

Computation of Regions of Attraction for Hybrid Limit Cycles Using Reachability: An Application to Walking Robots

Jason J. Choi, Ayush Agrawal, Koushil Sreenath, Claire J. Tomlin, and Somil Bansal

Abstract—Contact-rich robotic systems, such as legged robots and manipulators, are often represented as hybrid systems. However, the stability analysis and region-of-attraction computation for these systems are often challenging because of the discontinuous state changes upon contact (also referred to as *state resets*). In this work, we cast the computation of region-of-attraction as a Hamilton-Jacobi (HJ) reachability problem. This enables us to leverage HJ reachability tools that are compatible with general nonlinear system dynamics, and can formally deal with state and input constraints as well as bounded disturbances. Our main contribution is the generalization of HJ reachability framework to account for the discontinuous state changes originating from state resets, which has remained a challenge until now. We apply our approach for computing region-of-attractions for several underactuated walking robots and demonstrate that the proposed approach can (a) recover a bigger region-of-attraction than state-of-the-art approaches, (b) handle state resets, nonlinear dynamics, external disturbances, and input constraints, and (c) also provides a stabilizing controller for the system that can leverage the state resets for enhancing system stability.

Index Terms—Legged Robots, Optimization and Optimal Control, Hybrid Logical/Dynamical Planning and Verification.

I. INTRODUCTION

Practical legged locomotion is one of the fundamental problems in robotics. Legged robots can exhibit efficient, versatile, and highly dynamic gaits that allow them to operate in highly unstructured environments. Mathematically, legged robots are often represented as hybrid dynamical systems. Hybrid dynamical systems may exhibit both continuous and discrete dynamics, wherein swinging of a leg has continuous dynamics, whereas contacts with ground are well-modeled as discrete events [1]. Under hybrid models, walking gaits are represented as periodic *hybrid limit cycles*, i.e., a limit cycle

involving discontinuous jumps in the system state when the leg impacts the ground.

To obtain a stable walking behavior, it is often desirable to compute the *region-of-attraction (RoA)* for a walking gait – the robot configurations from which it can eventually converge to the corresponding hybrid limit cycle and follow the gait. Other than analyzing the stability of a gait, RoA estimation can also be useful for designing different candidate control laws for walking and switching between different gaits [2]. Despite significant progress over the last few years, stability analysis and RoA computation for legged robots is still quite challenging because the dynamics are highly nonlinear and intrinsically hybrid, requiring controlling through contact.

In this work, we cast the RoA computation as a Hamilton-Jacobi (HJ) reachability problem. In reachability analysis, one computes the *Backward Reachable Tube (BRT)* of a dynamical system – the set of states such that the trajectories that start from this set will eventually reach some given target set despite the worst case disturbance [3]. If we use the limit cycle as the target set, the corresponding BRT is the RoA of the limit cycle. The advantage of casting RoA computation as a reachability problem is the availability of numerical tools that can compute BRT for a variety of nonlinear systems in the presence of disturbances and control bounds. Even though computationally intensive compared to other existing methods that can solve for the underapproximations of RoAs, the reachability-based methods recover the maximal RoA without imposing any restriction on its shape.

However, one problem that still remains is that, so far, we lack a general framework for the reachability analysis of dynamical systems that involve discontinuous state changes (also called *state resets*). To overcome this challenge, we extend the reachability framework to dynamics with discontinuous state resets. Our key insight is that since HJ reachability analysis ultimately corresponds to solving a dynamic programming problem, we can reason about the effect of state resets on the BRT by applying the Bellman principle of optimality at the switching surface. We demonstrate that this corresponds to a value remapping between post and pre reset states during the BRT computation. Moreover, the existing numerical algorithms to compute BRTs can easily be extended to account for this remapping without incurring any additional computation cost. Thus, the obtained value function, the corresponding BRT, as well as the stabilizing

Manuscript received: Sep., 9, 2021; Revised Dec., 31, 2021; Accepted Jan., 24, 2022.

This paper was recommended for publication by Editor A. Kheddar upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by (organizations/grants which supported the work.)

This research is supported in part by NSF Grants CMMI-1931853, CMMI-1944722, an ONR Basic Research Challenge Program on Multibody Systems, NASA ULI program (through Stanford) on Safe Aviation Autonomy, and the DARPA Assured Autonomy Program. The work of Jason Choi received the support of a fellowship from Kwanjeong Educational Foundation. Jason Choi, Ayush Agrawal, Koushil Sreenath, and Claire Tomlin are with the University of California, Berkeley, and Somil Bansal is with the University of Southern California, Los Angeles. Contact info: {jason.choi, ayush.agrawal, koushils, tomlin}@berkeley.edu, somilban@usc.edu.

Project website: <https://hybridreachability.github.io/website-hybrid-roa/>
Digital Object Identifier (DOI): see top of this page.

controller implicitly reason about the effect of state resets on the stability of the system.

We apply the proposed approach on various “dynamic walking” systems and demonstrate that it can obtain a bigger RoA than most state-of-the-art approaches, while accounting for input constraints, state resets, and external disturbances at the same time. Finally, the proposed approach also provides a stabilizing controller for the system that can leverage state resets for a faster stabilization.

II. RELATED WORK

Poincaré Map-based methods. Stability analysis of limit cycles can be reduced to the problem of analyzing stability of the fixed point of the corresponding Poincaré map. In particular, if the fixed point of the Poincaré map is stable, then there exists a stable limit cycle for the original system that passes through this fixed point [4, 5]. However, finding an analytical expression for the Poincaré map is challenging; hence, a common approach is to compute a numerical approximation of its linearization [1, 6]. The local stability and RoA of the limit cycle are then computed using this discrete-time linear system [2]. This approach has also been used to analyze stability in the presence of external disturbances [7].

