ZonoReach: A Reachability-Guided Controller Using Zonotopes and Local Hamilton–Jacobi Analysis

El-Hajj, Isabelle; Beers, Jasper van; Solanki, Prashant

DOI
10.1007/s12555-025-0554-z

Publication date
2025

Document Version
Final published version

Published in
International Journal of Control, Automation and Systems

Citation (APA)
El-Hajj, I., Beers, J. V., & Solanki, P. (2025). ZonoReach: A Reachability-Guided Controller Using Zonotopes and Local Hamilton–Jacobi Analysis. *International Journal of Control, Automation and Systems*, *23*(11), 3199-3208. https://doi.org/10.1007/s12555-025-0554-z

International Journal of Control, Automation, and Systems 23(11) (2025) 3199-3208
http://dx.doi.org/10.1007/s12555-025-0554-z
ISSN:1598-6446 eISSN:2005-4092
http://www.springer.com/12555

ZonoReach: A Reachability-Guided Controller Using Zonotopes and Local Hamilton–Jacobi Analysis

Isabelle El-Hajj*, Jasper van Beers, and Prashant Solanki

Abstract: This paper presents a reachability-guided controller for nonlinear systems that synthesizes pseudo-optimal control using only local linear models. At each step, a forward reachable tube (FRT) is computed via zonotope-based set propagation; the closest point in the FRT to the target is chosen as an intermediate waypoint, around which a backward reachable tube (BRT) is solved using Hamilton–Jacobi (HJ) reachability. The resulting value function yields a locally optimal control action. This process is repeated iteratively to steer the system toward the target without requiring global nonlinear dynamics. We evaluate the method on the double integrator, inverted pendulum, and Dubins car, benchmarking against model predictive control baselines. For the double integrator, we additionally benchmark against its ground-truth time-optimal bang-bang solution. Our proposed ZonoReach controller achieves successful setpoint tracking and near time-optimal performance. Results highlight the influence of planning and control horizons, while limitations include reliance on local linear approximations and grid-based solvers for BRT computation. We conclude with directions for improving scalability toward real-world systems.

Keywords: Backward and forward reachable tubes, Hamilton–Jacobi reachability, reachability-based control, zonotopes.

1. INTRODUCTION

The problem of synthesizing goal-reaching controllers for nonlinear systems under input constraints is central to robotics [1], autonomous systems [2], and safety-critical control [3]. Hamilton–Jacobi (HJ) reachability analysis provides a rigorous mathematical framework for computing such controllers with formal guarantees [4,5]. In particular, backward reachable tubes (BRTs) yield not only the set of states from which a target can be reached [6], but also the corresponding optimal control. However, a major drawback of this approach lies in its computational cost: solving HJ partial differential equations (PDEs) on high-dimensional grids is infeasible for many real-world systems due to the curse of dimensionality [7]. This bottleneck is especially severe when the target is far, as the BRT must expand over a large domain to include the initial state.

In contrast, forward reachable tubes (FRTs) describe the evolution of reachable states over time from a given initial condition [5]. They can be computed efficiently, especially for linear systems, using set-propagation methods based on zonotopes [8,9]. However, FRTs are often limited to verification and open-loop planning, rather than feedback control synthesis [9].

Indeed, several approaches have been developed to efficiently approximate reachable sets for nonlinear and uncertain systems. Chen *et al*. [6] proposed decomposing reachable sets into lower-dimensional components to manage the computational burden of high-dimensional systems. Herbert *et al*. [4] introduced efficient initializations for reachability analysis to provide real-time safety guarantees. Zonotopes have emerged as a popular choice for representing reachable sets due to their closure under linear operations. Althoff and Frehse [10] and Althoff and Krogh [11] demonstrated how zonotopes can improve both accuracy and scalability. Bird *et al*. [12] extended these methods to hybrid zonotopes for use in explicit model predictive control (MPC). For stochastic systems, Gleason *et al*. [13] developed Lagrangian-based methods for underapproximating reach-avoid sets. Recent work has also leveraged data-driven and learning-based techniques: Alanwar *et al*. [14] proposed a robust data-driven predictive control scheme using matrix zonotopes, while Holaza *et al*. [15] revisited reachability-driven MPC for embedded systems. Other notable contributions include FaSTrack [16], a modular framework for fast and safe motion planning; classification-based reachability approximation [17]; and classical reachability analysis using support functions, ellipsoids [18], or optimal control [19].

Manuscript received June 16, 2025; revised August 22, 2025; accepted August 29, 2025. Recommended by Associate Editor Juhoon Back under the direction of Senior Editor Sangmoon Lee.

Isabelle El-Hajj, Jasper van Beers, and Prashant Solanki are with the Section of Control & Simulation at the Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, 2629HS, Delft, The Netherlands (e-mails: {i.z.el-hajj-1, j.j.vanbeers, p.solanki}@tudelft.nl).
* Corresponding author.

Collectively, these methods highlight a rich landscape of tools for reachability-based verification and control.

However, many of these approaches rely on a global (nonlinear) model of the system, which may not always be available. Moreover, for methods reliant on grid-based solvers, the (global) spatial and temporal resolution is constrained by the available computational resources—even for offline estimations of the reachable sets.

Therefore, we propose a hybrid reachability-based controller for nonlinear systems that combines the computational efficiency of FRTs with the control-theoretic power of BRTs. In contrast to existing work, our contribution lies in a practical controller synthesis framework that 1) avoids the need for global reachability computation, 2) operates using only local linear models circumventing the need for a global nonlinear model, and 3) generates a feedback control strategy that can steer the system to a target while respecting constraints. This is achieved by strategically combining zonotope-based FRT propagation with local BRT computation in a receding-horizon fashion. At each step, we use a local linear model of the dynamics at the current state and compute the FRT of this model using zonotope-based propagation. Among the states in the FRT, we identify the one closest to the final goal. We then compute a local BRT on a small grid around this intermediate state using the continuous-time linearized dynamics. If the current state lies within this local BRT, we determine the associated optimal control, apply it to move toward the intermediate state, and repeat the procedure. The result is a method that provides pseudo-optimal control for nonlinear systems reliant only on a family of linear representations of the system (e.g., linear parameter-varying models).

