

Reach-Avoid Problems with Time-Varying Dynamics, Targets and Constraints*

Jaime F. Fisac, Mo Chen, Claire J. Tomlin, and S. Shankar Sastry
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720, USA
{jfisac, mochen72, tomlin, sastry}@eecs.berkeley.edu

ABSTRACT

We consider a reach-avoid differential game, in which one of the players aims to steer the system into a target set without violating a set of state constraints, while the other player tries to prevent the first from succeeding; the system dynamics, target set, and state constraints may all be time-varying. The analysis of this problem plays an important role in collision avoidance, motion planning and aircraft control, among other applications. Previous methods for computing the guaranteed winning initial conditions and strategies for each player have either required augmenting the state vector to include time, or have been limited to problems with either no state constraints or entirely static targets, constraints and dynamics. To incorporate time-varying dynamics, targets and constraints without the need for state augmentation, we propose a modified Hamilton-Jacobi-Isaacs equation in the form of a double-obstacle variational inequality, and prove that the zero sublevel set of its viscosity solution characterizes the capture basin for the target under the state constraints. Through this formulation, our method can compute the capture basin and winning strategies for time-varying games at virtually no additional computational cost relative to the time-invariant case. We provide an implementation of this method based on well-known numerical schemes and show its convergence through a simple example; we include a second example in which our method substantially outperforms the state augmentation approach.

1. INTRODUCTION

Dynamic reach-avoid games have received growing interest in recent years and have many important applications in engineering problems, especially concerning the control of

*This work has been supported in part by NSF under grant CNS-931843; by ONR under MURIs N0014-08-0696, N00014-09-1-1051, N00014-13-1-0341, and grant N00014-12-1-0609; by AFOSR under MURI FA9550-10-1-0567. J. F. Fisac has received funding from “la Caixa” Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
HSCC '15, April 14 - 16, 2015, Seattle, WA, USA
Copyright is held by owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3433-4/15/04 ...\$15.00
<http://dx.doi.org/10.1145/2728606.2728612>

strategic or safety-critical systems: in many scenarios, one must find a control action that will guarantee reaching a desired state while respecting a set of constraints, often in the presence of an unknown disturbance or adversary. Practical applications include collision avoidance, surveillance, energy management and safe reinforcement learning, in which targets typically describe desired waypoints or operating conditions and constraints can model obstacles in the environment or forbidden configurations.

In the two-player reach-avoid formulation, one seeks to determine the set of states from which one of the players (the attacker) can successfully drive the system to some target set, while keeping the state within some state constraint set at all times, regardless of the opposing actions of the other player (the defender)—this set is commonly referred to as the *capture basin*, *backwards reachable set* or simply *reach-avoid set* of the target under the constraints. In the absence of state constraints, reachability problems involving possibly time-varying target sets can be posed as a minimum (maximum) cost game where the players try to optimize the pointwise minimum over time of some metric to the target. In this case, the backwards reachable set can be obtained by finding the viscosity solution to the corresponding Hamilton-Jacobi-Isaacs (HJI) equation in the form of a variational inequality: this value function captures the minimum distance to the target that will be achieved by the optimal trajectory starting at each point, so the capture basin is characterized by the region of the state space where this minimum future distance is equal to, or less than, zero. Maximum cost control problems were studied in detail in [3], and extended to the two-player setting in [2]; these problems require separate treatment and are in general harder to analyze than ordinary Bolza problems with running and terminal cost. While computationally intensive, Hamilton-Jacobi approaches are practically appealing nowadays due to the availability of modern numerical tools such as [14, 17, 19, 23], which are able to solve the associated equations for these problems when the dimensionality is low.

If the game is played under state constraints, then the value function generally becomes discontinuous [20, 21], which leads to numerical issues. In the case of systems with time-invariant dynamics, targets and constraints, the approach in [6] characterizes the capture basin through an auxiliary value function that solves a modified Hamilton-Jacobi variational inequality. Although the new value function no longer captures the minimum distance from a trajectory to the target, the reach-avoid set is still given by the value function's subzero region. This makes it possible to effectively turn a

constrained final cost problem into an unconstrained problem with a maximum cost.

For problems with time-varying dynamics, targets and constraints, the approach proposed in [5] as an extension of [6] requires augmenting the state space with an additional dimension accounting for time; one can then transform time dependence into state dependence and apply the above methods to solve the fixed problem in the space-time state space. Unfortunately, this approach presents a significant drawback, since the complexity of numerical computations grows exponentially with problem dimensionality.

The main contribution of this paper is an extension of the Hamilton-Jacobi reach-avoid formulation to the case where the target set, the state constraint set and the dynamics are all allowed to be time-varying, enabling computation of the reach-avoid set at no significant additional cost relative to the time-invariant case. To this end, we formulate a double-obstacle HJI variational inequality, and prove that the zero sublevel set of its viscosity solution characterizes the desired reach-avoid set. We also provide a numerical scheme based on [17, 18] and implementation based on [13] to solve the variational inequality and verify the numerical solution using a simple example. We finish by showing that our method vastly outperforms techniques requiring state augmentation.

It should be noted that other authors have recently studied Hamilton-Jacobi equations with a double obstacle applied to different settings involving Bolza problems [7, 10]. To our knowledge, however, the present work constitutes the first analysis of double-obstacle Hamilton-Jacobi equations in the context of reachability problems. We also note that the results presented in this paper for differential games are readily applicable to optimal control problems.

2. PROBLEM FORMULATION

2.1 System Dynamics

Let $\mathcal{A} \subset \mathbb{R}^{n_a}$ and $\mathcal{B} \subset \mathbb{R}^{n_b}$ be nonempty compact sets and, for $t \leq T$, let \mathbb{A}_t and \mathbb{B}_t denote the collections of measurable¹ functions $a : [t, T] \rightarrow \mathcal{A}$ and $b : [t, T] \rightarrow \mathcal{B}$ respectively. We consider a dynamical system with state $x \in \mathbb{R}^n$, and two agents, player I and player II, with inputs $a(\cdot) \in \mathbb{A}_t$ and $b(\cdot) \in \mathbb{B}_t$ respectively. The system dynamics are given by the flow field $f : \mathbb{R}^n \times \mathcal{A} \times \mathcal{B} \times [0, T] \rightarrow \mathbb{R}^n$, which is assumed to be uniformly continuous, with

$$\begin{aligned} |f(x, a, b, t)| &< L, \\ |f(x, a, b, t) - f(\tilde{x}, a, b, t)| &\leq L|x - \tilde{x}|, \end{aligned} \quad (1)$$

for some $L > 0$ and all $t \in [0, T]$, $x, \tilde{x} \in \mathbb{R}^n$, $a \in \mathcal{A}$, $b \in \mathcal{B}$. Then for any initial time $t \in [0, T]$ and state x , under input signals $a(\cdot) \in \mathbb{A}_t$, $b(\cdot) \in \mathbb{B}_t$, the evolution of the system is determined (see for example [9], Chapter 2, Theorems 1.1, 2.1) by the unique continuous trajectory $\phi_{x,t}^{a,b} : [t, T] \rightarrow \mathbb{R}^n$ solving

$$\begin{aligned} \dot{x}(s) &= f(x(s), a(s), b(s), s), \text{ a.e. } s \in [t, T], \\ x(t) &= x. \end{aligned} \quad (2)$$

Note that this is a solution in the *extended sense*, that is, it satisfies the differential equation *almost everywhere* (i.e.