Lyapunov-based methods. One of the limitations of Poincaré analysis is that the verifiable RoA is restricted to the set of states on the chosen Poincaré section. The works in [8–10] use linearized transverse dynamics to estimate the RoA by constructing a (time-varying) Lyapunov function via Sum-of-Squares (SOS) programs. As noted in [8], a limitation of this approach is that the size of the computed RoA depends on the choice of the transverse coordinates. The work in [11] also proposes a Lyapunov based approach to compute RoA for nonlinear hybrid systems with state resets using the *complementarity* formulation of contact.

In several of the aforementioned methods, Lyapunov functions are computed using SOS programming. The key advantage of SOS-based methods is that they can leverage convex optimization to compute fast underapproximations of the RoA. However, as a tradeoff, the verified RoA can be a significant underapproximation of the true RoA. First, the SOS methods restrict the Lyapunov functions to the class of (piecewise) SOS polynomials. Second, although one can typically design a convex SOS program when the feedback controller is fixed, solving for the controller and the Lyapunov function together to maximize the verifiable RoA results in a nonconvex program which induces another source of underapproximation when we obtain a local optimum. These limitations are exacerbated further in the presence of state resets as the Lyapunov condition at the switching surface leads to a nonconvex program [8]. Ultimately, this results in recovering a smaller RoA. In contrast, the proposed approach does not impose any specific structure on the value function and hence can recover the maximal RoA (subjected to numerical errors), even in the presence of state resets.

Stabilizing controllers for hybrid systems with state resets. The *Hybrid Zero Dynamics* (HZD) framework [1, 12, 13] was one of the earliest works on creating stabilizable hybrid

limit cycles and stabilizing controllers for walking robots with discrete impacts. The method of Poincaré can also be incorporated into the HZD framework [14–16]. Building on the HZD framework, [17] proposes a *rapidly exponentially stabilizing* control Lyapunov function based quadratic program (CLF-QP) to exponentially stabilize hybrid limit cycles. There is also a rich literature on contact implicit trajectory optimization [18, 19] and model predictive control [20, 21] that can simultaneously handle multiple contact phases. Some of these methods have also leveraged value remapping to account for the state resets. In this work, we extend this principle to reachability analysis and focus on the RoA computation, as opposed to trajectory optimization.

III. PROBLEM SETUP

We consider the following class of hybrid systems:

$$\begin{aligned} \dot{x} &= f(x, u, d), & x \notin S \\ x^+ &= \Delta(x^-), & x^- \in S, \end{aligned} \quad (1)$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathcal{U}$ is the control input, and $d \in \mathcal{D}$ is the disturbance. S indicates the switching surface where the reset map Δ is applied; under the reset event at time t , the state $x^- := \lim_{\tau \nearrow t} x(\tau)$ instantaneously shifts to x^+ . Therefore, the resulting trajectory of this hybrid system is composed of a set of continuous trajectories along the vector flow f , and discontinuous jumps between them whenever the state hits the switching surface S .

Hybrid limit cycle: Let $x^*(\cdot)$ represents a non-trivial T -periodic solution of the system that undergoes N impacts at times $\{t_1, t_2, \dots, t_N\} + kT$ for non-negative integer k . We call $x^*(\cdot)$ a hybrid limit cycle of the system.

We assume that this limit cycle is given along with a baseline control law $\pi_0(\cdot) : \mathbb{R}^n \rightarrow \mathcal{U}$ that makes the limit cycle forward invariant and stable in the small neighborhood around it. There are many existing works that study approaches for designing such stable hybrid limit cycles [12, 13] as well as the locally-stabilizing control laws [1, 17].

The overall objective is twofold. First, we seek to compute a region of state space $\Omega(x^*) \subseteq \mathbb{R}^n$ around $x^*(\cdot)$, from which there exists an admissible control that can asymptotically stabilize the system to $x^*(\cdot)$. We call $\Omega(x^*)$ the (asymptotically) stabilizable region or RoA for $x^*(\cdot)$. Second, we also want to design a corresponding control law that stabilizes the system trajectory to $x^*(\cdot)$.

Remark 1: Note that depending on the baseline controller, the RoA under $\pi_0(\cdot)$ can be very small. Our goal in this work is to find a stabilizing control law that is beyond the capability of the baseline controller and provides us with the *maximal* RoA $\Omega(x^*)$.

Running example (Teleporting Dubins Car). We now introduce a simple hybrid system example that we will use throughout the paper to illustrate our method. We introduce state resets to a well-known Dubins car system which makes the car “teleport” (Fig. 1). It is chosen because the key principles of our method are well exposed in its design and results; more practical examples will be introduced in Sec VI. The dynamics are given by




Fig. 1: Teleporting Dubins car system. The car “teleports” to the other side of the x -axis whenever it hits the x -axis from a positive y -direction.

$$\dot{p}_x = v \cos \theta, \quad \dot{p}_y = v \sin \theta, \quad \dot{\theta} = u, \quad (2)$$

where $x = [p_x, p_y, \theta]$ is the state of the system consisting of car’s position (p_x, p_y) and its heading θ . v is the constant velocity of the car, and $u \in [-\bar{\omega}, \bar{\omega}]$ is the control input, which is the angular speed. Next, we introduce a switching surface and the reset map as the following:

$$S = \{x \mid p_y = 0, \sin \theta < 0\} \quad (3)$$

$$x^+ = \Delta(x^-) = [-R \cdot (|p_x|/R)^\alpha \cdot \text{sign}(p_x), p_y, \pi + \theta]$$

Intuitively, the car undergoes a discrete jump whenever it hits the x -axis from a positive y -direction. Upon the reset, the car’s x -position and heading change discontinuously. The parameter $\alpha > 0$ in the reset map dictates the nonlinearity of the reset map. For instance, $\alpha = 1$ makes $p_x^+ = -p_x$, i.e., the car position is mirrored around the y -axis. For $\alpha < 1$, the reset map has a contraction property—everytime the state resets, the value of $|p_x|$ gets closer to R . On the other hand, value of $\alpha > 1$ induces an expanding reset map.