We note that there are many conceptual similarities between our proposed approach and MPC formulations. However, there are some nuanced differences between the two. The most important of these distinctions lies in the treatment of the cost/value function. MPC relies on a known (user-tunable) cost function in order to derive a sequence of control actions that minimize the cost. In contrast, our approach attempts to find the initially unknown reach cost value function which, once found, can be used to determine the time-optimal control sequence needed to reach the target set/state from the current point. Thus, the control strategy is obtained directly from the value function, rather than from a trajectory optimization routine.

The structure of the paper is as follows: Section 2 formally states the control problem of interest. Section 3 reviews key relevant concepts. Section 4 introduces our proposed ZonoReach controller. Section 5 describes the benchmark systems used to evaluate our method and Section 6 presents the results. Finally, Section 7 offers concluding remarks and highlights opportunities for extending the approach to real-world systems.

A preliminary version of this work was presented at the ICCAS 2025 conference [20]. This paper provides a more detailed exposition with expanded explanations.

2. PROBLEM STATEMENT

We consider a general nonlinear control system

$$\dot{x}(t) = f(x(t), u(t)),\ x(t) \in \mathbb{R}^n,\ u(t) \in \mathcal{U} \subset \mathbb{R}^m, \quad (1)$$

where $x(t)$ is the state vector, $u(t)$ is the control input, and $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ denotes the system dynamics. Moreover, the control inputs are constrained within a compact, convex set $\mathcal{U}$ that is box-bounded

$$\mathcal{U} = \{u \in \mathbb{R}^m \,|\, \underline{u} \le u \le \bar{u}\}, \quad (2)$$

where $\underline{u},\ \bar{u} \in \mathbb{R}^m$ are element-wise lower/upper input bounds.

Given an initial state $x_0 \in \mathbb{R}^n$ and a compact target set $\mathcal{T} \subset \mathbb{R}^n$, the objective is to synthesize a feedback control strategy $u^*(t)$ that steers the system from x_0 to $\mathcal{T}$ within a finite time horizon T, while respecting input constraints and system dynamics.

Assumptions: To facilitate tractable control synthesis, we make the following assumptions:

1) **(Lipschitz continuity)**: We assume that $f(x,u)$ is locally Lipschitz continuous in x.
2) **(Smoothness)**: $f(x,u) \in \mathcal{C}^1$ in both arguments.
3) **(Bounded inputs)**: The admissible input set $\mathcal{U}$ is a convex zonotope or hyperrectangle centered at the origin.
4) **(Reachable goal)**: The target set $\mathcal{T}$ is reachable from x_0 under some admissible control input.
5) **(Availability of local linear models)**: We assume access to representative local linear dynamics of the nonlinear system about the current operating point: $x_0,\ u_0$. This may be obtained, for example, via a first order Taylor expansion.

Assumptions 1, 2, and 4 are standard assumptions in nonlinear control and reachability analysis problems [21-24]. As for Assumption 3, actuator limits are inevitable for real-world systems [25]. Assumption 4 rules out trivial infeasible problems. Finally, Assumption 5 allows us to investigate our approach without introducing the additional complexity of model identification.

3. BACKGROUND

We briefly introduce the tools that form the basis of our approach: Hamilton–Jacobi reachability, forward and backward reachable tubes, and zonotope-based set propagation.

3.1. Hamilton–Jacobi reachability analysis

HJ reachability provides a rigorous framework for computing reachable tubes. For instance, the BRT characterizes the sets of initial states from which a target set can be reached within a finite time horizon under admissible control inputs. Likewise, the FRT describes the sets of future states that are reachable over a finite time horizon from a given initial condition under admissible control inputs. These computations are typically framed as differential games and solved using dynamic programming principles.

Let $l(x)$ be a continuous function such that the target set $\mathcal{T} \subset \mathbb{R}^n$ is defined as the sub-zero level set of $l(x)$

$$\mathcal{T} = \{x \in \mathbb{R}^n \mid l(x) \leq 0\}. \tag{3}$$

Given a time horizon $\tau \in [t, T]$, the BRT of (1) is

$$\mathcal{B}(t;\mathcal{T}) = \Big\{x_0 \in \mathbb{R}^n \,\Big|\, \exists u(\cdot):[t,T] \to \mathcal{U}, \\ \exists \tau \in [t,T] \text{ s.t. } \phi^u_{x_0,t}(\tau) \in \mathcal{T}\Big\}, \tag{4}$$

where $\phi^u_{x_0,t}(\tau)$ denotes the state trajectory under control $u(\cdot)$, starting from x_0 at time t. In the context of reach problems, we are concerned with finding the trajectories that reach the target set by the terminal time, T. Accordingly, the BRT at T may be derived from the sub-zero level set of a value function, $V(x,t)$, that describes the reach cost of the trajectories which terminate in the target set within the time horizon $\tau \in [t, T]$

$$\mathcal{B}(t;\mathcal{T}) = \{x \in \mathbb{R}^n \mid V(x,t) \leq 0\}, \tag{5}$$

where $V(t,x)$ is given by

$$V(x,t) = \inf_{u(\cdot)} \min_{\tau\in[t,T]} l\big(\phi^u_{x,t}(\tau)\big). \tag{6}$$

This value function satisfies the final-time Hamilton–Jacobi–Isaacs variational inequality (HJI-VI)

$$\min\left\{ \frac{\partial V}{\partial t}(x,t) + H(x,t),\ l(x) - V(x,t) \right\} = 0, \\ V(x,T) = l(x), \tag{7}$$

where the Hamiltonian $H(x,t)$ is defined as

$$H(x,t) = \min_{u\in\mathcal{U}} \langle \nabla V(x,t), f(x,u)\rangle. \tag{8}$$

The BRT contains all states from which the system can be driven into the target set $\mathcal{T}$ within the time horizon $[t,\ T]$, under some admissible control signal. Moreover, the time-optimal feedback control (on the BRT boundary) is given by

$$u^*(x,t) = \arg\min_{u\in\mathcal{U}} \langle \nabla V(x,t), f(x,u)\rangle. \tag{9}$$

HJ analysis can also be used to compute FRT by time-reversing the dynamics and target condition, yielding

$$\mathcal{F}(t;x_0) = \Big\{x(t) \,\Big|\, x(0) = x_0,\ \dot{x}(s) = f(x(s),u(s)), \\ u(s) \in \mathcal{U},\ s \in [0,t]\Big\}. \tag{10}$$

However, while the FRT characterizes which states are reachable, it does not provide a controller to steer the system to a specific point within the set. In general, determining the input signal required to reach a particular state in the FRT is a nontrivial inverse problem, and may not admit a unique or easily synthesized solution. For this reason, the FRT is typically used for reachability certification and safety analysis, but not for direct control synthesis.

