on the relaxed strip in the coordinates $(\zeta, p) \in G \times [0, \varepsilon]$ by

$$\tilde{F}^\varepsilon((\zeta, p), u) = (1 - \alpha^\varepsilon(p))\tilde{F}_1((\zeta, p), u) + \alpha^\varepsilon(p)\tilde{F}_2((\zeta, p), u),$$

where for each $\varepsilon > 0$ we define $\alpha^\varepsilon \colon [0, \varepsilon] \to [0, 1]$ by $\alpha^\varepsilon(\tau) = \alpha(\frac{\tau}{\varepsilon})$ where $\alpha \colon [0, 1] \to [0, 1]$ is any smooth

function such that *i)* $\alpha(0) = 0$, *ii)* $\alpha(1) = 1$, *iii)* α is monotonically increasing and *iv)* α has bounded first and second derivatives. That is, $\tilde{F}^\varepsilon$ smoothly interpolates between $\tilde{F}_1$ and $\tilde{F}_2$ (which represent the dynamics of F on either side of the strip).

Remark 1. It was shown in (Westenbroek et al., 2018, Theorem 6) that F^ε is a smooth vector field.

Remark 2. We can represent an execution $\xi\colon [0,T] \to D$ of $\mathcal{H}$ with initial condition $\xi(0) = \xi_0$ as a piecewise continuous curve $x = \pi^\varepsilon \circ \xi\colon [0,T] \to \mathcal{M}^\varepsilon$. The corresponding solution to the relaxed dynamics is given by the curve $x^\varepsilon\colon [0,T] \to \mathcal{M}^\varepsilon$ which is the integral curve of F^ε starting from the initial condition $x^\varepsilon(0) = \pi^\varepsilon(\xi_0)$ under the application of the same input used to generate ξ. In (Westenbroek et al., 2018, Theorem 7) it was demonstrated that $\sup_{t\in[0,T]} \tilde{d}_{\mathcal{M}^\varepsilon}(x(t), x^\varepsilon(t)) \to 0$ as $\varepsilon \to 0$, where $\tilde{d}_{\mathcal{M}^\varepsilon}$ is the metric for $\mathcal{M}^\varepsilon$ introduced in Burden et al. (2015). In other words, the solutions to the smooth dynamics converge to the executions of the original hybrid system when they are both represented on $\mathcal{M}^\varepsilon$ and the relaxation parameter is taken to zero.

3. OPTIMAL CONTROL

3.1 Hybrid Optimal Control Problems

We consider optimal control problems with a fixed initial condition $\xi_0 \in D$, and for each $u \in L^2([0,T], U)$ we will let $\xi^u\colon [0,T] \to D$ denote the hybrid execution corresponding to this data. We define the *Hybrid Optimal Control Problem* (HOCP):

$$\inf_{u\in L^2([0,T],U)} \ell(u) = \int_0^T \ell_r(\xi^u(t))dt + \ell_f(\xi^u(T)),$$

where $\ell_r\colon D \times U \to \mathbb{R}$ is the running cost and $\ell_f\colon D \to \mathbb{R}$ is the terminal cost. For simplicity, we assume that ξ^u exists for each $u \in L^2([0,T], U)$ so that $\ell(u)$ is always well-defined and assume that both cost functions are continuously differentiable.

It is well-known Caines et al. (2007) that the transversality conditions in Assumption 3 ensure that $\ell(\cdot)$ is continuous and differentiable in a neighborhood of each $u \in L^2([0,T], U)$ such that $\xi^u(T^-) \notin G$. However, when $\xi^u(T^-) \in G$ nearby executions may undergo a different number of discrete jumps in which case ℓ may be discontinuous at this point.

3.2 Recasting the Problem on the Hybrid Manifold

We may reformulate the HOCP as an optimal control problem defined on the hybrid manifold as follows:

$$\inf_{u\in L^2([0,T],U)} \mathcal{L}(u) = \int_0^T L_r(x^u(t), u(t))dt + L_f(x^u(T))$$

where $x^u\colon [0,T] \to \mathcal{M}$ is the integral curve of F corresponding to the data $(\pi(\xi_0), u)$, the running cost $L_r\colon \mathcal{M} \times U \to \mathbb{R}$ is defined by $L_r(x,u) = \ell_r(\pi^{-1}(x), u)$ for each $(x,u) \in D \times U$ and the terminal cost $L_f\colon \mathcal{M} \to \mathbb{R}$ is defined by $L_f(x) = \ell_f(\pi^{-1}(x))$ for each $x \in D$. Since π^{-1} is multi-valued for each $x \in \Sigma$, the terminal and final costs will in general be multi-valued and discontinuous at these points. Thus, recasting the HOCP on $\mathcal{M}$ highlights the underlying degeneracy of the original HOCP. The relaxed optimal control problems we formulate below resolve these issues by producing smooth approximations to these costs over our family of smooth vector fields.