¹ A function $f : X \rightarrow Y$ between two measurable spaces (X, Σ_X) and (Y, Σ_Y) is said to be measurable if the preimage of a measurable set in Y is a measurable set in X , that is: $\forall V \in \Sigma_Y, f^{-1}(V) \in \Sigma_X$, with Σ_X, Σ_Y σ -algebras on X, Y .

except on a subset of Lebesgue measure zero). It will be useful to denote by \mathbb{X}_t the collection of all trajectories $\phi_{x,t}^{a,b} : [t, T] \rightarrow \mathbb{R}^n$ that solve (2) for some initial condition and input signals:

$$\begin{aligned} \mathbb{X}_t &:= \{ \phi_{x,t}^{a,b} : [t, T] \rightarrow \mathbb{R}^n, \text{ for } x \in \mathbb{R}^n, a(\cdot) \in \mathbb{A}_t, b(\cdot) \in \mathbb{B}_t \} \\ \phi_{x,t}^{a,b}(t) &= x, \frac{d}{dt} \phi_{x,t}^{a,b} = f(\phi_{x,t}^{a,b}(s), a(s), b(s), s) \text{ a.e. } s \in [t, T] \}. \end{aligned}$$

2.2 Target and Constraint Sets

Let $t = 0$ be the start time of the reachability game and $t = T > 0$ be the end time. Following the notation in [5, 6], let $d : \mathbb{R}^n \times 2^{\mathbb{R}^n} \rightarrow \mathbb{R}$ give the distance between a point x and a set \mathcal{M} under some norm $|\cdot|$ on \mathbb{R}^n , that is, $d(x, \mathcal{M}) := \inf_{y \in \mathcal{M}} |x - y|$. Further, for every $\mathcal{M} \subseteq \mathbb{R}^n$, let $d_{\mathcal{M}} : \mathbb{R}^n \rightarrow \mathbb{R}$ be the signed distance function to \mathcal{M} :

$$d_{\mathcal{M}}(x) := \begin{cases} d(x, \mathcal{M}), & x \in \mathbb{R}^n \setminus \mathcal{M}, \\ -d(x, \mathbb{R}^n \setminus \mathcal{M}), & x \in \mathcal{M}. \end{cases}$$

A set-valued map $\mathcal{M} : [0, T] \rightarrow 2^{\mathbb{R}^n}$ is upper hemicontinuous² if for any open neighborhood V of $\mathcal{M}(t)$ there is an open neighborhood U of t such that $\mathcal{M}(\tau) \subseteq V \forall \tau \in U$. We define the upper hemicontinuous set-valued maps $\mathcal{T}, \mathcal{K} : [0, T] \rightarrow 2^{\mathbb{R}^n}$ which respectively assign a target set $\mathcal{T}_t \subset \mathbb{R}^n$ and a constraint set $\mathcal{K}_t \subset \mathbb{R}^n$ to each time $t \in [0, T]$. Requiring that $\mathcal{T}_t, \mathcal{K}_t$ are closed for all t , we can construct the space-time sets

$$\mathbb{T} := \bigcup_{t \in [0, T]} \mathcal{T}_t \times \{t\}, \quad \mathbb{K} := \bigcup_{t \in [0, T]} \mathcal{K}_t \times \{t\},$$

which are then closed subsets of $\mathbb{R}^n \times [0, T]$ by the following lemma, stated as Lemma 2 in [5] (a proof, based on elementary topology, is presented here for completeness).

LEMMA 1. *Let $\mathcal{M} : [0, T] \rightarrow 2^{\mathbb{R}^n}$ be an upper hemicontinuous set-valued map with $\mathcal{M}(t) = \mathcal{M}_t$ closed in \mathbb{R}^n for all $t \in [0, T]$. Then the set $\mathbb{M} = \bigcup_{t \in [0, T]} \mathcal{M}_t \times \{t\}$ is closed in $\mathbb{R}^n \times [0, T]$.*

PROOF. *We prove the lemma by contradiction, recalling that a set is closed if and only if it contains all of its limit points. Suppose \mathbb{M} is not closed: then there exists a limit point (x, t) of \mathbb{M} that is not in \mathbb{M} , i.e. $x \notin \mathcal{M}_t$. Since \mathcal{M}_t is closed, $x \notin \mathcal{M}_t$ cannot be a limit point of \mathcal{M}_t , that is, under any metric on \mathbb{R}^n , there exists $r > 0$ such that \mathcal{M}_t and the open ball $B(x, r)$ are disjoint. On the other hand, as (x, t) is a limit point of \mathbb{M} , every open neighborhood of (x, t) must meet \mathbb{M} . In particular, any neighborhood of the form $B(x, r/2) \times U$, with U any open neighborhood of t , contains a point $(y, s) \in \mathbb{M}$, i.e. $y \in \mathcal{M}_s$. But then for the open neighborhood of \mathcal{M}_t given by the Minkowski sum $V := \mathcal{M}_t + B(0, r/2)$, we have that for all open neighborhoods U of t , $\exists s \in U$ such that there is $y \in \mathcal{M}_s \cap B(x, r/2)$, so $y \notin V$ and therefore $\mathcal{M}_s \not\subseteq V$. This directly contradicts upper hemicontinuity of the set-valued map \mathcal{M} . \square*

The closed sets \mathbb{T} and \mathbb{K} can then be implicitly characterized as the subzero regions of two Lipschitz continuous functions $l : \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}$ and $g : \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}$ respectively,

²Sometimes also called upper semicontinuous.

that is, $\exists L_l, L_g > 0 : \forall (x, t), (\tilde{x}, \tilde{t}) \in \mathbb{R}^n \times [0, T]$,

$$\begin{aligned} |l(x, t) - l(\tilde{x}, \tilde{t})| &\leq L_l |x, t) - (\tilde{x}, \tilde{t})|, \\ |g(x, t) - g(\tilde{x}, \tilde{t})| &\leq L_g |x, t) - (\tilde{x}, \tilde{t})|, \end{aligned} \quad (3)$$

so that

$$(x, t) \in \mathbb{T} \iff l(x, t) \leq 0, \quad (x, t) \in \mathbb{K} \iff g(x, t) \leq 0.$$

These functions always exist, since we can simply choose the signed distance functions $l(x, t) = d_{\mathbb{T}}(x, t)$ and $g(x, t) = d_{\mathbb{K}}(x, t)$, which are Lipschitz continuous by construction, i.e. they are the infimum of point-to-point distances.

Note that this definition of targets and constraints is flexible and allows one to formulate a variety of target and constraint behaviors, including changing topologies over time (e.g. a target splitting into multiple separate sets or disappearing entirely).

We say that a trajectory $\phi \in \mathbb{X}_t$ is *admissible* on $[t, t + \delta]$ for some $\delta > 0$ if for all $t \leq \tau \leq t + \delta$, it satisfies $\phi(\tau) \in \mathcal{K}_\tau$. The minimum value of l achieved by an admissible state trajectory in the course of the game determines its outcome (it will be zero or negative if the trajectory ever enters the target \mathbb{T}); we therefore refer to l as the *payoff function*. On the other hand, the maximum value of g reached by any trajectory determines whether or not it is admissible (it will be positive if the trajectory ever breaches the constraints \mathcal{K}_t); we call g the *discriminator function*.

2.3 Value and Strategies

We will adopt the arbitrary convention that player I seeks to minimize the outcome of the game, while player II tries to maximize it: that is, I is trying to drive the system into the target set, and II wants to prevent I from succeeding, possibly by driving the system out of the constraint set; we will refer to I as the *attacker* and II as the *defender*. For each trajectory $\phi_{x,t}^{a,b} \in \mathbb{X}_t$ we define the outcome of the game as the functional

$$\mathcal{V}(x, t, a(\cdot), b(\cdot)) = \min_{\tau \in [t, T]} \max \left\{ l(\phi_{x,t}^{a,b}(\tau), \tau), \max_{s \in [t, \tau]} g(\phi_{x,t}^{a,b}(s), s) \right\}. \quad (4)$$

The above expression is considering, for each time τ , the maximum between the current value of l and the greatest value of g reached so far by the trajectory; therefore this term will be less than or equal to zero for a given τ if and only if the system is in the target at time τ without ever having left the constraint set on $[t, \tau]$. If this situation takes place for any $\tau \in [t, T]$, player I wins the game; therefore, the minimum for all τ reflects whether player I wins at any point between t and the end of the game. We summarize this through the following proposition.