The desired hybrid limit cycle is a semi-circular trajectory of radius R centered at the origin in the counterclockwise direction. Therefore, the orbit can be expressed as

$$x^*(\cdot) = \left\{ x \mid \|(p_x, p_y)\| = R, \theta = \frac{\pi}{2} + \tan^{-1}\left(\frac{p_y}{p_x}\right), \frac{\pi}{2} \leq \theta < \frac{3\pi}{2} \right\}.$$

We are interested in computing the RoA for $x^*(\cdot)$, i.e., the states from which the car can be stabilized to follow the semi-circular trajectory eventually, as well as the corresponding stabilizing controller. A baseline stabilizing controller $\pi_0(\cdot)$ is designed by feedback linearization, with the output defined as $y = \|(p_x, p_y)\| - R$.

IV. BACKGROUND

A. Hamilton-Jacobi Reachability

We present an overview of Hamilton-Jacobi (HJ) reachability analysis, first for continuous systems without the disturbance: the dynamics of the system is given by $\dot{x} = f(x, u)$. Let $\xi_{x,t}^u(\tau)$ denote the state achieved at time τ by starting at initial state x and initial time t , and applying input functions $u(\cdot)$ over $[t, \tau]$. Given a target set $\mathcal{L} \subset \mathbb{R}^n$, we are interested in computing the Backward Reachable Tube (BRT) of \mathcal{L} :

Backward Reachable Tube (BRT) is the set of initial states from which the agent acting optimally will eventually reach the target set \mathcal{L} within the time horizon $[t, T]$:

$$\mathcal{V}(t) = \{x \mid \exists u(\cdot) : [t, T] \rightarrow \mathcal{U}, \exists \tau \in [t, T], \xi_{x,t}^u(\tau) \in \mathcal{L}\}. \quad (4)$$

In HJ reachability, the computation of BRT is formulated as an optimal control problem, which can be solved by using the principle of dynamic programming.

First, a Lipschitz continuous target function $l(x)$ is defined whose zero sub-level set is the target set \mathcal{L} , i.e. $\mathcal{L} = \{x \mid l(x) \leq 0\}$. Typically, $l(x)$ is defined as a signed distance function to \mathcal{L} . The BRT seeks to find all states that could enter \mathcal{L} at any point in the time horizon. This is computed by finding the minimum distance to \mathcal{L} over time:

$$J(x, t, u(\cdot)) = \min_{\tau \in [t, T]} l(\xi_{x,t}^u(\tau)). \quad (5)$$

Our goal is to capture this minimum distance for *optimal trajectories* of the system. Thus, we compute the optimal control that minimizes this distance (and eventually drives the system into the target set when $l(\cdot) \leq 0$). The value function corresponding to this optimal control problem is:

$$V(x, t) = \inf_{u(\cdot)} \left\{ J(x, t, u(\cdot)) \right\}. \quad (6)$$

The value function in (6) can be computed using dynamic programming, which results in the following final value Hamilton-Jacobi Isaacs Variational Inequality (HJI-VI):

$$\min \left\{ D_t V(x, t) + H(x, t), l(x) - V(x, t) \right\} = 0, \quad (7)$$

with the terminal value function $V(x, T) = l(x)$. D_t and ∇ represent the time and spatial gradients of the value function. H is the Hamiltonian that encodes the role of dynamics and the optimal control inputs:

$$H(x, t) = \min_u \langle \nabla V(x, t), f(x, u) \rangle. \quad (8)$$

Intuitively, the term $l(x) - V(x, t)$ in (7) enforces the value function to “memorize” the best record (closest instance to the target set) of the optimal trajectories. Once $V(x, t)$ is obtained, the BRT is given as the zero sub-level set of the value function $\mathcal{V}(t) = \{x \mid V(x, t) \leq 0\}$. The corresponding optimal control for reaching the target set \mathcal{L} is derived as

$$u^*(x, t) = \arg \min_u \langle \nabla V(x, t), f(x, u) \rangle. \quad (9)$$

Finally, note that every state that is contained in a *finite-time* BRT of \mathcal{L} , represents an initial condition from which the system can be controlled to \mathcal{L} within finite time.

B. Casting RoA computation as a HJ reachability problem

We now explain how HJ reachability can be used to estimate the RoA for a hybrid limit cycle, $\Omega(x^*)$. Let $h(x)$ denote distance to the limit cycle $x^*(\cdot)$, i.e., $h(x) := \inf_{y \in x^*(\cdot)} \|x - y\|$. We define the target function as $l(x) := h(x) - \epsilon$ for some $\epsilon > 0$. Thus, the target set \mathcal{L} is an ϵ -tube around the orbit $x^*(\cdot)$. Note that for a stable hybrid limit cycle, we can pick ϵ such the baseline controller π_0 can stabilize to the limit cycle from anywhere inside \mathcal{L} .

Given the value function, $V(x, t)$ and the BRT, $\mathcal{V}(t)$, corresponding to this target set, we can stabilize to the limit cycle from any state within the BRT using the following control law:

$$u(x, t) = \begin{cases} \arg \min_u \langle \nabla V(x, t), f(x, u) \rangle, & \text{if } x \notin \mathcal{L} \\ \pi_0(x), & \text{otherwise} \end{cases} \quad (10)$$

Intuitively, we can use the controller provided by HJ reachability to steer the system to \mathcal{L} from any state within the BRT; upon reaching \mathcal{L} , we can switch to the baseline controller. Therefore, for every state that can be verified as an element

of a finite-time BRT of \mathcal{L} , we can conclude that it is an element of $\Omega(x^*)$. In other words, $\mathcal{V}(t)$ can be used as an estimation of $\Omega(x^*)$. The corresponding stabilizing controller is given by (10).