3.3 Relaxed Optimal Control Problems

We now construct a smooth approximation to $\mathcal{L}(\cdot)$ using smooth approximations to L_r and L_f defined on $\mathcal{M}^\varepsilon$ and the smooth vector field F^ε. For each $\varepsilon > 0$ the smooth approximate running cost $L_r^\varepsilon\colon \mathcal{M}^\varepsilon \times U \to \mathbb{R}$ and terminal cost $L_f^\varepsilon\colon \mathcal{M}^\varepsilon \to \mathbb{R}$ are of the form

$$L_r^\varepsilon(x,u) = \begin{cases} L_r(x,u) & \text{if } x \in \mathcal{M} \setminus \Sigma \\ \tilde{L}_r^\varepsilon(x,u) & \text{if } x \in \Sigma^\varepsilon \end{cases} \tag{2}$$

$$L_f^\varepsilon(x) = \begin{cases} L_f(x) & \text{if } x \in \mathcal{M} \setminus \Sigma \\ \tilde{L}_f^\varepsilon(x) & \text{if } x \in \Sigma^\varepsilon \end{cases} \tag{3}$$

where $\tilde{L}_r^\varepsilon\colon \Sigma^\varepsilon \times U \to \mathbb{R}$ and $\tilde{L}_f^\varepsilon\colon \Sigma^\varepsilon \to \mathbb{R}$ are constructed analogously to how $\tilde{F}^\varepsilon$ was constructed in Section 2.3 by interpolating between the discontinuities in L_r and L_f. With these smooth costs in hand, for each $\varepsilon > 0$ we define the *relaxed hybrid optimal control problem* (RHOCP)

$$\min_{u\in L^2([0,T],U)} \mathcal{L}^\varepsilon(u) = \int_0^T L_r^\varepsilon(x^{\varepsilon,u}(t), u(t))dt + L_f^\varepsilon(x^{\varepsilon,u}(T)),$$

where $x^{\varepsilon,u}\colon [0,T] \to \mathcal{M}^\varepsilon$ is the integral curve of F^ε corresponding to the data $(\pi(\xi_0), u)$. The following Lemma demonstrates that the relaxed cost functional well-approximates the original hybrid cost functional at points where it is continuous, and largely follows from (Westenbroek et al., 2018, Theorem 7):

Lemma 2. Let our standing assumptions hold. If $u \in L^2([0,T], U)$ is s.t. $x^u(T) \notin \Sigma$ then $\lim_{\varepsilon\to 0^+} \mathcal{L}^\varepsilon(u) = \mathcal{L}(u)$.

Moreover, since F^ε, L_r^ε and L_f^ε are smooth the directional derivatives of $\mathcal{L}^\varepsilon$ are well-defined and characterized by a well-behaved adjoint equation:

Lemma 3. Let our standing assumptions hold. For each $u \in L^2([0,T], U)$ and $\delta u \in L^2([0,T], \mathbb{R}^m)$ we have

$$D\mathcal{L}^\varepsilon(u;\delta u) = \int_0^T \langle \lambda(t)B(t)^T + \tfrac{\partial}{\partial u}\ell_r^\varepsilon(x^{\varepsilon,u}(t), u(t)), \delta u(t)\rangle dt,$$

where the adjoint process $\lambda\colon [0,T] \to T^*\mathcal{M}^\varepsilon$ satisfies the co-variational equation defined, in local coordinates, by

$$\dot{\lambda}(t) = -\lambda(t)A(t) - \tfrac{\partial}{\partial x}\ell_r^\varepsilon(x^{\varepsilon,u}(t), u(t)), \tag{4}$$

with the terminal condition $\lambda(T) = \frac{d}{dx}\ell_f^\varepsilon(x^{\varepsilon,u}(T))$, where $A(t) = \frac{\partial}{\partial x} f^\varepsilon(x^{\varepsilon,u}(t), u(t))$ and $B(t) = \frac{\partial}{\partial u} f^\varepsilon(x^{\varepsilon,u}(t), u(t))$, with f^ε the local representation of F^ε, ℓ_r^ε the local representation of L_r^ε and ℓ_f^ε the local representation of L_f^ε.