PROPOSITION 1. *The set of points x at time $t \in [0, T]$ from which the system trajectory $\phi_{x,t}^{a,b}(\cdot)$ under given controls $a(\cdot) \in \mathbb{A}_t, b(\cdot) \in \mathbb{B}_t$ will enter the target set at some time $\tau \in [t, T]$ without violating the constraints at any $s \in [t, \tau]$ is equal to the zero sublevel set of $\mathcal{V}(x, t, a(\cdot), b(\cdot))$. That is:*

$$\begin{aligned} &\{(x, t) \in \mathbb{R}^n \times [0, T] : \exists \tau \in [t, T], \\ &\quad \phi_{x,t}^{a,b}(\tau) \in \mathcal{T}_\tau \wedge \forall s \in [t, \tau], \phi_{x,t}^{a,b}(s) \in \mathcal{K}_s\} \\ &= \{(x, t) \in \mathbb{R}^n \times [0, T] : \mathcal{V}(x, t, a(\cdot), b(\cdot)) \leq 0\}. \end{aligned} \quad (5)$$

Note that this value function is negative when the trajectory starting at (x, t) reaches the target without *previously* breaching the constraints: it is agnostic to whether constraints are breached *after* the target has been reached. One

could formulate an alternative problem requiring that trajectories remain feasible for the entire duration of the game: in that case the game's outcome would instead be

$$\mathcal{W}(x, t, a(\cdot), b(\cdot)) = \max \left\{ \min_{\tau \in [t, T]} l(\phi_{x,t}^{a,b}(\tau), \tau), \max_{\tau \in [t, T]} g(\phi_{x,t}^{a,b}(s), s) \right\}. \quad (6)$$

This alternative problem is not the object of this paper, and we will restrict our attention to the problem described by (4).

Following [11, 12, 22, 25], we define the set of nonanticipative strategies for player II containing the functionals $\Lambda_t = \{\beta : \mathbb{A}_t \rightarrow \mathbb{B}_t \mid \forall s \in [t, T], \forall a(\cdot), \hat{a}(\cdot) \in \mathbb{A}_t, (a(\tau) = \hat{a}(\tau) \text{ a.e. } \tau \in [t, s]) \Rightarrow (\beta[a](\tau) = \beta[\hat{a}](\tau) \text{ a.e. } \tau \in [t, s])\}$. By allowing II to use nonanticipative strategies, we are giving it a certain advantage, since at each instant it can adapt its control input to the one declared by I. This information pattern leads to the upper value of the game, given by:

$$V^+(x, t) := \sup_{\beta(\cdot) \in \Lambda_t} \inf_{a(\cdot) \in \mathbb{A}_t} \mathcal{V}(x, t, a(\cdot), \beta[a](\cdot)). \quad (7a)$$

Analogously, we can decide to give I the advantage by defining its set of nonanticipative strategies as $\Gamma_t := \{\alpha : \mathbb{B}_t \rightarrow \mathbb{A}_t \mid \forall s \in [t, T], \forall b(\cdot), \hat{b}(\cdot) \in \mathbb{B}_t, (b(\tau) = \hat{b}(\tau) \text{ a.e. } \tau \in [t, s]) \Rightarrow (\alpha[b](\tau) = \alpha[\hat{b}](\tau) \text{ a.e. } \tau \in [t, s])\}$. This determines the lower value of the game as:

$$V^-(x, t) := \inf_{\alpha(\cdot) \in \Gamma_t} \sup_{b(\cdot) \in \mathbb{B}_t} \mathcal{V}(x, t, \alpha[b](\cdot), b(\cdot)). \quad (7b)$$

Naturally, it follows that $V^-(x, t) \leq V^+(x, t)$ everywhere. In those cases in which equality holds, the game is said to have value and $V(x, t) := V^+(x, t) = V^-(x, t)$ is simply referred to as the *value of the game*.

Given an information pattern, we say that a point (x, t) is in the *capture basin* (or *reach-avoid set*) $\mathcal{C}_{\mathbb{T}, \mathbb{K}}$ of the target \mathbb{T} under constraints \mathbb{K} when the system trajectory $\phi_{x,t}^{a,b}$, with both players acting optimally, reaches \mathbb{T} at some time $\tau \in [t, T]$ while remaining in \mathbb{K} for all time $s \in [t, \tau]$. In particular, when player II uses nonanticipative strategies,

$$\begin{aligned} \mathcal{C}_{\mathbb{T}, \mathbb{K}}^+ &:= \{(x, t) \in \mathbb{R}^n \times [0, T] : \exists a(\cdot) \in \mathbb{A}_t, \forall \beta(\cdot) \in \Lambda_t, \\ &\quad \exists \tau \in [t, T], \phi_{x,t}^{\alpha, \beta[a]}(\tau) \in \mathcal{T}_\tau \wedge \forall s \in [t, \tau], \phi_{x,t}^{\alpha, \beta[a]}(s) \in \mathcal{K}_s\}, \end{aligned} \quad (8a)$$

and similarly, when player I uses nonanticipative strategies,

$$\begin{aligned} \mathcal{C}_{\mathbb{T}, \mathbb{K}}^- &:= \{(x, t) \in \mathbb{R}^n \times [0, T] : \exists \alpha(\cdot) \in \Gamma_t, \forall b(\cdot) \in \mathbb{B}_t, \\ &\quad \exists \tau \in [t, T], \phi_{x,t}^{\alpha[b], b}(\tau) \in \mathcal{T}_\tau \wedge \forall s \in [t, \tau], \phi_{x,t}^{\alpha[b], b}(s) \in \mathcal{K}_s\}. \end{aligned} \quad (8b)$$

Given Proposition 1 and the above definitions, we have an important result expressed by the following proposition.

PROPOSITION 2. *The capture basin of the space-time target set \mathbb{T} when the defender (attacker) is allowed to use nonanticipative strategies is given by the zero sublevel set of the upper (resp. lower) value function V^\pm . That is:*

$$\mathcal{C}_{\mathbb{T}, \mathbb{K}}^+ = \{(x, t) \in \mathbb{R}^n \times [0, T] : V^+(x, t) \leq 0\}, \quad (9a)$$

$$\mathcal{C}_{\mathbb{T}, \mathbb{K}}^- = \{(x, t) \in \mathbb{R}^n \times [0, T] : V^-(x, t) \leq 0\}. \quad (9b)$$

COROLLARY 1. *The capture basin when the defender is allowed to use nonanticipative strategies is a subset of that resulting from the attacker using nonanticipative strategies:*

$$\mathcal{C}_{\mathbb{T}, \mathbb{K}}^+ \subseteq \mathcal{C}_{\mathbb{T}, \mathbb{K}}^-. \quad (10)$$

$$V^+(x, t) = \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \left\{ \min \left[\min_{\tau \in [t, t+\delta]} \max \left(l(\phi_{x,t}^{a, \beta[a]}(\tau), \tau), \max_{s \in [t, \tau]} g(\phi_{x,t}^{a, \beta[a]}(s), s) \right), \right. \right. \\ \left. \left. \max \left(V^+(\phi_x^{a, \beta[a]}(t+\delta), t+\delta), \max_{\tau \in [t, t+\delta]} g(\phi_{x,t}^{a, \beta[a]}(\tau), \tau) \right) \right] \right\}. \quad (11a)$$

$$V^-(x, t) = \inf_{\alpha \in \Gamma_t} \sup_{b \in \mathbb{B}_t} \left\{ \min \left[\min_{\tau \in [t, t+\delta]} \max \left(l(\phi_{x,t}^{a, \beta[a]}(\tau), \tau), \max_{s \in [t, \tau]} g(\phi_{x,t}^{a, \beta[a]}(s), s) \right), \right. \right. \\ \left. \left. \max \left(V^-(\phi_x^{a, \beta[a]}(t+\delta), t+\delta), \max_{\tau \in [t, t+\delta]} g(\phi_{x,t}^{a, \beta[a]}(\tau), \tau) \right) \right] \right\}. \quad (11b)$$

3. THE DOUBLE-OBSTACLE ISAACS EQUATION

It has been shown that the value function for minimum payoff games can be characterized as the unique viscosity solution to a variational inequality involving an appropriate Hamiltonian [2, 3], which has commonly been referred to as a Hamilton-Jacobi equation with an obstacle. We now extend the results for minimum cost problems to the category of problems with a more complex cost in the form of (4).

We first state the particular form of Bellman's principle of optimality [4] for the problem at hand.