Remark 2: Note that if we pick $\epsilon = 0$, only the states that can achieve finite-time convergence to $x^*(\cdot)$ can be verified from HJ reachability analysis. However, this would exclude many states that are not finite-time stabilizable but still asymptotically stabilizable to $x^*(\cdot)$. In fact, even for a state where stronger notion like exponential stabilizability [17] is achievable, its time-to-reach to the orbit is not necessarily a finite value. Therefore, by using $\epsilon > 0$ together with the baseline policy π_0 which makes sure that every state in \mathcal{L} is stabilizable to the actual orbit, we can capture the states that are not necessarily finite-time but are still asymptotically stabilizable to $x^*(\cdot)$ and are inside the finite-time BRT of \mathcal{L} . For more detailed analysis and a constructive way to choose ϵ , please refer to [22]. Finally, defining \mathcal{L} as the ϵ -neighborhood of $x^*(\cdot)$ is also necessary for numerical methods to compute $V(x, t)$ and to tackle the presence of disturbance, which will be explained in the next subsection.

Running example (Teleporting Dubins Car). The target function $l(x)$ for the running example is $l(x) := \|(p_x, p_y)\| - R - \epsilon$. The Hamiltonian can be computed analytically in this case (see [23]):

$$H(x, t) = \frac{\partial V}{\partial p_x} v \cos \theta + \frac{\partial V}{\partial p_y} v \sin \theta - \bar{\omega} \left| \frac{\partial V}{\partial \theta} \right|. \quad (11)$$

C. Robustness to bounded disturbance

For systems with bounded disturbance $d \in \mathcal{D}$, HJ reachability can again be used to synthesize optimal controllers that are robust to the disturbance. The definition of BRT is now the set of initial states for which, *under worst-case disturbances*, the agent acting optimally will eventually reach the target set \mathcal{L} within time $[t, T]$:

$$\mathcal{V}(t) = \{x : \exists u(\cdot), \forall d(\cdot) : [t, T] \rightarrow \mathcal{D}, \exists \tau \in [t, T], \xi_{x,t}^{u,d}(\tau) \in \mathcal{L}\}. \quad (12)$$

Computation of the BRT can be formulated as a zero-sum game between the control and the disturbance, and can be solved in a similar way by applying the dynamic programming [23]. $V(x, t)$ is governed by the same variational inequality as (7), where the Hamiltonian is as follows:

$$H(x, t) = \max_u \min_d \langle \nabla V(x, t), f(x, u, d) \rangle. \quad (13)$$

Finally, the optimal control signal that is robust to the disturbance can be synthesized by the following rule:

$$u^*(x, t) = \arg \max_u \min_d \langle \nabla V(x, t), f(x, u, d) \rangle. \quad (14)$$

For the specific problem of stabilization to the hybrid limit cycle, it must be noted that any non-zero disturbance signal can break the forward-invariance of the baseline controller $\pi_0(\cdot)$ on the limit cycle $x^*(\cdot)$. Therefore, the notion of $x^*(\cdot)$ being *robustly asymptotically stable* is defined as an existence of an ϵ -neighborhood of the limit cycle to which the system can be asymptotically stabilized under the worst-case disturbance [24]. Under a mild assumption that the baseline controller $\pi_0(\cdot)$ itself renders the set \mathcal{L} locally robustly stable,

any elements of the finite-time BRT for \mathcal{L} can be verified as robustly asymptotically stabilizable to $x^*(\cdot)$.

V. HAMILTON-JACOBI REACHABILITY FOR DYNAMICS WITH STATE RESETS

We now present the main contribution of this paper, which is an extension of the reachability framework above (for continuous dynamics) to account for discontinuous state resets. First, the following theorem proves that the value function in the presence of state resets can be obtained by solving a constrained version of HJI-VI in (7).

Theorem 1: $V(x, t)$ in the presence of state resets is given by the solution to the following constrained VI:

$$\begin{aligned} \min \left\{ D_t V(x, t) + H(x, t), l(x) - V(x, t) \right\} &= 0, \quad x \notin S, \\ V(x, t) &= V(\Delta(x), t), \quad x \in S, \end{aligned} \quad (15)$$

with the terminal value function

$$\begin{aligned} V(x, T) &= l(x), \quad x \notin S, \\ V(x, T) &= l(\Delta(x)), \quad x \in S. \end{aligned} \quad (16)$$

Intuitively, Theorem 1 states that the value function can be obtained by solving the usual HJI-VI for the states that are not on the switching surface. For the states that are on the switching surface, the value is given by that of the corresponding post reset state. This is not surprising because if the state is on the switching surface, it will instantaneously change to the post reset state. Consequently, how “close” the state is to the target set is completely characterized by how close the post reset state is to the target set. The proof of Theorem 1 formalizes this intuition using the Bellman principle of optimality and is presented in Appendix A in the supplementary material.

The interesting aspect about the HJI-VI in (15) is that since it reasons about the state resets, the obtained value function and the stabilizing controller implicitly exploit the reset map to reach the target set as quickly as possible. We will demonstrate this aspect further in Sec. VI-A.

Numerical Implementation. We now present a numerical method to compute the value function based on the result in Theorem 1. The baseline of the numerical algorithm is a standard level set-based method that computes the value function over a discretized computational domain of a compact subset $\mathcal{X} \subset \mathbb{R}^n$ [25]. Let $\{p(x)\}_G$ denote the set of values of a function $p(x)$ evaluated over a state space grid G . Since our computation will proceed backward in time, we will numerically solve the variational inequality in (15) for each discrete time step in a finite interval $[0, T]$ using the three-step update rule described in Algorithm 1.

The numerical algorithm is implemented by means of the computational methods provided in [25]. It is important to stress the remarkable computational similarity of this new algorithm to its non-reset counterpart in (7). Indeed, the only computational overhead is introduced by the step in (18). As a result, as will be demonstrated in the following section, our method can compute the BRT at essentially no additional cost compared to the case where there are no state resets.

Algorithm 1: Value Function Computation for Dynamics with State Resets.