It is well-known (see e.g. Caines et al. (2007)) that the adjoint process associated to the original HOCP is discontinuous at time instances when the nominal hybrid execution undergoes a discontinuous jump. The proof of the following theorem demonstrates that the relaxed adjoint equations recover these jumps as $\varepsilon \to 0^+$ at points where the original hybrid cost is differentiable (i.e. when $x^u(T) \notin \Sigma$). This result, which is significantly stronger than Lemma 2, indicates that the first-order variational properties of F^ε agree with those of the original hybrid system. Indeed, Lemma 2 is insufficient to establish that local

minimizers of the RHOCP coincide with local minimizers of HOCP as the relaxation parameter tends to zero, as results of this sort generally require the stronger condition that the optimality conditions of the two problems agree in the limit (Polak (2012)). Thus the following result is foundational for the systematic design of algorithms which provably find approximate local minimizers of the original hybrid optimal control problem using our relaxations:

Theorem 1. Let our standing assumptions hold. Further assume that $u \in L^2([0,T],U)$ is such that $x^u(T) \notin \Sigma$. Then for each $\delta u \in L^2([0,T],\mathbb{R}^m)$ we have:

$$\lim_{\varepsilon \to 0^+} D\mathcal{L}^\varepsilon(u;\delta u) = D\mathcal{L}(u;\delta u).$$

Future work will aim to extend well-established numerical frameworks (Polak, 2012, Chapter 4) for finding approximate minimizers of continuous-time optimal control problems to the current setting using this result.

4. APPLICATIONS

We now apply our approach to two mechanical systems which repeatedly undergo impacts. A more in depth discussion of the models and practical numerical implementations of the smoothing approach is provided in the technical report.

4.1 Robotic Bouncing Ball

We first consider a robotic bouncing ball with horizontal thrust forces to drive the ball to different locations, which is bouncing on the sinusoidal surface depicted in Figure 3. The ball has two degrees of freedoms (DoFs). Its states are described as $x = [b_z, b_x, \dot{b}_z, \dot{b}_x]^T$, with b_z, b_x being the vertical and horizontal positions. The continuous dynamics are defined by $f(x,u) = [\dot{b}_z, \dot{b}_x, g, u]^T$, where $u \in \mathbb{R}^1$ is the horizontal thrust force on the ball and g is the gravitational constant. The ground profile yields a unilateral constraint of the form $h(b_x, b_z) = 0$. We use the impact model from Or and Ames (2010) so the positions remain unchanged

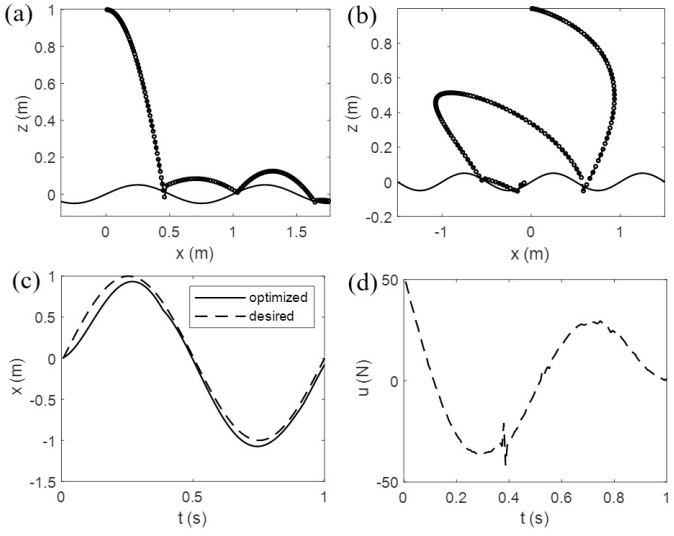

Fig. 3. (a) The trajectory of the ball with $u = 0$, (b) the optimized trajectory of the ball, (c) the horizontal trajectories, and (d) the optimized input.

during impacts but the velocities $\dot{b} = [\dot{b}_x, \dot{b}_z]$ are reset according to $\dot{b}^+ = \dot{b}^- - (1 + r_c)\frac{\dot{b}^- J_h^T}{J_h J_h^T} J_h^T$, where J_h is the Jacobian of h and $r_c \in (0,1)$ is the coefficient of restitution.

We consider the problem of driving the ball to follow a desired horizontal trajectory $b_x^{\text{desired}}(t)$ from an initial state $x_0 = [1,0,0,1]^T$. Thus, the cost function is: $\ell(x,t) = \int_{t=0}^{T} |b_x(t) - b_x^{\text{desired}}(t)|^2 dt$, where $T = 1$s is the time horizon. We set $r_c = 0.4$, $\varepsilon = 10^{-3}$, and use forward Euler integration with a time step of 10^{-3}s for the discretization. The initial guess of the inputs u is set to be zeros. Fig. 3 shows the results, where the optimization produces a markedly different trajectory than the initialization and closely tracks the desired behavior while undergoing a different number of impacts than the solution produced by the initial guess.