LEMMA 2 (DYNAMIC PROGRAMMING PRINCIPLE). *Let $0 \leq t < T$ and $0 < \delta \leq T - t$. Then equation (11) holds.*

PROOF. The correctness of this lemma can be verified by inspection of (11), considering how the value in (4) is propagated back in time as per (7) along the *characteristic* (optimal trajectory) in all possible cases. The first term in the outer minimum of (11) is the local application of the definition in (4) restricted to the interval $[t, t + \delta]$,

$$V_{[t, t+\delta]} := \min_{\tau \in [t, t+\delta]} \max \left(l(\phi_{x,t}^{a, b}(\tau), \tau), \max_{s \in [t, \tau]} g(\phi_{x,t}^{a, b}(s), s) \right). \quad (12)$$

The minimum outcome achieved in the whole of $[t, T]$, however, will also be a function of the future value of (4) throughout the remainder of the game after $[t + \delta]$, captured by

$$V_{[t+\delta, T]} := V(\phi_{x,t}^{a, b}(t+\delta), t+\delta). \quad (13)$$

Now, if for all $\tau \in [t, t+\delta]$, $g(\phi_{x,t}^{a, b}(\tau), \tau) \leq V_{[t+\delta, T]}$, then from (4) it will clearly be that $V(x, t) = \min\{V_{[t, t+\delta]}, V_{[t+\delta, T]}\}$. The future value that is propagated along the characteristic, however, will be altered if anywhere on $[t, t + \delta]$, g exceeds $V_{[t+\delta, T]}$, in which case the maximum of g along the characteristic between t and $t + \delta$ will be propagated instead. Thus the second term in the outer minimum of (11) is

$$V_{t \leftarrow [t+\delta, T]} := \max \left(V_{[t+\delta, T]}, \max_{\tau \in [t, t+\delta]} g(\phi_{x,t}^{a, b}(\tau), \tau) \right). \quad (14)$$

The resulting value at (x, t) is therefore determined by the minimum of the local element and this last term, that is:

$$V^+(x, t) = \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \min \{ V_{[t, t+\delta]}^+, V_{t \leftarrow [t+\delta, T]}^+ \}, \quad (15a)$$

$$V^-(x, t) = \inf_{\alpha \in \Gamma_t} \sup_{b \in \mathbb{B}_t} \min \{ V_{[t, t+\delta]}^-, V_{t \leftarrow [t+\delta, T]}^- \}. \quad (15b)$$

The statement in (15) is a more compact form of (11). \square

We introduce the *upper* and *lower* Hamiltonians H^\pm :

$$H^+(x, p, t) = \min_{a \in \mathbb{A}} \max_{b \in \mathbb{B}} f(x, a, b, t) \cdot p, \quad (16a)$$

$$H^-(x, p, t) = \max_{b \in \mathbb{B}} \min_{a \in \mathbb{A}} f(x, a, b, t) \cdot p. \quad (16b)$$

The following theorem constitutes the main theoretical contribution of this paper; it shows that the value function V^\pm is the viscosity solution of a particular variational inequality that has the form of a Hamilton-Jacobi-Isaacs equation with a double obstacle.

THEOREM 1. *Assume f satisfies (1), and that $l(x, t), g(x, t)$ are globally Lipschitz continuous. Then the value function $V^\pm(x, t)$ for the game with outcome given by (4) is the unique viscosity solution of the variational inequality*

$$\max \left\{ \min \left\{ \partial_t V + H^\pm(x, D_x V^\pm, t), l(x, t) - V^\pm(x, t) \right\}, \right. \\ \left. g(x, t) - V^\pm(x, t) \right\} = 0, \quad t \in [0, T], x \in \mathbb{R}^n, \quad (17a)$$

with terminal condition

$$V^\pm(x, T) = \max \{ l(x, T), g(x, T) \}, \quad x \in \mathbb{R}^n. \quad (17b)$$

To prove this main result, we will make use of the following important continuity argument (stated and proven in [12] as Lemma 4.3).

LEMMA 3. *Let $\psi \in C^1(\mathbb{R}^n \times (0, T))$.*

(a) *If $\psi_t(x_0, t_0) + H^+(x_0, D_x \psi, t_0) \leq -\theta < 0$ then, for sufficiently small $\delta > 0$, there exists an input $a \in \mathbb{A}_{t_0}$ such that for all strategies $\beta \in \Lambda_{t_0}$,*

$$\int_{t_0}^{t_0+\delta} f(s, \phi_{x_0, t_0}^{a, \beta[a]}(s), a(s), \beta[a](s)) \cdot D_x \psi(\phi_{x_0, t_0}^{a, \beta[a]}(s), s) \\ + \psi_t(\phi_{x_0, t_0}^{a, \beta[a]}(s), s) ds \leq -\frac{\theta}{2} \delta.$$

(b) *If $\psi_t(x_0, t_0) + H^+(x_0, D_x \psi, t_0) \geq \theta > 0$ then, for sufficiently small $\delta > 0$, there exists a strategy $\beta \in \Lambda_{t_0}$ such that for all inputs $a \in \mathbb{A}_{t_0}$,*

$$\int_{t_0}^{t_0+\delta} f(s, \phi_{x_0, t_0}^{a, \beta[a]}(s), a(s), \beta[a](s)) \cdot D_x \psi(\phi_{x_0, t_0}^{a, \beta[a]}(s), s) \\ + \psi_t(\phi_{x_0, t_0}^{a, \beta[a]}(s), s) ds \geq \frac{\theta}{2} \delta.$$

PROOF OF THEOREM 1. The structure of the proof follows the classical approach in [12] and draws from viscosity solution theory. In every case, we start by assuming that V^\pm is not a viscosity solution of the HJI equation and derive a contradiction of the dynamic programming principle stated in Lemma 2. We will prove the theorem for V^+ with Hamiltonian H^+ ; the proof for V^- with H^- is analogous.

First, applying the definition of V^+ (4),(7) to the terminal case $t = T$, it is seen to satisfy the boundary condition (17b).

A continuous function is a viscosity solution of a partial differential equation if it is both a *subsolution* and a *supersolution* (defined below). We will first prove that V^+ is a viscosity subsolution of (17a). Let $\psi \in C^1(\mathbb{R}^n \times (0, T))$ such that $V^+ - \psi$ attains a local maximum at (x_0, t_0) ; without loss of generality, assume that this maximum is 0. We say that V^+ is a subsolution of (17a) if, for any such ψ ,

$$\max \left\{ \min \left\{ \partial_t \psi(x_0, t_0) + H^+(x_0, D_x \psi, t_0), l(x_0, t_0) - \psi(x_0, t_0) \right\}, g(x_0, t_0) - \psi(x_0, t_0) \right\} \geq 0. \quad (18)$$

Suppose (18) is false. Then it must be that

$$g(x_0, t_0) \leq \psi(x_0, t_0) - \theta_1, \quad (19)$$

and, in addition, at least one of the following holds:

$$l(x_0, t_0) \leq \psi(x_0, t_0) - \theta_2, \quad (20a)$$

$$\partial_t \psi(x_0, t_0) + H^+(x_0, D_x \psi, t_0) \leq -\theta_3, \quad (20b)$$

for some $\theta_1, \theta_2, \theta_3 > 0$. If (19) and (20a) are true, then by continuity of g, l and system trajectories there is a sufficiently small $\delta > 0$ such that for all $a(\cdot), b(\cdot), \tau \in [t_0, t_0 + \delta]$,

$$g(\phi(\tau), \tau) \leq \psi(x_0, t_0) - \frac{\theta_1}{2} = V^+(x_0, t_0) - \frac{\theta_1}{2},$$

$$l(\phi(\tau), \tau) \leq \psi(x_0, t_0) - \frac{\theta_2}{2} = V^+(x_0, t_0) - \frac{\theta_2}{2}.$$

For conciseness, we write $\phi_{x_0, t_0}^{a, b}(\cdot)$ as simply $\phi(\cdot)$ whenever statements hold for all inputs $a(\cdot), b(\cdot)$. Incorporating the above into the dynamic programming principle (11a) gives

$$\begin{aligned} V^+(x_0, t_0) &\leq \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \left\{ \min_{\tau \in [t_0, t_0 + \delta]} \max \left[l(\phi_{x_0, t_0}^{a, \beta[a]}(\tau), \tau), \right. \right. \\ &\quad \left. \left. \max_{s \in [t_0, \tau]} g(\phi_{x_0, t_0}^{a, \beta[a]}(s), s) \right] \right\} \\ &\leq V^+(x_0, t_0) - \min \left\{ \frac{\theta_1}{2}, \frac{\theta_2}{2} \right\}, \end{aligned}$$

which is a contradiction, since $\theta_1, \theta_2 > 0$.