Input : Target function $\{l(x)\}_G$, Time horizon T
Output: Value function $\{V(x, t)\}_G$ for $t = 0, \delta, 2\delta, \dots, T$

 Initialization: $\{V(x, T)\}_{G \setminus S} = \{l(x)\}_{G \setminus S}$,
 $\{V(x, T)\}_S = l(x^+)$,
 $t = T$;

while $t > 0$ **do**

 For all $x \in G \setminus S$,

$$V(x, t - \delta) = V(x, t) + \min_{u \in \mathcal{U}} \max_{d \in \mathcal{D}} D_x V(x, t) \cdot f(x, u) \delta; \quad (17)$$

$$V(x, t - \delta) = \min\{V(x, t - \delta), l(x)\};$$

 For all $x \in S$,

$$V(x, t - \delta) = V(\Delta(x), t - \delta); \quad (18)$$

 $t = t - \delta$;

end

Remark 3: Note that Theorem 1 and the numerical method in Algorithm 1 are not specific to the RoA computation problem, which is the main focus of this paper. In general, these can be applied to any kind of target set \mathcal{L} that can be expressed as a zero-sublevel set of a Lipschitz continuous function $l(x)$ and has a non-empty volume.

VI. CASE STUDIES

In this section, we will demonstrate how our extended HJ reachability framework can be used to compute RoA for various hybrid systems, including a simple biped walker.

A. Running example: Teleporting Dubins car

We first apply our method to the running example. We use the following parameters for our simulations: $v = 1 \text{ m s}^{-1}$, $R = 2 \text{ m}$, $\bar{\omega} = 1 \text{ rad s}^{-1}$, $\alpha = 0.5$, which induces a contracting reset map, and $\epsilon = 0.2$.

To confirm that our new method is able to capture states that exploit discrete jumps to reach the limit cycle, we compare our method to a simple baseline – we compute the set of states that are stabilizable to the limit cycle only by following the continuous part of the dynamics. Thus, for this baseline, we “freeze” the dynamics on the switching surface, i.e. set $f(x, u) \equiv 0$ whenever $x \in S$.

Fig. 2 visualizes the BRTs computed for $T = 6.3 \text{ s}$ from our method (blue), and under the frozen dynamics (brown). It shows the 2D slices of the BRTs at a fixed heading angle $\theta = -3\pi/4 \text{ rad}$. Note that the BRT computed by freezing the dynamics is strictly contained in the BRT computed by our method. The gap between the two sets uncovers the states that can reach the target set within 6.3 s only via leveraging discrete jumps. Fig. 3 displays the resulting optimal trajectory (blue) from a particular initial state that




Fig. 2: p_x - p_y slices of the Backward Reachable Tubes for the teleporting Dubins car at $\theta = -3\pi/4$, computed by value remapping (blue), and by freezing the dynamics (brown) on the switching surface ($T = 6.3 \text{ s}$, $\alpha = 0.5$). The gap between the two sets captures the states that can reach the target set within 6.3 s only by using the “teleports” (reset map). We evaluate an initial state included in the BRT (cross) in Fig.3.




Fig. 3: The optimal trajectory simulated for 6.3 s which reaches the target set after 4.3 s by exploiting two discrete transitions. For comparison, the feedback-linearization-based baseline controller is simulated for 14 s , which reaches the target set after 11.9 s .

is contained in this gap (marked with the red cross). It exploits two state jumps to quickly reach the target set after 4.3 s ; at every discrete jump, the distance to the orbit reduces significantly. In contrast, the feedback linearization-based baseline-controller (grey), although doing its best to get closer to the orbit by heading towards the radial direction, is able to reach the target set only after 11.9 s . Furthermore, the discrete jump near the limit cycle is not “intended” by the baseline controller, since the control law is not aware of the reset map. In contrast, the proposed method is *explicitly* able to reason that since the reset map is a contraction in this case, it can be exploited to reduce the distance to the orbit and stabilize to the limit cycle faster.

In Fig. 2, we can observe that many states adjacent to the switching surface are captured as the stabilizable region to the limit cycle. This is because $\alpha = 0.5$ induces the




Fig. 4: p_x - p_y slices of the Backward Reachable Tubes at $\theta = -3\pi/4$ (top) and $\theta = \pi/4$ (bottom), with various values of the reset map nonlinearity parameter α under the value remapping, and under dynamics freezing, applied to the switching surface ($T = 6.3 \text{ s}$). Note that the reset map is contracting when $\alpha < 1$, and expanding when $\alpha > 1$.




Fig. 5: (Left) Rimless wheel system. (Right) Compass-gait walker.

contracting property of the reset map, which is beneficial for the stabilization. Fig. 4 visualizes BRTs computed for the same T under three different values of α . Note that $\alpha = 3.0$ is unfavorable for the stabilization as it is a divergent reset map; as a result, most of the region adjacent to the switching surface is not captured in the BRT. In fact, with even larger α , the BRT eventually shrinks to the BRT computed by freezing the dynamics on the switching surface, i.e., the reset map is so divergent that it is no longer possible to stabilize the system after hitting the switching surface.

Computation details: The grid used for the computation is constructed in the polar coordinate system of the dynamics ($r = \|(p_x, p_y)\|$, $\alpha = \tan^{-1}(p_y/p_x)$, θ); it is a $81 \times 41 \times 81$ grid over the state space of $[0, 16] \times [0, \pi] \times [-\pi, \pi]$. It takes under an hour to compute the BRT of $T = 6.3$ s on a computer with a 2.6GHz 6-core 9th-gen. Intel Core i7 processor.

B. Rimless wheel system

We next apply the proposed approach to the rimless wheel system. The rimless wheel is a popular model of a passive-dynamic walker. One of the reasons for its popularity is that RoA can be computed analytically for the rimless wheel [26], and hence it can be used to analyze the efficacy of a RoA computation algorithm.