4.2 Bipedal Robotic Walker

Next we use our approach to generate a stable multi-step walking behavior on the planar under-actuated five-linkage bipedal robot shown in Fig. 4 (a). We use the planar single-domain hybrid model from Grizzle et al. (2014) in which during each phase of continuous evolution the *swing foot* is above the ground while the *stance foot* is assumed to remain motionless on the ground. At the end of each step (when the swing foot impacts the ground) the swing foot and stance foot are swapped and the new swing foot is assumed to lift off from the ground instantaneously.

Continuous Dynamics: The configuration variables for the system are $q = (q_1, q_2, q_3, q_4, \theta) \in Q \subset R^5$ whose dynamics can be derived using the method of Lagrange:

$$M(q)\ddot{q} + C(q,\dot{q}) + G = Bu$$

where M is the mass matrix, C is the Coriolis and centrifugal term, G is the gravitational term, B and $u \in \mathbb{R}^4$ are the actuation matrix and the motor torque vector. The overall state variable is $x = (q, \dot{q}) \in TQ$. The continuous domain for the system is defined by $D = \{x \in TQ \times \mathbb{R}^2 \colon z_{sw}(q) \geq 0\}$, where $z_{sw}(q)$ gives the height of the swing foot.

Impact Assumption: The impact between the foot and the ground is assumed to be plastic. The guard for the system is defined by $G = \{x \in D \colon z_{sw}(q) = 0, \frac{d}{dq} z_{sw}(q) \cdot \dot{q} < 0\}$, the set of states where the swing foot impacts the ground. At impact the joint velocities are reset via

$$\dot{q}^+ = \Delta(q)\dot{q}^-,$$

where the exact expression of $\Delta(q)$ can be found in Grizzle et al. (2014). The reset map $R \colon G \to R(G)$ is given by $R(q,\dot{q}) = (\hat{R}q, \hat{R}\Delta(q)\dot{q})^T$, where $\hat{R}$ is a relabeling matrix used to swap the stance foot and swing foot.

Optimal Control Problem: The goal of the optimization is for the robot to smoothly accelerate from rest and terminate the optimized trajectory on a pre-designed periodic walking gate which was constructed using the Hybrid Zero Dynamics approach (Grizzle et al. (2014)). Methods for generating periodic walking motions generally do not prescribe a method for reaching the periodic gate from rest. Thus, solving this task immediately compliments established tools from the dynamic walking

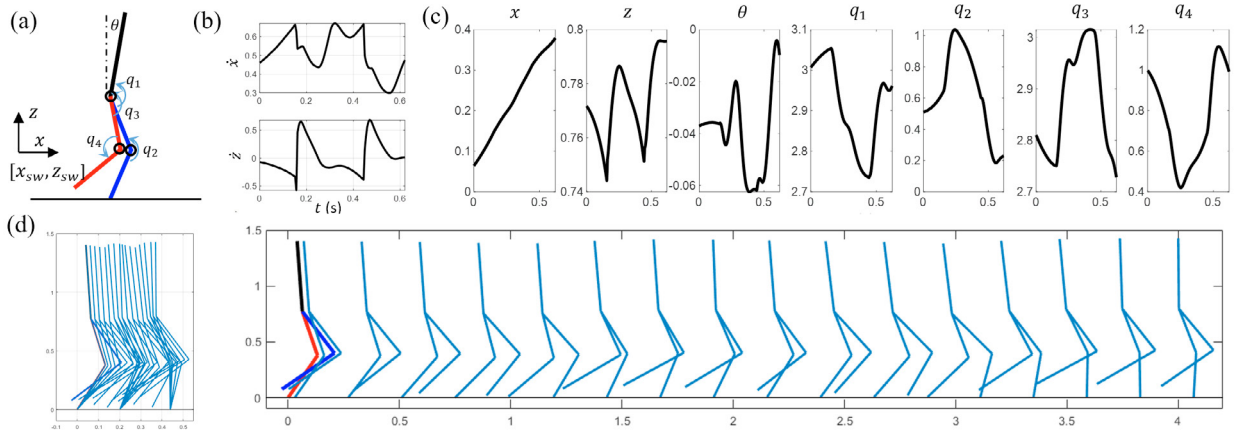

Fig. 4. (a) The five-linkage biped model. (b, c) Trajectories of the optimized walking. (d) Snapshots of the walking.

literature. Figure 4 displays the results of applying our method. The optimization is able to reach the periodic gate without needing to pre-specifiy the number of steps ahead of time. A description of the cost functions used to achieve this behavior can be found in Westenbroek et al. (2021). The simulation is conducted with a relaxation parameter $\varepsilon = 10^{-3}$, time step of $10^{-3}s$ and forward Euler integration.