If (19) and (20b) are true, then by Lemma 3, for small enough $\delta > 0$, there will exist some input $a \in \mathbb{A}(t_0)$ such that for all strategies $\beta \in \Lambda_{t_0}$,

$$\psi(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta) - \psi(x_0, t_0) \leq -\frac{\theta_3}{2} \delta,$$

and, recalling that $V^+ - \psi$ has a local maximum at (x_0, t_0) ,

$$V^+(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta) \leq V^+(x_0, t_0) - \frac{\theta_3}{2} \delta.$$

Inspecting (11a) in this case, we obtain

$$\begin{aligned} V^+(x_0, t_0) &\leq \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \left\{ \max \left[V^+(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta), \right. \right. \\ &\quad \left. \left. \max_{\tau \in [t_0, t_0 + \delta]} g(\phi_{x_0, t_0}^{a, \beta[a]}(\tau), \tau) \right] \right\} \\ &\leq V^+(x_0, t_0) - \min \left\{ \frac{\theta_1}{2}, \frac{\theta_3}{2} \delta \right\}, \end{aligned}$$

which again is a contradiction, since $\theta_1, \theta_3, \delta > 0$. Therefore, we conclude that (18) must be true and hence V^+ is indeed a subsolution of (17a).

We now proceed to show that V^+ is also a viscosity supersolution of (17a), that is, for all $\psi \in C^1(\mathbb{R}^n \times (0, T))$ such that $V^+ - \psi$ attains a local minimum at (x_0, t_0) (again, we can assume for convenience that this minimum is 0), it holds that

$$\max \left\{ \min \left\{ \partial_t \psi(x_0, t_0) + H^+(x_0, D_x \psi, t_0), l(x_0, t_0) - \psi(x_0, t_0) \right\}, g(x_0, t_0) - \psi(x_0, t_0) \right\} \leq 0. \quad (21)$$

If we suppose that (21) is false, then either it holds that

$$g(x_0, t_0) \geq \psi(x_0, t_0) + \theta_1, \quad (22)$$

or both of the following are true:

$$l(x_0, t_0) \geq \psi(x_0, t_0) + \theta_2, \quad (23a)$$

$$\partial_t \psi(x_0, t_0) + H^+(x_0, D_x \psi, t_0) \geq \theta_3, \quad (23b)$$

for some $\theta_1, \theta_2, \theta_3 > 0$.

If (22) holds, then there is a small enough $\delta > 0$ such that for all trajectories starting at (x_0, t_0) and all $t_0 \leq \tau \leq t_0 + \delta$

$$g(\phi(\tau), \tau) \geq \psi(x_0, t_0) + \frac{\theta_1}{2} = V^+(x_0, t_0) + \frac{\theta_1}{2}.$$

Then the dynamic programming principle (11a) yields

$$\begin{aligned} V^+(x_0, t_0) &\geq \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \left\{ \min \left[\min_{\tau \in [t_0, t_0 + \delta]} \max_{s \in [t_0, \tau]} g(\phi_{x_0, t_0}^{a, \beta[a]}(s), s), \right. \right. \\ &\quad \left. \left. \max_{\tau \in [t_0, t_0 + \delta]} g(\phi_{x_0, t_0}^{a, \beta[a]}(\tau), \tau) \right] \right\} \\ &\geq V^+(x_0, t_0) + \frac{\theta_1}{2}, \end{aligned}$$

which is a contradiction, as $\theta_1 > 0$.

If, on the other hand, (23) holds, then there is a small enough $\delta > 0$ such that, by (23a),

$$l(\phi(\tau), \tau) \leq \psi(x_0, t_0) + \frac{\theta_2}{2} = V^+(x_0, t_0) + \frac{\theta_2}{2},$$

and by (23b) and Lemma 3, there exists a strategy $\beta \in \Lambda_{t_0}$ such that for all inputs $a \in \mathbb{A}(t_0)$,

$$\begin{aligned} \frac{\theta_3}{2} \delta &\leq \psi(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta) - \psi(x_0, t_0) \\ &\leq V^+(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta) - V^+(x_0, t_0), \end{aligned}$$

by the local minimum condition. With this, (11a) gives

$$\begin{aligned} V^+(x_0, t_0) &\geq \sup_{\beta \in \Lambda_t} \inf_{a \in \mathbb{A}_t} \left\{ \min \left[\min_{\tau \in [t_0, t_0 + \delta]} l(\phi_{x_0, t_0}^{a, \beta[a]}(\tau), \tau), \right. \right. \\ &\quad \left. \left. V^+(\phi_{x_0, t_0}^{a, \beta[a]}(t_0 + \delta), t_0 + \delta) \right] \right\} \\ &\geq V^+(x_0, t_0) + \min \left\{ \frac{\theta_2}{2}, \frac{\theta_3}{2} \delta \right\}, \end{aligned}$$

resulting in another contradiction, as $\theta_2, \theta_3, \delta > 0$. We thus conclude that (21) holds and V^+ is a supersolution of (17a).

Since we have shown that V^+ is both a viscosity subsolution and a viscosity supersolution of the variational inequality, this completes the proof that V^+ is a viscosity solution of (17) with Hamiltonian H^+ . Uniqueness follows from the classical comparison and uniqueness theorems for viscosity solutions (see Theorem 4.2 in [3]). \square

4. NUMERICAL IMPLEMENTATION

We present in this section a numerical method to compute the value function (7) for the time-varying reach-avoid problem, based on the result in Theorem 1. For conciseness, we drop the distinction between upper and lower values and Hamiltonians, as the method is equally applicable to either.

Let $\mathbf{i} \in I$ denote the index of the grid point in a discretized computational domain of a compact subset $\mathcal{X} \subset \mathbb{R}^n$ and let $k \in \{1, \dots, n\}$ denote the index of each discrete time step in a finite interval $[0, T]$. Since our computation will proceed in backward time, we will let $T = t_0 > t_1 > \dots > t_n = 0$. To numerically solve the variational inequality (17), we use the following procedure, based on a three-step update rule:

Algorithm 1: Numerical Double-Obstacle HJI Solution

Data: $\hat{l}(x_i, t_k), \hat{g}(x_i, t_k)$
Result: $\hat{V}(x_i, t_k)$

Initialization
for $\mathbf{i} \in I$ **do**
Init $\hat{V}(x_i, t_0) \leftarrow \max\{\hat{l}(x_i, t_0), \hat{g}(x_i, t_0)\};$

Value propagation
for $k \leftarrow 1$ **to** n **do**
for $\mathbf{i} \in I$ **do**
U1 $\hat{V}(x_i, t_k) \leftarrow \hat{V}(x_i, t_{k-1})$
 $\quad + \int_{t_k}^{t_{k-1}} \hat{H}(x_i, D_x^+ \hat{V}(x_i, \tau), D_x^- \hat{V}(x_i, \tau), t_{k-1}) d\tau;$
U2 $\hat{V}(x_i, t_k) \leftarrow \min\{\hat{V}(x_i, t_k), l(x_i, t_k)\};$
U3 $\hat{V}(x_i, t_k) \leftarrow \max\{\hat{V}(x_i, t_k), g(x_i, t_k)\};$

The method uses discretized values of the payoff function $\hat{l}(x_i, t_k)$ and the discriminator function $\hat{g}(x_i, t_k)$; \hat{V} denotes the numerical approximation to V . The integral in the first update step (U1) is computed numerically using time derivative approximations. As an illustrative example, with a first order forward Euler scheme, we would have

$$\hat{V}(x_i, t_k) = \hat{V}(x_i, t_{k-1}) + (t_{k-1} - t_k) \hat{H}(x_i, D_x^+ \hat{V}(x_i, t_{k-1}), D_x^- \hat{V}(x_i, t_{k-1}), t_{k-1}). \quad (24)$$

The numerical scheme of Algorithm 1 is consistent with (17). $D_x^+ \hat{V}, D_x^- \hat{V}$ represent the ‘‘right’’ and ‘‘left’’ approximations of spatial derivatives. For the numerical Hamiltonian \hat{H} , we use the Lax-Friedrich approximation [15, 18]:

$$\hat{H}(x_i, D_x^+ \hat{V}, D_x^- \hat{V}, t_k) = H\left(x_i, \frac{D_x^- \hat{V} + D_x^+ \hat{V}}{2}, t_k\right) - \frac{1}{2} \alpha^\top (D_x^+ \hat{V} - D_x^- \hat{V}). \quad (25)$$

The components of α are given by $\alpha_i = \max_{p \in \mathcal{I}} \left| \frac{\partial H}{\partial p_i} \right|$, where \mathcal{I} is a hypercube containing all the values that p takes over the computational domain. With this choice of α for the Hamiltonian, the numerical scheme is stable [16, 18].