Therefore, in this work, we use forward reachability for intermediate target selection and backward reachability to construct locally optimal control actions.

3.2. Zonotope-based set propagation

For locally linear models, the FRT can be efficiently computed using zonotope-based set propagation. A zonotope $\mathcal{Z}$ is defined as

$$\mathcal{Z} = \left\{ c + \sum_{i=1}^{m} \alpha_i g_i \,\middle|\, \alpha_i \in [-1,\ 1] \right\}, \tag{11}$$

where $c \in \mathbb{R}^n$ is the center of the zonotope, $g_i \in \mathbb{R}^n$ are the generator vectors of the zonotope, and α_i are the coefficients of the generators. Zonotopes are closed under linear transformations and Minkowski sums, enabling efficient computation of reachable sets for linear systems.

Continuous-time dynamics: Consider an affine linear time-invariant (LTI) system of the form

$$\dot{x} = Ax + Bu + E,\ u \in \mathcal{U}, \tag{12}$$

where $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ are the state and input matrices, respectively, and $\mathcal{U}$ is a compact control input set. $E \in \mathbb{R}^n$ is an offset matrix. Though $E = 0$ in many linear systems, such an offset matrix may appear by linearizing a nonlinear system about an operating point (x_0, u_0). Consider, for example, a first-order Taylor expansion of the system dynamics given in (1)

$$f(x,u) \approx f(x_0,u_0) + A(x - x_0) + B(u - u_0), \tag{13}$$

yielding a local affine model of the form

$$\dot{x} = Ax + Bu + E, \tag{14}$$

with

$$E = f(x_0,u_0) - Ax_0 - Bu_0. \tag{15}$$

Discrete-time dynamics: For discrete-time systems of the form

$$x_{k+1} = A_d x_k + B_d u_k + E_d,\ u_k \in \mathcal{U}, \tag{16}$$

the forward reachable set (FRS) over a finite horizon can be computed iteratively using the recurrence

$$\mathcal{X}_{k+1} = A_d\mathcal{X}_k \oplus B_d\mathcal{U} \oplus E_d,\ \mathcal{X}_0 = \text{initial set}, \tag{17}$$

where $\mathcal{X}_k$ is the reachable state set at step k.

Since zonotopes are closed under affine maps and Minkowski sums, each $\mathcal{X}_k$ remains a zonotope. This property allows for fast and scalable propagation of reachable sets in discrete-time, especially when combined with zonotope order reduction techniques to mitigate complexity growth [9]. Note that, in our current implementation as a proof-of-concept, such order reduction is not conducted and is reserved for future work.

4. ZONOREACH CONTROLLER

In this section, we present ZonoReach, a control framework that synthesizes locally optimal control actions for nonlinear systems through forward and backward reachability analysis. The key idea is to enable nonlinear control using only local linear models and short-horizon reachability computations. The advantages are twofold: first, we do not require explicit knowledge of the nonlinear dynamics. Second, ZonoReach operates on local spaces, avoiding the estimation of the value function along a global grid, which is computationally prohibitive in high dimensions. Our method, illustrated in Fig. 1, consists of four main components: 1) local FRT construction via zonotope propagation on discrete-time linearized dynamics, 2) geometric selection of an intermediate goal along the contour of the FRT closest to the final target, 3) BRT computation about the intermediate goal, and 4) greedy control selection using the BRT gradients.

4.1. Forward reachable tube using zonotope-based set propagation

To approximate the FRT, we assume access to local linear approximations of the system dynamics around the current state and input, given by A , B, and E matrices, as discussed in Subsection 3.2.

With zero-order hold, the discretized matrices of (16) are obtained as shown in (18)

$$\begin{aligned} A_d &= e^{A\Delta t}, \\ B_d &= A^{-1}(e^{A\Delta t} - I)B, \\ E_d &= A^{-1}(e^{A\Delta t} - I)E, \end{aligned} \tag{18}$$

where Δt is the time step.

We propagate the dynamics over k steps using (17) to obtain the FRS. The FRT is then approximated by the union of these FRSs.

4.2. Intermediate target based on geometric construction

The FRT contains many states, but not all are equally useful for reaching the final target. To focus the control effort, we identify the point on the contour of the FRT closest to the target as an intermediate subgoal, which we denote by x_{int}. To determine it in practice, we evaluate the proximity of the target to each FRS composing the FRT. We denote the boundary of the i-th FRS by Ω_i. For each Ω_i, we identify the point that yields the smallest Euclidean distance between the target, denoted by x_{target}, and the points on Ω_i. Then x_{int} corresponds to the boundary point that results in the smallest overall distance

$$x_{\text{int}} = \arg \min_{x\in\{\Omega_i\}} \|x - x_{\text{target}}\|_2. \tag{19}$$

This construction ensures that the intermediate point lies on the contour of one of the FRSs which compose the FRT and can be reached by the current dynamics and control constraints.

4.3. Local value estimation using backward reachability

We compute the BRT about the intermediate target x_{int} to determine the locally time-optimal control input that steers the system toward this point. We use continuous-time HJ reachability for this step (i.e., (7) and (8)), applying it to the linearized dynamics about x_{int}. For linear dynamics, the Hamiltonian (8) becomes

$$\begin{aligned} H(x,t) &= \min_{u\in\mathcal{U}} p^\top(Ax + Bu) \\ &= p^\top Ax + \min_{u\in\mathcal{U}} \left(u^\top B^\top p\right)^\top. \end{aligned} \tag{20}$$

Since the control set $\mathcal{U}$ is a box, the result is bang-bang control

$$u^*(x) = -u_{\max} \cdot \text{sign}(B^\top \nabla V(x)). \tag{21}$$

Note that though (21) is the theoretical optimal control for linear systems with box-bounded inputs, the presence of the sign$(\cdot)$ function may induce chattering in practice.