5. DISCUSSION

There are several important avenues for future work. It will be important to understand how the state constraints on the original HOCP translate onto the hybrid manifold, as many HOCPs solved by practitioners involve such constraints. Future work will further characterize the numerics of the RHOCP, and seek to find numerical algorithms with can find approximate minimizers of the original HOCP. Finally, future research will examine directly formulating costs on the hybrid manifold.

REFERENCES

Ames, A.D. and Sastry, S. (2005). A homology theory for hybrid systems: Hybrid homology. In *International Workshop on Hybrid Systems: Computation and Control*, 86–102. Springer.

Branicky, M.S., Borkar, V.S., and Mitter, S.K. (1998). A unified framework for hybrid control: Model and optimal control theory. *IEEE trans. on autom. contl.*, 43(1), 31–45.

Burden, S.A., Gonzalez, H., Vasudevan, R., Bajcsy, R., and Sastry, S.S. (2015). Metrization and Simulation of Controlled Hybrid Systems. *IEEE Trans. Automat. Contr. (TAC)*, 60(9), 2307–2320.

Caines, P., Egerstedt, M., Malhame, R., and Schöllig, A. (2007). A hybrid bellman equation for bimodal systems. In *International Workshop on Hybrid Systems: Computation and Control*, 656–659. Springer.

Grizzle, J.W., Chevallereau, C., Sinnet, R.W., and Ames, A.D. (2014). Models, feedback control, and open problems of 3d bipedal robotic walking. *Automatica*, 50(8), 1955–1988.

Hereid, A., Hubicki, C.M., Cousineau, E.A., and Ames, A.D. (2018). Dynamic humanoid locomotion: A scalable formulation for hzd gait optimization. *IEEE Transactions on Robotics*, 34(2), 370–387.

Hiskens, I.A. and Pai, M. (2000). Trajectory sensitivity analysis of hybrid systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(2), 204–220.

Johansson, K.H., Egerstedt, M., Lygeros, J., and Sastry, S. (1999). On the regularization of zeno hybrid automata. *Systems & control letters*, 38(3), 141–150.

Kelley, J.L. (2017). *General topology*. Courier Dover Publications.

Lygeros, J., Johansson, K.H., Simic, S.N., Zhang, J., and Sastry, S.S. (2003). Dynamical properties of hybrid automata. *IEEE Trans. Automat. Contr.*, 48(1), 2–17.

Lygeros, J., Tomlin, C., and Sastry, S. (1999). Controllers for reachability specifications for hybrid systems. *Automatica*, 35(3), 349–370.

Or, Y. and Ames, A.D. (2010). Stability and completion of zeno equilibria in lagrangian hybrid systems. *IEEE Trans. on Autom. Control*, 56(6), 1322–1336.

Pakniyat, A. and Caines, P.E. (2017). On the relation between the minimum principle and dynamic programming for classical and hybrid control systems. *IEEE Trans. Automat. Contr.*, 62(9), 4347–4362.

Polak, E. (2012). *Optimization: algorithms and consistent approximations*, volume 124. Springer.

Shaikh, M.S. and Caines, P.E. (2007). On the hybrid optimal control problem: theory and algorithms. *IEEE Trans. Automat. Contr.*, 52(9), 1587–1603.

Simic, S.N., Johansson, K.H., Lygeros, J., and Sastry, S. (2005). Towards a geometric theory of hybrid systems. *Dynamics of Continuous, Discrete and Impulsive Systems Series B: Applications and Algorithms*, 12(5-6), 649–687.

Sotomayor, J. (1996). Regularization of discontinuous vector fields. World Scientific.

Sussmann, H.J. (1999). A maximum principle for hybrid optimal control problems. In *Proceedings of the 38th IEEE Conf Decis Control*, volume 1, 425–430.

Westenbroek, T., Gonzalez, H., and Sastry, S.S. (2018). A new solution concept and family of relaxations for hybrid dynamical systems. In *2018 IEEE Conf Decis Control (CDC)*, 743–750. IEEE.

Westenbroek, T., Xiong, X., Sastry, S.S., and Ames, A.D. (2021). Smooth approximations for hybrid optimal control problems with application to robotic walking. Arxiv Preprint.