In the numerical examples in Section 5, we use a fifth-order accurate weighted essentially non-oscillatory scheme [17, 18] for the spatial derivatives $D_x^\pm \hat{V}$; for the time derivative $D_t \hat{V}$, we use a third-order accurate total variation diminishing Runge-Kutta scheme [17, 24]. These methods are

implemented by means of the computational tools provided in [13]. It should be noted that lower order spatial and time derivative approximations can also yield a numerically stable (although less accurate) solution to (17) at lower computational expense [15, 18].

It is important to stress the remarkable computational similarity of this new method to its time-invariant counterpart. Indeed, the only computational overhead is introduced by step (U3) in Algorithm 1, and the need to allow \hat{l} and \hat{g} to depend on time³. As a result, as will be demonstrated in the following section, our method can compute the backwards reachable set for time-varying problems at essentially no additional cost compared to the time-invariant case.

Lastly, the optimal action for each player is implicitly obtained in solving the minimax to compute the Hamiltonian \hat{H} in step (U1). It follows from Algorithm 1 that, starting inside a player’s winning region (the reach-avoid set for the attacker and its complement for the defender), applying this optimal action at each state as a feedback policy yields a guaranteed winning strategy for the reach-avoid game, to an arbitrary degree of precision determined by the discretization used.

5. NUMERICAL EXAMPLES

To illustrate our proposed method for computing reach-avoid sets, we present two numerical examples. The first shows the computational procedure in a simple optimal control scenario with a moving target and a moving obstacle, and the obtained capture basin is validated against the analytical result. The second example presents a two-player reach-avoid game with moving target and constraint sets; our method is benchmarked against the approach proposed in [5], producing the same computed set (within one grid cell of accuracy) at substantially lower computational cost.

5.1 Example 1: Reachability Problem

Consider the simple optimal control problem below, consisting of a vehicle that can move in any direction at some maximum speed trying to reach a moving target set while avoiding a moving obstacle, modeled as the complement of the constraint set. The time span of the problem is $[0, T]$ with $T = 0.5$. The system is described by the state vector $\mathbf{x}(t) = (p_x(t), p_y(t))$, which represents the vehicle’s position on a plane, with the following dynamics:

$$\dot{\mathbf{x}} = v_{\text{veh}} u(t), \quad u(t) \in \mathcal{U}, \quad (26)$$

with $v_{\text{veh}} = 0.5$ the vehicle’s speed and \mathcal{U} the unit disk.

The target set is a square with side length 0.4 moving down in forward time with velocity $v_{\text{tar}} = 1.5$. The center of the target set is at $(0, 0.75)$ at $t = 0$, and $(0, 0)$ at $t = 0.5$. The set is given mathematically as follows:

$$\mathcal{T}_t = \{(p_x, p_y) : \max(|p_x|, |p_y - 0.75 + v_{\text{tar}} t|) \leq 0.2\}. \quad (27)$$

We represent this moving target set using a signed distance function, $l(p_x, t) : l(p_x, t) \leq 0 \iff p_x \in \mathcal{T}_t$.

³Note that Algorithm 1 can be performed without storing the value of \hat{V} for all time steps, but only keeping track of the previous iteration. Just as in the finite-horizon time-invariant case, one may generally want to study the evolution of the reach-avoid set in time, but can discard intermediate iterations of \hat{V} if only the initial set is of interest.

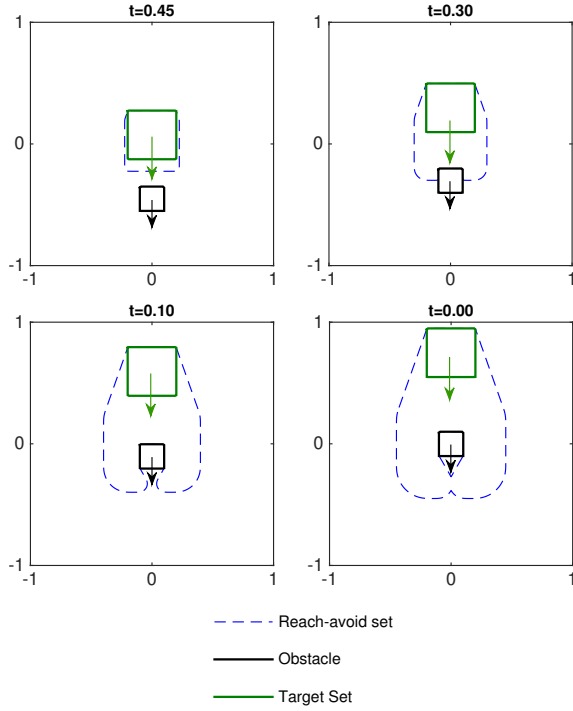


Figure 1: Time evolution of the reach-avoid set for a problem with a target (large square) moving down at speed 1.5, and an obstacle (small square) moving down at speed 1. The inside of the dashed boundary represents the set of states that can reach the target set while avoiding the obstacle.

The obstacle is a square with side length 0.2 moving down (in forward time) with velocity $v_{\text{obs}} = 1$. The center of the obstacle is at $(0, 0)$ at $t = 0$, and $(0, -0.5)$ at $t = 0.5$. The set is expressed as:

$$\mathcal{O}_t = \{(p_x, p_y) : \max(|p_x|, |p_y + v_{\text{obs}}t|) < 0.1\}. \quad (28)$$

We represent this moving obstacle by defining the constraint set $\mathcal{K}_t = \mathbb{R}^2 \setminus \mathcal{O}_t$, through the discriminator function $g(p_x, t)$: $g(p_x, t) \leq 0 \iff p_x \notin \mathcal{O}_t$.

Figure 1 shows the time evolution of the numerical solution for the example problem described above.

The obtained reach-avoid set is not unlike what one might expect out of intuition: because the target set is moving down at a speed greater than that of the vehicle, the lower boundary of the capture basin for $t = 0.45$ consists of states from which the vehicle can meet the target set at its final position. This lower boundary directly below the target moves down in backward time (as the vehicle has more time to get to this final position), but eventually gets “blocked” by the obstacle ($t = 0.3$). For earlier times ($t = 0.1$), the boundary is “pinched inwards” again, including nearby states from which the vehicle can move around the obstacle to get to the target; yet, there remains a triangular region directly below the obstacle, shown in the $t = 0$ subplot, that is not part of the reach-avoid set, because starting from those states the vehicle is unable to avoid the obstacle that is moving down. The diagonal boundaries of the capture basin at its upper region are formed by those states from which the vehicle can meet the target set between its initial and final positions.

5.1.1 Analytic Solution

The reach-avoid set boundary for this example problem can be computed analytically, and thereby compared against the numerically obtained boundary. Because the problem is symmetric about the p_y axis, we will consider the capture basin in the region $p_x \leq 0$. We now derive the analytic boundary by considering several different segments separately; it will be convenient to refer to Figure 2 below.