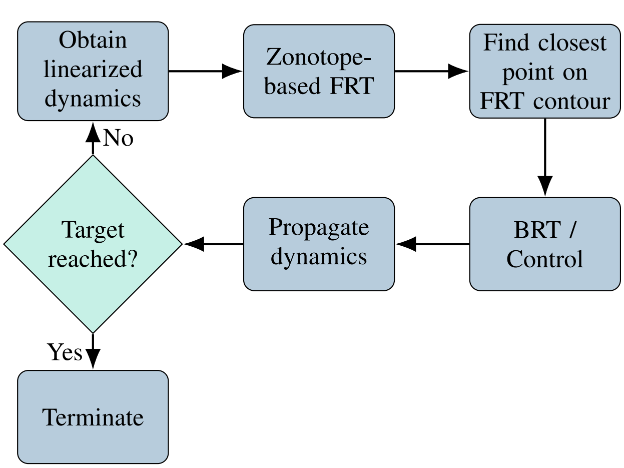

Fig. 1. Flowchart showing the main steps of ZonoReach.

Algorithm 1: ZonoReach.

```
while current state has not reached target set do
    Obtain local state and input matrices (A, B, E)
    Discretize dynamics to obtain (A_d, B_d, E_d)
    Compute FRT over time horizon T = kΔt
    if target ∈ FRT then
        Set x_int = x_target
    else
        Compute x_int as closest point in FRT to target
    end if
    Solve BRT from x_int using continuous-time HJ on
  linearized model
    for i = 1 to n_c do
        Compute ∇V(x) at current state
        Determine u* using (21)
        Apply control u* and update state
    end for
end while
```

Chattering may arise due to noise, uncertainties, or numerical errors within the representation of the value function. Chattering poses several disadvantages that can be especially detrimental for real-world applications including wear and tear on the hardware, energy inefficiency, and degraded performance. Chattering can be mitigated through, for instance, the use of low-pass filters applied to the control signal or hysteresis [26].

4.4. Greedy action selection and replanning

Once the BRT contains the current state, we determine the control action based on the spatial gradient of $V(x,t)$ at the current state. Rather than re-computing the reachable tubes at every time step, the BRT is used over a certain control horizon, which we denote by n_c. During this period, at each step, the gradient of the $V(x,t)$ is evaluated at the current state to determine the optimal control input u^*, which is then applied to the system. This process is repeated for n_c steps before proceeding to the next iteration of the algorithm. Algorithm 1 summarizes the steps of our ZonoReach controller.

5. CASE STUDIES

To demonstrate the effectiveness and versatility of our proposed controller, we apply it to three representative dynamical systems: the double integrator, Dubins car, and the inverted pendulum. These systems are chosen to highlight different challenges in nonlinear control, such as linear versus nonlinear dynamics, holonomic versus nonholonomic constraints, and stability around equilibrium points. Each example showcases a distinct aspect of the controller's capabilities and sheds light on some limitations. We also compare our approach against standard MPC designs to highlight the differences in performance between MPC and ZonoReach. In what follows, we describe the system models that serve as testbeds for our controller.

5.1. Double integrator

As a baseline, we consider the double integrator system by virtue of its simplicity and the availability of an analytical bang-bang solution. This serves as a ground-truth benchmark for assessing our method's ability to approximate time-optimal behavior.

The equations of the double integrator where the control is bounded are shown in (22).

$$\dot{x}_1 = x_2,\ \dot{x}_2 = u,\ |u| \leq u_{\max}. \tag{22}$$

In the case of bounded control, the time-optimal control law is a bang-bang controller. That is, the control takes on its extreme values and switches once at a critical time t_s, which may be computed by solving a quadratic equation [27] that depends on the initial and final states. The optimal control is expressed as

$$u(t) = \begin{cases} u_{\max}, & t < t_s, \\ -u_{\max}, & t \geq t_s, \end{cases} \quad \text{or the reverse.} \tag{23}$$

5.2. Inverted pendulum

The inverted pendulum is a canonical benchmark in nonlinear control. It models a rigid rod pivoted at one end, free to rotate in the vertical plane, with the state defined by the angular position $\theta \in [-\pi, \pi)$ and angular velocity $\dot{\theta}$. The system is underactuated, with a single control input in the form of a torque applied at the pivot, which cannot directly actuate all state variables.

The system's nonlinear dynamics, under gravitational and damping effects, are given by

$$\ddot{\theta} = \frac{T - b\dot{\theta} + mgl\sin\theta}{ml^2/3}, \tag{24}$$

where m denotes the mass of the pendulum, l is its length, g is the gravitational acceleration, b is a viscous damping coefficient, and T is the torque. The pendulum mass is set to $m = 1$ kg, the length is $l = 1$ m, the viscous damping coefficient is $b = 0.05$ kg m^2/s, and T is constrained to $[-2, 2]$ Nm, reflecting realistic actuator limitations. In our experiments, we focus on stabilization around the upright equilibrium.

5.3. Dubins car

Furthermore, we consider Dubins car, a classical example of a nonholonomic and nonlinear system. The state of the Dubins car is described by the position $(x, y) \in \mathbb{R}^2$ and the heading angle $\theta \in [-\pi, \pi)$ rad. The control input $u \in \mathcal{U} = [-1, 1]$ rad/s represents the bounded turning rate, while the forward velocity is held constant at $V = 1$ m/s. The system evolves according to the following dynamics:

$$\begin{aligned}\dot{x} &= V\cos(\theta),\\ \dot{y} &= V\sin(\theta),\\ \dot{\theta} &= u(t),\end{aligned} \tag{25}$$

where $u(t)$ is a measurable function constrained to lie within the control bounds for all time. These dynamics are subject to a nonholonomic constraint as the vehicle cannot move laterally. Despite its low dimensionality, the Dubins car presents notable challenges for long-horizon planning and feedback control, particularly under state and input constraints. This makes it a valuable testbed for validating reachability-based control methods.