The system consists of a central mass with several spokes extending radially outward (see Fig. 5); it has no control input. At any given time, one of the spokes is pinned at the ground, and the system follows the dynamics of a simple pendulum, $f(\theta, \dot{\theta}) = [\dot{\theta}, \sin(\theta)]^T$, where $x := (\theta, \dot{\theta})$ is the state of the system. When an unpinned spoke contacts the ground, an inelastic collision takes place and there is an instantaneous change in θ . Upon collision, the pinned spoke leaves the ground and the new spoke becomes the pinned one. The switching surface and the reset map are given as:

$$S = \{x \mid \theta = \alpha + \gamma\}, \quad x^+ = \Delta(x^-) = \begin{bmatrix} 2\gamma - \theta \\ \cos(2\alpha)\dot{\theta} \end{bmatrix}.$$

For a sufficiently inclined slope, the rimless wheel exhibits several stable hybrid limit cycles that are given as

$$x^*(\cdot) = \{x \mid \cos(\theta) + \frac{1}{2}\dot{\theta}^2 = E\},$$

where E is the total initial energy in the system. One such limit cycle is shown in black in Fig. 6 (left). The figure is reproduced from [8]. The right boundary of the plot is the switching surface—whenever the system state reaches the right boundary, it resets back to the left boundary as per the reset map. The light pink region is the exact RoA of the




Fig. 6: Regions of Attraction (RoA) for the rimless wheel limit cycle (black). (Left) The true RoA is shown in light pink. RoA obtained using SOS programming is shown in purple. The figure is reproduced from [8]. (Right) RoA computed using the proposed approach. The proposed approach is able to recover the entire RoA.

system [26]. The purple region is the RoA obtained using a time-varying Lyapunov function, which is synthesized using SOS programming [8]. As evident from the figure, using the SOS programming, we are able to recover only a subset of the RoA. In Fig. 6 (right), we demonstrate the RoA obtained using our method (dark pink). The proposed approach is able to recover the entire RoA of the system.

Computation details: We use a grid of size 201×201 over the state space $[-0.2, 0.6] \times [-0.6, 1.3]$. It takes within a few minutes to compute the BRT on a standard laptop.

C. Compass-gait walker

We next consider a compass-gait walker, which consists of two links with an actuated joint between them. Similar to the rimless wheel example, we consider a *pinned* model of the robot, with the configuration variable $q := [q_1, q_2]^T$ as illustrated in Fig. 5. We consider a hybrid model of walking, with alternating phases of a continuous-time single-support phase followed by an instantaneous inelastic impact of the swing leg with the ground. The dynamics in single-support can be represented by the standard manipulator equation, and the reset map Δ is derived from an inelastic impact model [1]. By defining the state as $x := [q, \dot{q}]^T$, the continuous dynamics can be expressed as control-affine: $\dot{x} = f(x) + g(x)u$. The switching surface is defined as the set of states where the swing foot strikes the ground with a negative velocity and the stance leg angle crosses a predefined threshold \bar{q}_1 ,

$$S := \{x \mid q_1 \leq \bar{q}_1, 2q_1 + q_2 = 0, 2\dot{q}_1 + \dot{q}_2 < 0\}. \quad (19)$$

The target hybrid limit cycle we wish to stabilize to is designed as a reference walking gait [12, 13, 27], which is tracked by a baseline controller designed based on feedback linearization and CLF-QP [17]. The configuration of the swing leg on the gait, $q_{2,r}$, can be parametrized as a polynomial function of the stance leg angle q_1 , and we define the output as $y := q_2 - q_{2,r}$ for feedback linearization. Each walking step takes 0.73 s, where q_1 starts with $q_{1,d} = 0.13$ and ends at $-q_{1,d}$. The phase plot of the reference gait in q_1 - q_2 space is visualized in Fig. 7(a) black dashed line.

In Fig. 7, the BRTs computed from our method for $T = 1.5$ s is visualized as the projection of the 4D set to the q_1 - q_2 space (colored region). Note that the chosen time




Fig. 7: Backward Reachable Tubes (BRTs) for the compass-gait walker visualized by their projection on the q_1 - q_2 space. ($T = 1.5$ s for all three cases) (a) BRT for the closed-loop dynamics under the CLF-QP. (b) BRT with the control bound $u \in [-4, 4]$, without the disturbance. (c) BRT with the control bound $u \in [-4, 4]$, and the disturbance bound $d \in [-0.3, 0.75]$. The color indicates the thickness of the BRT in the 4D space; it gets thicker as yellow \rightarrow blue.




Fig. 8: q_2 - \dot{q}_2 slices of the BRT corresponding to Figure 7(a), along the hybrid periodic orbit (reference gait). (a) Initial phase: $[q_1, \dot{q}_1] = [0.13, -0.67]$. (b) Intermediate phase: $[q_1, \dot{q}_1] = [0, -0.36]$. (c) Terminal phase (imminent of reset): $[q_1, \dot{q}_1] = [-0.13, -0.91]$. The black dots indicate the corresponding $[q_2, \dot{q}_2]$ values on the orbit, which are always contained in the BRT. The color indicates the value of $V(x, T)$; it is yellow when $V(x, T) = 0$ and gets blue as the value get smaller. The diagrams on the top right corners visualize the swing leg (blue) configuration range that is contained in the BRT, together with the stance leg (magenta) for each phase.

horizon encompasses more than two nominal walking steps. First, we apply our method to the closed loop dynamics of the baseline controller (Fig. 7(a)). Although it is designed to be locally stabilizing, the verified RoA comprises only a small neighborhood of the reference gait. By contrast, Fig. 7(b) shows the verified stabilizable region for the gait by allowing full control authority defined by the control bound $u \in [-4, 4]N$. The bound is selected such that on the limit cycle, the CLF-QP barely saturates this bound, so that we do not give additional advantage to the optimal controller. However, by exploiting full control capacity and the reset map, we are able to verify that the actual stabilizable region is much bigger than what can be achieved by the CLF-QP controller. Fig. 8 shows the q_2 - \dot{q}_2 slices of the BRT along the reference gait. Note that the gait itself is actually contained in the BRT, and that the BRT gets smaller as it gets closer to the end of the step (Fig. 8(c)), because of the imminent reset.