The optimal path for a vehicle with initial position on segment 1 of the capture basin boundary is a straight trajectory, perpendicular to the segment, that reaches the upper corner of the moving target at some intermediate position. Segment 1 is continued by a short arc 2 comprising initial states from which the vehicle can follow a straight path reaching this top corner exactly at the target’s final position. For a vehicle starting on segments 3, 4 and 5 the optimal action is to take the shortest path to the closest point of the target’s final position, which will be reached at exactly the final time. The optimal action for a vehicle with initial position on segment 6 is similar: it must follow a straight line to barely miss the obstacle, before redirecting its path to the target, reaching it at the final time. Finally, a vehicle initially within the triangular region enclosed by segment 7 and the obstacle cannot avoid being hit by the obstacle. Based on these considerations, the expression for each of these segments can be geometrically derived, leading to the following:

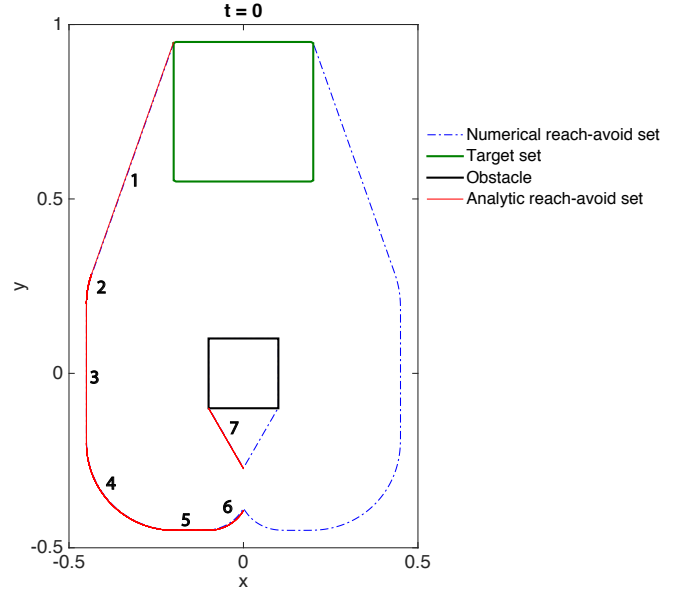


Figure 2: Analytic and numeric reach-avoid set.

1. Upper diagonal segment:

$$\{(p_x, p_y) : p_y = m(p_x + 0.2) + 0.95, p_x \in [p_x^*, -0.2]\}, \quad (29)$$

where

$$p_x^* = \frac{-v_{\text{tar}}T}{m^{-1} + m} - 0.2, \quad m = \sqrt{\frac{v_{\text{tar}}^2 - v_{\text{veh}}^2}{v_{\text{veh}}^2}}. \quad (30)$$

$$p_y^* = \frac{-v_{\text{tar}}T}{m^{-2} + 1} + 0.95,$$

2. Upper transition arc

$$\begin{aligned} \{(p_x, p_y) : (p_x + 0.2)^2 + (p_y - 0.2)^2 &= (v_{\text{veh}}T)^2, \\ p_x \in [-0.45, p_x^*], p_y \in [0.2, p_y^*]\}, \end{aligned} \quad (31)$$

3. Side straight segment:

$$\{(p_x, p_y) : p_y \in [-0.2, 0.2], p_x = -0.45\} \quad (32)$$

4. Outer bottom rounded corner:

$$\begin{aligned} \{(p_x, p_y) : p_x = \cos(\theta) - 0.2, p_y = \sin(\theta) - 0.2, \\ \theta \in [\pi, 3\pi/2]\} \end{aligned} \quad (33)$$

5. Bottom straight segment:

$$\{(p_x, p_y) : p_x \in [-0.2, -0.1], p_y = -0.45\} \quad (34)$$

6. Bottom rounded corner under obstacle:

$$\{(p_x, p_y) : p_y = y_{om} - \sqrt{d^2 - (p_x - x_{om})^2}, p_x \in [x_{om}, 0]\} \quad (35)$$

where (d, y_{om}) solves

$$\begin{aligned} d + y_{tf} - y_{om} &= C \\ \frac{d}{v_{\text{veh}}} &= \frac{y_{oi} - y_{om}}{v_{\text{obs}}} \end{aligned} \quad (36)$$

7. Obstacle's "shadow":

$$\{(p_x, p_y) : p_y = mp_x + 0.95, p_x \in [-0.1, 0]\}, \quad (37)$$

where

$$m = \sqrt{\frac{v_{\text{obs}}^2 - v_{\text{veh}}^2}{v_{\text{veh}}^2}}. \quad (38)$$

5.1.2 Convergence

Using the scheme described in Section 4, we numerically solved the double-obstacle Hamilton-Jacobi variational inequality (17) on a computation domain consisting of $N \times N$ grid points for $N = 51, 101, 151, 201, 251, 301$. We compared each of the numerical solutions to the analytic solution derived in Section 5.1.1 by the following procedure:

1. Construct signed distance functions with the zero level set corresponding to the boundary of the numerically computed reach-avoid set (for instance using [13]).
2. Evaluate the signed distance functions at approximately 20 000 points distributed on the analytically determined boundary of the reach-avoid set.

The values of the signed distance function correspond to the distance between the analytically computed reach-avoid set boundary points to the numerically computed boundary. These values are used as the error metric for the numerical approximation.

Figure 3 shows, in logarithmic scale, the mean error and maximum error over all analytic points plotted against the size of spatial discretization or grid spacing. Consistently across the different grid spacings, the mean error is approximately one-tenth of the grid spacing, and the maximum error is approximately half of the grid spacing. The numerical scheme therefore converges both in terms of the mean error and the maximum error.




Figure 3: Convergence of our numerical implementation for Example 1 with different grids. Average error is consistently an order of magnitude smaller than the grid spacing, with the maximum error being roughly half of the grid size.

5.2 Example 2: Reach-Avoid Game

Consider a reach-avoid game in which the attacker moves in a two-dimensional space while the defender moves on the vertical line $x = 0.05$. Let $p_A = (x_A, y_A)$ be the position of the attacker, and y_D be the position of the defender, with $\mathbf{x} = (x_A, y_A, y_D)$ the state of the system, governed by the following time-varying dynamics:

$$\begin{aligned} \dot{p}_A &= (1-t)v_A a(t), \quad \|a\|_2 \leq 1, \\ \dot{y}_D &= v_D b(t), \quad b \in [-1, 1]. \end{aligned} \quad (39)$$

In this reach-avoid game, the attacker wishes to reach a target set that is moving upwards at speed $v_T = 1.5$, while the defender tries to prevent the attacker from succeeding by intercepting or delaying its advance. The attacker is free to move in any direction at a limited speed, anywhere in a square domain with the exception of a growing obstacle whose lower edge is expanding downwards at a speed of $v_O = 0.5$. The attacker has a time-varying maximum speed that decreases linearly from $v_A = 3$ at $t = 0$ to 0 at $t = 1$. The defender has a time-invariant maximum speed of $v_D = 3$. Here, interception is defined as the two players being within a radius of $R = 0.1$ of each other. Figure 4 shows the initial configuration of the moving target and the moving obstacle, as well the interception set centered at four different defender positions.

For this reach-avoid game, we seek to compute the reach-avoid set, comprised by the set of joint positions from which the attacker is guaranteed to be able to reach the target while avoiding interception by the defender as well as collision with the obstacle. To compute the reach-avoid set, we solve (17) with the following Hamiltonian:

$$H(\mathbf{x}, \nabla_{\mathbf{x}} V, t) = \min_{a \in \mathcal{A}} \max_{b \in [-1, 1]} \nabla_{p_A} V (1-t)v_A a(t) + \nabla_{y_D} V v_D b(t). \quad (40)$$

Solving the minimax in the Hamiltonian, we get

$$H(\mathbf{x}, \nabla_{\mathbf{x}} V, t) = -(1-t)v_A \|\nabla_{p_A} V\|_2 + v_D |\nabla_{y_D} V|. \quad (41)$$

Since the state space of the reach-avoid game is three-dimensional ($n = 3$), we visualize two-dimensional cross sec-

tions of the three-dimensional reach-avoid set at $t = 0$, taken at various defender initial positions. Figure 4 compares the two-dimensional slices of the reach-avoid sets computed by the state augmentation method in [5] and by our newly proposed augmentation-free method. Given the defender positions shown in each of the subplots, the attacker will be able to reach the target if it is on the side of the reach-avoid set boundary containing the target. As can be appreciated, the capture basin boundaries computed by the two methods are very similar (well within a grid cell of distance); however, computation⁴ using the state augmentation method took approximately 3 hours on a 51^4 grid. With our proposed augmentation-free method, computation only took about 3 minutes on a 51^3 grid. Our computation was two orders of magnitude faster and provided essentially identical results.