6. RESULTS AND DISCUSSION

This section presents and analyzes the results obtained from applying our proposed ZonoReach controller to each of the case study systems. Table 1 summarizes the experimental configuration for each benchmark system. Note that for the double integrator, some units are omitted since its state variables admit versatile interpretations.

Across all experiments, FRTs were computed using custom routines that leverage the zonotope class from the CORA toolbox [28], while BRTs were obtained using the HelperOC toolbox [29] with additional integration into our control loop. As such, the current implementation cannot be used in real-time as it still relies on a grid-based solver to estimate the BRT, even though this grid is over a smaller domain. However, we emphasize that the ZonoReach algorithm itself does not require a grid-solver to run. Thus, if, for example, the BRT is computed via the Hopf–Lax formula [22,30-32], the value function could be evaluated without discretizing the state space, replacing grid-based PDE solvers with variational or closed-form characterizations and thereby improving scalability.

Double integrator: For the double integrator, the ZonoReach controller successfully drives the state from $[0.5, 0.25]$ m to the target set, as shown in Fig. 2. With a FRT horizon of 50 steps, the trajectory converges swiftly to the target, demonstrating behavior consistent with the known time-optimal bang-bang solution, which is also plotted in Fig. 2.

Table 1. Experimental setup for each benchmark system.

	Double integrator	Inverted pendulum	Dubins car
x_0	[0.5, 0.25]	$[\frac{\pi}{12}$ rad, 0 $\frac{\text{rad}}{\text{s}}]$	$[-2$ m, 0.5 m, $\frac{\pi}{6}$ rad$]$
x_{target}	[0, 0]	[0 rad, 0 $\frac{\text{rad}}{\text{s}}$]	[0 m, 0 m, 0 rad]
Target radius	0.1	0.05	0.5
k	50 steps	15 steps	40 steps
Δt	0.01 s	0.01 s	0.01 s
$\mathcal{U}$	$[-1, 1]$	$[-2, 2]$ Nm	$[-1, 1]$ $\frac{\text{rad}}{\text{s}}$
n_c	5 steps	5 steps	10 steps

Moreover, as a baseline comparison, we also task a standard linear MPC with the same goal. To make the comparison fair, we do not penalize the control effort in the MPC cost function and only specify the standard penalty for state deviations through xQx^T. Moreover, the same planning and control horizons employed by ZonoReach (see Table 1) are used by the MPC baselines. For one MPC, we choose Q as the identity matrix, and for another, we take Q as the diagonal matrix of $[1.0, 0.1]$. The resultant trajectories and the control effort are shown in Fig. 2 (solid and dashed purple lines).

Although MPC with $Q = I$ produces a smoother trajectory toward the target set, it is the slowest to reach the target set with $t = 2.51$ s. ZonoReach arrives almost twice as fast ($t = 1.58$ s), which is near the time-optimal value ($t = 1.54$ s). However, this comes at the cost of chattering in the control signal, which arises due to numerical errors in the value function representation for near-zero gradients. Such chattering may be circumvented by using a low-pass filter on the control signal, at the cost of a longer time-to-reach. It is interesting to note that all approaches initially follow the same trajectory before diverging. The divergence for ZonoReach occurs due to the aforementioned numerical errors at near-zero gradients. In contrast, the MPC breakaway is governed by the cost function. If the cost penalizes positional deviations more heavily than velocity as is the case for MPC 2, then the time-optimal trajectory gets followed more closely, as seen in Fig. 2.

Inverted pendulum: Here, the goal of the controller is to stabilize the pendulum near the (unstable) upright position. Again, we compare ZonoReach to a standard nonlinear model predictive control (NMPC) design (implemented via the `nlmpc` function in MATLAB R2021b). Furthermore, we provide the NMPC design with perfect knowledge of the nonlinear dynamics. As the prediction horizon can affect the performance of the NMPC, especially for a system like the inverted pendulum with a nontrivial trajectory to the desired state, we compare two different planning horizons. The first mirrors that of ZonoReach (15 steps); whereas, the second enjoys a longer horizon (100 steps).

Fig. 3 depicts the resultant trajectories of the NMPC designs and ZonoReach alongside their control sequences. Although all controllers initially follow the same trajectory, ZonoReach not only arrives at the target set faster than either NMPC formulation, but also does not overshoot the origin. While it is possible to augment the cost function of the NMPC to penalize overshoot and promote time-to-reach performance, defining such a cost function is not always trivial. Alternatively, NMPC parameters can be tuned in such a way that reduces the overshoot and improves the time-to-reach performance; however, this can involve considerable tuning effort and such additional tun-