In Fig. 9, we show some trajectories that start at initial states which are contained in the verified BRT. These initial states have a large deviation from the reference gait, corresponding to a large perturbation. We start by applying the reachability-based optimal control (blue), and then switch to the CLF-QP controller (orange) when we reach the target




Fig. 9: Phase plots of the reachability-based optimal trajectories (blue \rightarrow orange) and the CLF-QP trajectories (grey) in q_1 - q_2 space for various initial states from which the CLF-QP fails to stabilize to the gait. The trajectories are evaluated for 4 walking steps. Each start of the step is indicated by the index of the step in italic type. The evolution of time is indicated by the color fading out. The optimal trajectory switches from HJ optimal control to CLF-QP once it reaches the target set (indicated by the transition from blue to orange curve).

set. In Fig. 9(b,c), the optimal controller is exploiting the reset map to converge to the limit cycle faster. Physically, this means that the robot steps on the ground first, and use this ground force to converge to the nominal gait at the next walking step. For comparison, the trajectories from the same initial states where the CLF-QP is used from the very beginning are visualized together (grey), which all fail to recover from the ground impact and cannot stabilize to the limit cycle. In particular, in Fig. 9(c), the CLF-QP fails to overcome a large momentum, resulting in the robot to fall backwards.

Finally, we introduce disturbance to the robot's dynamics: an unmodeled repulsive or stiction torque applied to the joint between the leg. This can be expressed as $\dot{x} = f(x) + g(x)u + \dot{q}_2 g(x)d$. There could be various source of uncertainty for this disturbance, for instance, unmodeled motor dynamics or joint friction. The bound of the disturbance is $d \in [-0.75, 0.3]$; under this bound, on the limit cycle, the CLF-QP is still able to robustly stabilize the gait. The BRT computed under the robust optimal control setting (as explained in Sec. IV-C) is visualized in Fig. 7(c). Since now we only verify states that are *robustly* stabilizable to the gait, the resulting BRT is smaller than the non-robust BRT in Fig. 7(b). However, it still shows that by applying the robust optimal controller, it is able to robustly stabilize a larger set of states than the CLF-QP. For instance, the CLF-QP fails to stabilize a state slightly perturbed from the gait under the presence of disturbance $d = -0.75$ (Fig. 10(b)), because the CLF-QP cannot reason about the additional torque required to compensate the effect of the disturbance. In contrast, the robust stabilizing controller obtained using the proposed approach is simultaneously able to counter the disturbances and stabilize to the reference gait (Fig. 10(c)).

Computation details: We use a grid of size $41 \times 81 \times 81 \times 81$ over the state space of $[-0.52, 0.52] \times [-1.04, 1.04] \times [-4, 4] \times [-8, 8]$, taking roughly 36 hours to compute the BRT of 1.5 s for this demonstration using a 3.3GHz 12-core Intel Core i7-5820K CPU.




Fig. 10: Phase plots of the trajectories in q_1 - q_2 space evaluated for 4 walking steps for a same initial state: (a) CLF-QP under no model-plant mismatch. (b) CLF-QP under model-plant mismatch, which fails to stabilize to the orbit. (c) Optimal control, which is able to reach the orbit under the same model-plant mismatch.

VII. CONCLUSION AND FUTURE WORK

We present a reachability-based approach to compute regions-of-attraction for dynamical systems with state resets. We apply our approach to several underactuated walking robots and demonstrate that the proposed approach can synthesize a bigger RoA compared to the state-of-the-art approaches and also provides a stabilizing controller. However, the proposed approach has several limitations that we will like to address in future. First, the reachability computation scales exponentially with the number of states, limiting its direct use to relatively low-dimensional systems. Recently, there have been some promising results that use Neural PDE solvers for computing reachable set for high-dimensional systems [28]. We will explore using these methods to extend the proposed approach to more complex walking robots. For providing stability guarantees, it would be interesting to leverage recent sampling-based methods [29] that provide probabilistic guarantees on the learned value functions. Second, it will be interesting to extend the proposed approach to other types of reachability problems, such as reaching a target set while avoiding some undesirable states, e.g., obstacles. Applying the proposed approach to real-world robotic systems will be another exciting future direction.