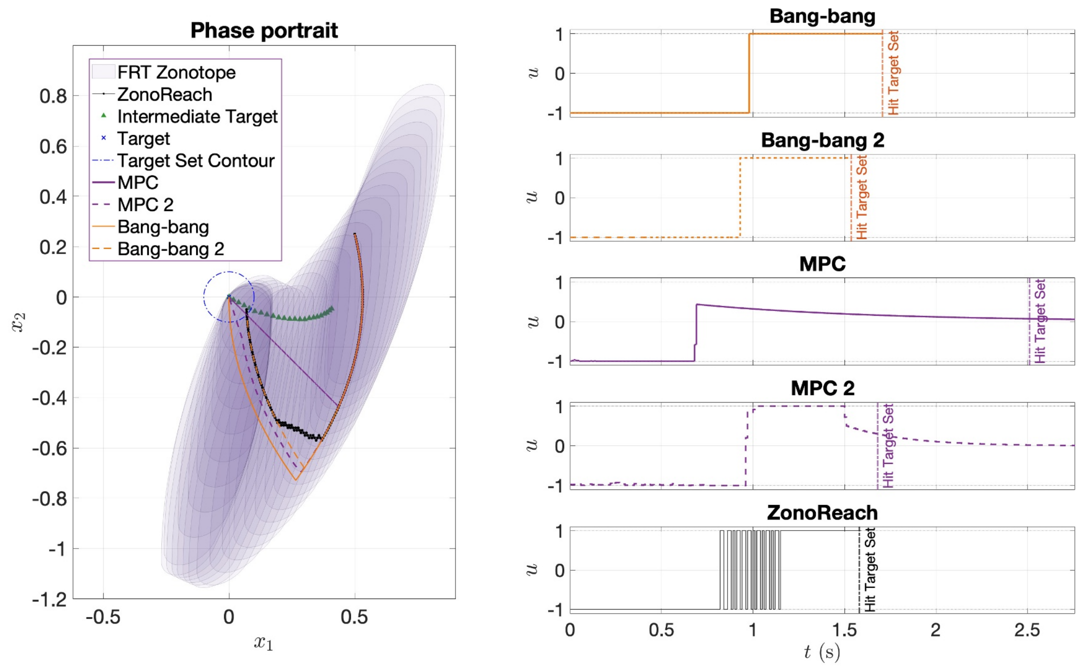

Fig. 2. Double integrator performance. The left panel depicts the target set, the ZonoReach & baseline controller trajectories, and FRT (for clarity, we plot only the FRS containing the closest point to the target among all FRSs pertaining to the FRT). The right panel illustrates the control effort and time-to-reach performance of the ZonoReach (bottom plot) and baseline controllers (top four plots). These baseline controllers include: 1) the analytically derived time-optimal trajectory to the origin (first row) and point of entry of ZonoReach (second row), and 2) a model predictive control design with Q as identity matrix (third row) and model predictive control with Q with diagonal entries $[1, 0.1]$ (fourth row).

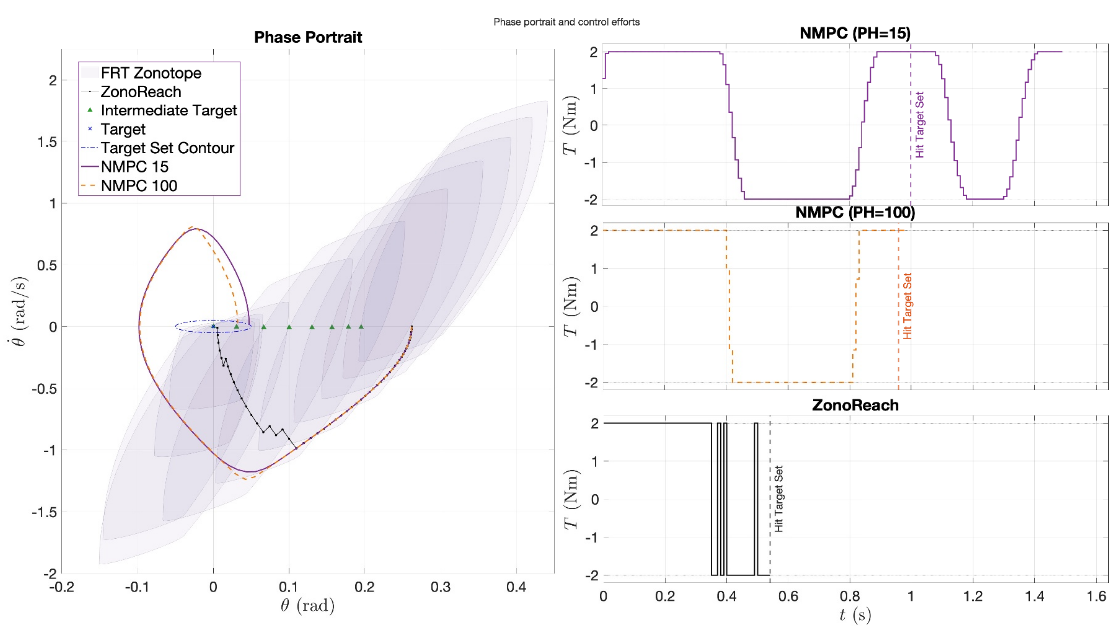

Fig. 3. Inverted pendulum performance. Shown is the evolution of the trajectories that bring the inverted system into the target set for the baseline NMPC and ZonoReach controllers. We also depict the FRT (for clarity, we plot only the FRS containing the closest point to the target among all FRSs pertaining to the FRT) and intermediate target employed by the ZonoReach controller. We also show two NMPC trajectories: one with planning horizon of 15 steps (same as for ZonoReach) and the other with a planning horizon of 100 steps.

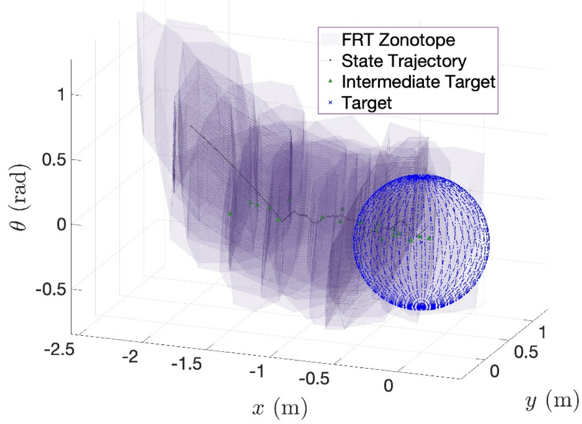

Fig. 4. Dubins car performance. Shown is the evolution of the state trajectories which bring the system into the target set with the ZonoReach controller. Also shown is the FRT (for clarity, we plot only the FRS containing the closest point to the target among all FRSs pertaining to the FRT).

ing is unnecessary for ZonoReach.

ZonoReach reaches the target set in $t = 0.54$ s; whereas, for the two NMPC controllers, the time needed is 0.96 s and 1.0 s. This is an impressive result given the fact that the ZonoReach only operates on local linear approximations of the system and uses comparatively short planning horizons. In fact, we observe that shorter FRT horizons lead to more consistent convergence for ZonoReach. This is because longer horizons tend to introduce greater discrepancies between the true nonlinear dynamics and the local linear approximations, degrading control performance. Thus, for highly nonlinear systems, it is favorable to keep the planning and control horizons short such that they remain within the domain of validity of the local linear approximations of the nonlinear dynamics. This limits the propagation of the approximation errors in the successive linearizations.

Dubins car: For the Dubins car, the controller generates curved trajectories consistent with Dubins car kinematics (Fig. 4). Similar to the inverted pendulum, shorter horizons result in more reliable tracking. Longer horizons increase the mismatch between the linear model and the true dynamics, making it harder to synthesize effective control actions.

Discussion: Through these experiments, we found that a key factor influencing performance was the choice of planning and control horizons for both forward and backward reachability computations. Preliminary observations suggest that: for linear systems, longer horizons lead to smoother and more direct trajectories. Conversely, for nonlinear systems, longer horizons may degrade performance due to the growing mismatch between the linearized and true dynamics. These observations highlight the potential importance of horizon selection in reachability-based planning. While longer horizons offer greater foresight, they are limited by model fidelity. Similarly, longer control horizon intervals reduce computational cost but may cause the controller to act on outdated reachable tubes, especially under rapidly changing dynamics.

Although these trends are consistent with intuition, a systematic study of horizon selection and model mismatch remains for future work.

Furthermore, when access to models is available, the control designer may decide between MPC and ZonoReach depending on their control task priorities. The examples of the double integrator and the inverted pendulum illustrate the nuanced difference between MPC and ZonoReach in their use of the cost and value functions, respectively: in MPC, the cost function is used to synthesize control actions; whereas, ZonoReach derives a reach-cost value function from which the time-optimal control can be determined. Therefore, if the designer prioritizes smooth and limited control effort, then MPC offers a viable path forward. However, note that augmenting a cost function and tuning parameters to produce the desired control may not be trivial. In contrast, if time-optimality is essential, then ZonoReach offers an attractive framework for reliably synthesizing such policies. However, ZonoReach is limited in its tendency to produce chattering. Despite this, chattering can be limited by employing mitigation strategies, such as low-pass filters.

Another limitation of the current proof-of-concept implementation is the reliance on a grid solver for the BRT computation. To this end, several recent works propose alternatives such as the Hopf–Lax formula [30-32], which provides closed-form or approximate characterizations that circumvent explicit gridding of the state space. Such extensions are the subject of future work.

7. CONCLUSION

We have presented a novel hybrid control approach for nonlinear systems that combines local linear models, zonotope-based forward reachability, and Hamilton–Jacobi-based backward reachability analysis. Our controller iteratively constructs a FRT using the local linear discrete-time dynamics and zonotope propagation, selects the point in the FRT closest to the final target, and solves a local backward reachability problem to determine the optimal control input.

Through numerical experiments on the double integrator, inverted pendulum, and Dubins car, we observe that our approach can drive the system toward the target under nonlinear and nonholonomic dynamics in a near time-optimal manner. Most notably, our approach can reach a target set of states faster than a comparable NMPC design which enjoys perfect model knowledge. This is achieved even with relatively short planning horizons. In fact, our results suggest that shorter time horizons may be beneficial for nonlinear systems, where linearization errors can accumulate over longer intervals, while linear systems tend to benefit from longer planning horizons as linearization error is absent.

This work opens several promising avenues for future research. A key assumption is perfect access to local linear models; relaxing this to allow estimation from observed data would improve applicability in real-world settings.

Likewise, the current method assumes a fixed target and full state observability, and does not model disturbances or stochasticity. Extending the framework to handle moving targets, partial observability, and uncertainty would significantly enhance robustness and scope.

Moreover, the computational cost of solving the Hamilton–Jacobi variational inequality on a grid in the current implementation remains a limiting factor, particularly in higher-dimensional systems. Since our backward reachability computations are performed on linearized dynamics, future work will explore scalable linear reachability techniques by leveraging, for example, the Hopf–Lax formula.

Finally, providing theoretical guarantees on convergence, robustness, and safety would be a critical step towards deploying this method in safety-critical domains.

Overall, this work lays a foundation for a scalable and modular control framework that leverages local model-based reachability to enable nonlinear control, while remaining agnostic to the underlying nonlinear dynamics.

DECLARATIONS

Conflict of Interest

The authors declare no competing interests.

Authors' Contributions

Isabelle El-Hajj proposed the main idea, developed the methodology, implemented the code, and contributed to the writing of the manuscript. Jasper van Beers and Prashant Solanki contributed to the development of the approach and the preparation and writing of the manuscript. Prashant Solanki also provided technical expertise on reachability analysis. All authors reviewed and approved the final version.

Funding

This work was conducted as part of the authors' Ph.D. studies and did not receive any external funding.

REFERENCES

[1] B. Lantos and L. Márton, *Nonlinear Control of Vehicles and Robots*, Berlin, Germany: Springer, 2010.

[2] J. Iqbal, M. Ullah, S. G. Khan, B. Khelifa, and S. Čuković, "Nonlinear control systems–a brief overview of historical and recent advances," *Nonlinear Engineering*, vol. 6, no. 4, pp. 301-312, 2017.

[3] J. Guiochet, M. Machin, and H. Waeselynck, "Safety-critical advanced robots: A survey," *Robotics and Autonomous Systems*, vol. 94, pp. 43-52, 2017.

[4] S. L. Herbert, S. Bansal, S. Ghosh, and C. J. Tomlin, "Reachability-based safety guarantees using efficient initializations," *Proc. of the 58th IEEE Conference on Decision and Control (CDC)*, Nice, France, pp. 4810–4816, December 2019.

[5] I. M. Mitchell, "Comparing forward and backward reachability as tools for safety analysis," *Proc. of the International Workshop on Hybrid Systems: Computation and Control (HSCC)*, Pisa, Italy, pp. 428–443, April 2007.

[6] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin, "Decomposition of reachable sets and tubes for a class of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 11, pp. 3675-3688, 2018.

[7] S. Bansal and C. J. Tomlin, "DeepReach: A deep learning approach to high-dimensional reachability," *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, Xi'an, China, pp. 1817–1824, May 2021.

[8] A. Girard, "Reachability of uncertain linear systems using zonotopes," *Proc. of the International Workshop on Hybrid Systems: Computation and Control (HSCC)*, Zurich, Switzerland, pp. 291–305, March 2005.

[9] S. Kousik, *Reachability-based Trajectory Design*, Ph.D. thesis, University of Michigan, 2020.

[10] M. Althoff and G. Frehse, "Combining zonotopes and support functions for efficient reachability analysis of linear systems," *Proc. of the 55th IEEE Conference on Decision and Control (CDC)*, Las Vegas, NV, USA, pp. 7439–7446, December 2016.