REFERENCES

- [1] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, et al. *Feedback control of dynamic bipedal robot locomotion*. CRC press, 2018.
- [2] M. S. Motahar, S. Veer, and I. Poulakakis. “Composing limit cycles for motion planning of 3D bipedal walkers”. *CDC*. 2016.
- [3] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-Jacobi Reachability: A Brief Overview and Recent Advances”. *CDC*. 2017.
- [4] J. Guckenheimer and P. Holmes. “Structurally stable heteroclinic cycles”. *Mathematical Proceedings of the Cambridge Philosophical Society*. Cambridge University Press. 1988.
- [5] J. Grizzle. “Remarks on event-based stabilization of periodic orbits in systems with impulse effects”. *Int. Symposium on Communications, Control and Signal Processing*. 2006.
- [6] B. Morris and J. W. Grizzle. “A restricted Poincaré map for determining exponentially stable periodic orbits in systems with impulse effects: Application to bipedal robots”. *CDC*. 2005.
- [7] S. Veer, I. Poulakakis, et al. “Input-to-state stability of periodic orbits of systems with impulse effects via Poincaré analysis”. *TAC* (2019).
- [8] I. R. Manchester, M. M. Tobenkin, M. Levashov, and R. Tedrake. “Regions of attraction for hybrid limit cycles of walking robots”. *IFAC Proceedings Volumes 44.1* (2011).
- [9] I. R. Manchester. “Transverse dynamics and regions of stability for nonlinear hybrid limit cycles”. *IFAC Proceedings Volumes 44.1* (2011).
- [10] L. B. Freidovich, A. S. Shiriaev, and I. R. Manchester. “Stability analysis and control design for an underactuated walking robot via computation of a transverse linearization”. *IFAC Proceedings Volumes 41.2* (2008).
- [11] M. Posa, M. Tobenkin, and R. Tedrake. “Lyapunov analysis of rigid body systems with impacts and friction via sums-of-squares”. *HSCC*. 2013.
- [12] J. W. Grizzle, G. Abba, and F. Plestan. “Asymptotically stable walking for biped robots: Analysis via systems with impulse effects”. *TAC* (2001).
- [13] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek. “Hybrid zero dynamics of planar biped walkers”. *TAC* (2003).
- [14] K. A. Hamed, B. G. Buss, and J. W. Grizzle. “Exponentially stabilizing continuous-time controllers for periodic orbits of hybrid systems: Application to bipedal locomotion with ground height variations”. *IJRR* (2016).
- [15] C.-L. Shih, J. Grizzle, and C. Chevallereau. “From stable walking to steering of a 3D bipedal robot with passive point feet”. *Robotica* 30.7 (2012).
- [16] A. Agrawal and K. Sreenath. “Discrete Control Barrier Functions for Safety-Critical Control of Discrete Systems with Application to Bipedal Robot Navigation.” *RSS*. 2017.
- [17] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle. “Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics”. *TAC* (2014).
- [18] M. Posa, C. Cantu, and R. Tedrake. “A direct method for trajectory optimization of rigid bodies through contact”. *IJRR* (2014).
- [19] H. Li and P. M. Wensing. “Hybrid systems differential dynamic programming for whole-body motion planning of legged robots”. *RA-L* (2020).
- [20] S. Cleac’h, T. Howell, M. Schwager, and Z. Manchester. “Linear Contact-Implicit Model-Predictive Control”. *arXiv* (2021).
- [21] T. Marcucci, R. Deits, M. Gabiccini, et al. “Approximate hybrid model predictive control for multi-contact push recovery in complex environments”. *Humanoids*. 2017.
- [22] F. Camilli, L. Grüne, and F. Wirth. “Control Lyapunov Functions and Zubov’s Method”. *SIAM Journal on Control and Optimization* (2008).
- [23] I. Mitchell, A. Bayen, and C. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. *TAC* (2005).
- [24] R. A. Freeman and P. V. Kokotovic. “Inverse optimality in robust stabilization”. *SIAM journal on control and optimization* (1996).
- [25] I. Mitchell. “A toolbox of level set methods”. <http://www.cs.ubc.ca/mitchell/ToolboxLS/ToolboxLS.pdf> (2004).
- [26] M. J. Coleman. *A stability study of a three-dimensional passive-dynamic model of human gait*. Cornell University, 1998.
- [27] A. Hereid and A. D. Ames. “Frost: Fast robot optimization and simulation toolkit”. *IROS*. 2017.
- [28] S. Bansal and C. Tomlin. “DeepReach: A deep learning approach to high-dimensional reachability”. *ICRA*. 2021.
- [29] N. Boffi, S. Tu, N. Matni, et al. “Learning stability certificates from data”. *CoRL* (2020).

APPENDIX A

A. Proof of Theorem 1

We will first prove the result for a state that is not on the switching surface. Let $x(t) \notin S$. The result for this case follows from the proof of standard HJI-VI. Here we present an informal proof based on the Taylor series expansion of the value function; a more rigorous proof of HJI-VI can be found in [23].

Consider a small $\delta > 0$ such that there is no state reset in the time interval $[t, t + \delta]$. The dynamic programming principle for the cost function in (5) implies that

$$\begin{aligned} V(x(t), t) &= \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \min\{l(x(t)), V(x(t + \delta), t + \delta)\}, \\ &= \min\{l(x(t)), \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} V(x(t + \delta), t + \delta)\}, \quad (20) \end{aligned}$$

On the other hand, the Taylor expansion of $V(x(t + \delta), t + \delta)$ can be written as:

$$V(x(t + \delta), t + \delta) = V(x(t), t) + D_t V(x(t), t) \delta + \nabla V(x(t), t) \delta x + \text{h.o.t.},$$

where δx is change in the state and can be approximated as $f(x, u, d) \delta$. Ignoring the higher order terms and plugging the Taylor approximation in (20), we have

$$\begin{aligned} V(x(t), t) &\approx \min\{l(x(t)), V(x(t), t) + D_t V(x(t), t) \delta \\ &\quad + \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \nabla V(x(t), t) \cdot f(x, u, d) \delta\}, \end{aligned}$$

where we have separated terms that do not depend on u and d . Subtracting $V(x(t), t)$ on both sides, we get

$$\begin{aligned} \min\{l(x(t)) - V(x(t), t), \\ \delta \left[D_t V(x(t), t) + \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \nabla V(x(t), t) \cdot f(x, u, d) \right]\} \approx 0, \end{aligned}$$

Since the above equation holds for all $\delta > 0$, we must have

$$\begin{aligned} \min\{l(x(t)) - V(x(t), t), \\ \left[D_t V(x(t), t) + \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \nabla V(x(t), t) \cdot f(x, u, d) \right]\} \approx 0, \end{aligned}$$

which results into the desired HJI VI.

We now turn our attention to a state $x(t) \in S$. Once again, the dynamic programming principle for the cost function in (5) implies that

$$V(x(t), t) = \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \min\{l(x(t)), V(x(t + \delta), t + \delta)\},$$

Since $x(t) \in S$, the system state will instantaneously change to $x^+(t)$ under any control and disturbance action. Thus, we have

$$\begin{aligned} V(x(t), t) &= \sup_{u \in \mathcal{U}} \inf_{d \in \mathcal{D}} \min\{l(x^+(t)), V(x^+(t + \delta), t + \delta)\}, \\ &= V(x^+(t), t), \end{aligned}$$

where the second equality follows from the definition of value function for $x^+(t)$.