[11] M. Althoff and B. H. Krogh, "Zonotope bundles for the efficient computation of reachable sets," *Proc. of the 50th IEEE Conference on Decision and Control (CDC) and European Control Conference (ECC)*, Orlando, FL, USA, pp. 6814–6821, December 2011.

[12] T. J. Bird, N. Jain, H. C. Pangborn, and J. P. Koeln, "Set-based reachability and the explicit solution of linear MPC using hybrid zonotopes," *Proc. of the American Control Conference (ACC)*, Atlanta, GA, USA, pp. 158–165, June 2022.

[13] J. D. Gleason, A. P. Vinod, and M. M. K. Oishi, "Underapproximation of reach-avoid sets for discrete-time stochastic systems via Lagrangian methods," *Proc. of the 56th IEEE Conference on Decision and Control (CDC)*, Melbourne, Australia, pp. 4283–4290, December 2017.

[14] A. Alanwar, Y. Stürz, and K. H. Johansson, "Robust data-driven predictive control using reachability analysis," *European Journal of Control*, vol. 68, 100666, 2022.

[15] J. Holaza, P. Bakaráč, and J. Oravec, "Revisiting reachability-driven explicit MPC for embedded control," *European Journal of Control*, vol. 78, 101019, 2024.

[16] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin, "FaSTrack: A modular framework for fast and guaranteed safe motion planning," *Proc. of the 56th IEEE Conference on Decision and Control (CDC)*, Melbourne, Australia, pp. 1517–1522, December 2017.

[17] V. Rubies-Royo, D. Fridovich-Keil, S. Herbert, and C. J. Tomlin, "A classification-based approach for approximate reachability," *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, Montreal, Canada, pp. 7697–7704, May 2019.

[18] A. B. Kurzhanski and P. Varaiya, "Ellipsoidal techniques for reachability analysis," *Proc. of the International Workshop on Hybrid Systems: Computation and Control (HSCC)*, Berkeley, CA, USA, pp. 202–214, March 2000.

[19] A. M. Bayen, I. M. Mitchell, M. M. K. Oishi, and C. J. Tomlin, "Aircraft autolander safety analysis through optimal control-based reach set computation," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 1, pp. 68-77, 2007.

[20] I. El-Hajj, J. van Beers, and P. Solanki, "Reachability-guided nonlinear control via zonotope propagation and local Hamilton–Jacobi analysis," *Proc. of the 25th International Conference on Control, Automation, and Systems (ICCAS 2025)*, pp. 1731-1736, 2025.

[21] H. K. Khalil and J. W. Grizzle, *Nonlinear Systems*, Prentice Hall, Upper Saddle River, NJ, 2002.

[22] L. C. Evans, *Partial Differential Equations*, American Mathematical Society, Providence, RI, 2022.

[23] J.-J. E. Slotine and W. Li, *Applied Nonlinear Control*, Prentice Hall, Englewood Cliffs, NJ, 1991.

[24] D. Bertsekas, *Dynamic Programming and Optimal Control: Volume I*, Athena Scientific, Belmont, MA, 2012.

[25] V. Kapila and K. Grigoriadis, *Actuator Saturation Control*, CRC Press, Boca Raton, FL, 2002.

[26] H. D. Taghirad, *Fundamentals of Robotics: Applied Case Studies with MATLAB® & Python*, CRC Press, Boca Raton, FL, 2025.

[27] A. Locatelli, *Optimal Control of a Double Integrator: A Primer on Maximum Principle*, Springer, Cham, Switzerland, 2016.

[28] M. Althoff and N. Kochdumper, "CORA 2016 manual," Technical Report, Technical University of Munich, December 2016.

[29] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton–Jacobi reachability: A brief overview and recent advances," *Proc. of the 56th IEEE Conference on Decision and Control (CDC)*, Melbourne, Australia, pp. 2242–2253, December 2017.

[30] J. Darbon and S. Osher, "Algorithms for overcoming the curse of dimensionality for certain Hamilton–Jacobi equations arising in control theory and elsewhere," *Research in the Mathematical Sciences*, vol. 3, no. 1, p. 19, March 2016.

[31] D. Lee and C. J. Tomlin, "A Hopf–Lax formula in Hamilton–Jacobi analysis of reach-avoid problems," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1055-1060, 2020.

[32] D. Lee and C. J. Tomlin, "Efficient computation of state-constrained reachability problems using Hopf–Lax formulae," *IEEE Transactions on Automatic Control*, vol. 68, no. 11, pp. 6481-6495, 2023.

Isabelle El-Hajj received her B.Sc. (*cum laude*) and M.Sc. (*cum laude*) degrees in aerospace engineering from Delft University of Technology, Delft, the Netherlands in 2017 and 2020, respectively. She is currently a Ph.D. candidate in the Section of Control & Simulation at Delft University of Technology. Her research explores synergies between model-based reinforcement learning, approximate dynamic programming, and reachability analysis. Her broader research interests lie at the intersection of machine learning and control theory.

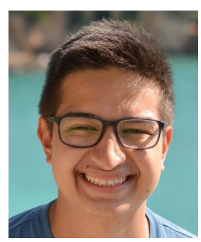

Jasper van Beers received the B.Sc. and M.Sc. degrees in aerospace engineering from Delft University of Technology, Delft, the Netherlands in 2018 and 2021 respectively. He is currently a Ph.D. candidate in the section of control & simulation at Delft University of Technology. His research focuses on unmanned aerial vehicle safety, in particular, the prediction of loss-of-control. His research interests broadly include robot safety, non-linear and complex systems, machine learning, and cognitive science.

Prashant Solanki received his B.Tech. (aerospace engineering) and M.Tech. (avionics) degrees Amity university, Noida, Uttar pradesh, India in 2017. Then he received an M.Sc. degrees in aerospace engineering (control & simulation) from Delft University of Technology, Delft, the Netherlands in 2020. He is currently a Ph.D. candidate in the section of control & simulation at Delft University of Technology. His research focuses on reachability analysis and safe flight envelope estimation of damaged unmanned aerial vehicle. His research interests broadly include robot safety, non-linear control, and machine learning.

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