At this point it should be mentioned that there exist local level set methods developed to speed up computation in time-invariant problems with monotonically propagating level sets [19, 23]; these methods update the value function only in a neighborhood of its zero level set, which can partially mitigate the cost of iterating over an $(n + 1)$ -dimensional grid. We did not implement such fast methods here, since the goal was to compare the accuracy and performance of the full implementation of both approaches. It is important to stress that our method inherently avoids the increase in computational cost altogether, because it uses an n -dimensional grid, and in addition returns the entire value function and not just its zero level set, which can provide useful additional information [1].

The effect of the different defender initial positions on the reach-avoid set is apparent in Figure 4. If the defender starts the game near the bottom of the domain (a), the defender would be able to block the attacker from going through the gap between the bottom edge of the obstacle and that of the domain. Thus we see that the reach-avoid set boundary does not extend into the bottom left quadrant of the domain. However, in this case, the attacker is free to cross the gap above the top edge of the obstacle, which leads to a large area of the top left quadrant being inside the capture basin.

Similarly, if the defender starts near the top (d), the reach-avoid set extends into the bottom left quadrant of the domain. Yet it does not propagate as far as in (a), due to the fact that the passage under the obstacle is closing and the target is moving away from it: an attacker not starting close enough to the bottom opening will not be able to make it through in time to reach the target. The remaining plots (b),(c) show the capture basin at $t = 0$ for intermediate defender positions.

Figure 5 shows the backward time evolution of the reach-avoid set for a single defender position. The subplots show the capture basin at various times. At $t = 0.90$, there is a relatively small region in the state space from which the attacker can reach the target by the end of the game ($t = 1$), as there is little time left. As the starting time t considered decreases, the attacker has more time to reach the target and thus the reach-avoid set grows; however, this growth is inhibited by both the defender’s interception set and by the presence of the obstacle. Furthermore, near $t = 1$, the attacker has a slow speed, so the growth of the capture basin depends primarily on the motion of the target set; as t decreases, the attacker’s motion becomes more relevant.

⁴Computations were run using [13] on desktop computer with a Core i7-2600K processor.

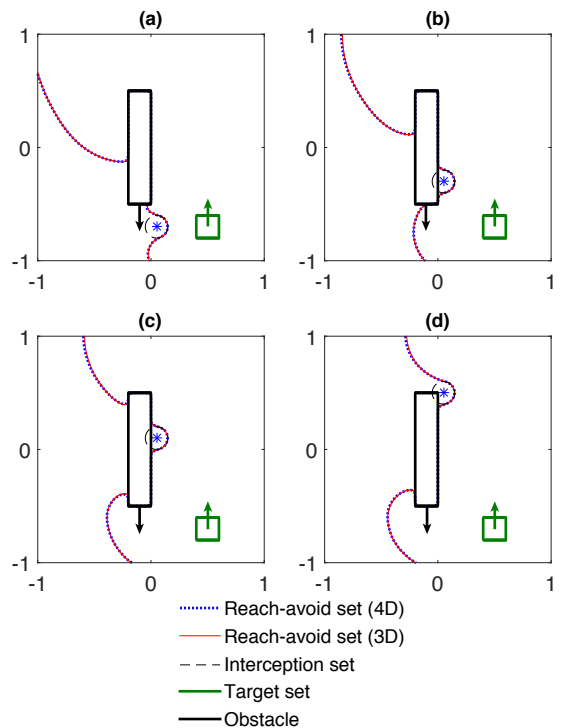


Figure 4: Reach-avoid set computed through the state augmentation method (4D) and our proposed augmentation-free method (3D). 2D cross-sections of the set are shown at the initial time for four different defender positions.

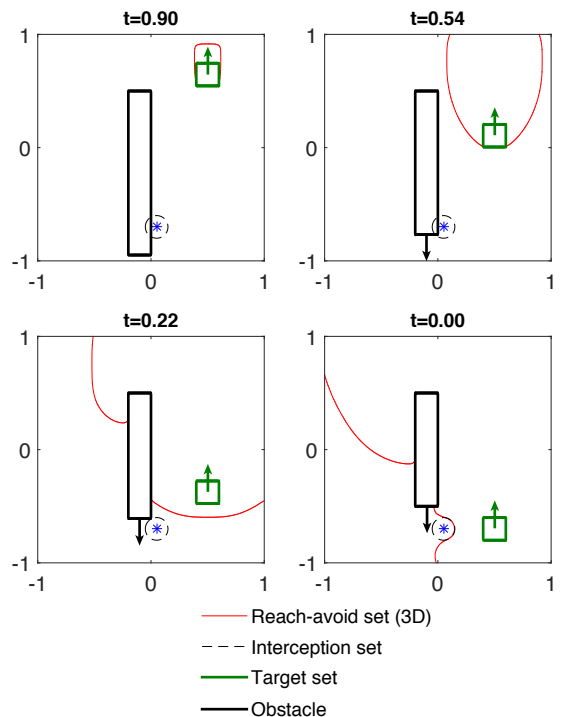


Figure 5: Backward time evolution of the reach-avoid set. As t decreases, the attacker has more time to reach the target, so the reach-avoid set grows. The growth of the reach-avoid set is inhibited by the defender’s interception set and the obstacle.

6. CONCLUSION

We have presented here a novel extension of Hamilton-Jacobi methods to reach-avoid problems with time-varying dynamics, targets, and constraints. This result enables the analysis of many relevant problems in game theory and optimal control, including pursuit-evasion, differential games, and safety certificates for dynamical systems. In particular, our result can provide guarantees for collision avoidance in dynamic environments with multiple moving obstacles.

Importantly, numerical implementations of our method have computational complexity equivalent to that of already existing techniques for time-invariant systems. This sets our method apart from previously proposed approaches that work around time variation by incorporating time as an additional variable in the state. In many important application contexts, such as online safety analysis in dynamical systems [1], the substantial reduction in computational cost introduced by our technique can allow timely obtention of results that would otherwise entail an impractical computational effort.

In the future, we intend to develop applications of this new formulation to large-scale multi-agent systems in both cooperative and adversarial contexts: a first result is presented in [8] for safe multi-vehicle path planning. By leveraging the possibilities of time-varying targets and constraints to encode the trajectories of other agents, we aim to incrementally build solutions that scale linearly, and not exponentially, with the complexity of the multi-agent network.

References

- [1] A. K. Akametalu, J. F. Fisac, et al. “Reachability-Based Safe Learning with Gaussian Processes”. *Proceedings of the 53rd IEEE Conference on Decision and Control* (2014).
- [2] E. Barron. “Differential Games with Maximum Cost”. *Nonlinear analysis: Theory, methods & applications* (1990), pp. 971–989.
- [3] E. Barron and H. Ishii. “The Bellman equation for minimizing the maximum cost”. *Nonlinear Analysis: Theory, Methods & Applications* (1989).
- [4] R. Bellman. *Dynamic Programming*. 1st ed. Princeton, NJ, USA: Princeton University Press, 1957.
- [5] O. Bokanowski and H. Zidani. “Minimal time problems with moving targets and obstacles”. *18th IFAC World Congress* (2011).
- [6] O. Bokanowski, N. Forcadel, and H. Zidani. “Reachability and Minimal Times for State Constrained Nonlinear Problems without Any Controllability Assumption”. *SIAM Journal on Control and Optimization* 48.7 (2010), pp. 4292–4316.
- [7] P. Cardaliaguet. “A double obstacle problem arising in differential game theory”. *Journal of Mathematical Analysis and Applications* 360.1 (2009), pp. 95–107.
- [8] M. Chen, J. F. Fisac, S. Sastry, and C. J. Tomlin. “Safe Sequential Path Planning of Multi-Vehicle Systems via Double-Obstacle Hamilton-Jacobi-Isaacs Variational Inequality”. *Proceedings of the 14th European Control Conference (to appear)* (2015).
- [9] E. A. Coddington and N. Levinson. *Theory of ordinary differential equations*. Tata McGraw-Hill Education, 1955.
- [10] A. Cosso. “Stochastic Differential Games Involving Impulse Controls and Double-Obstacle Quasi-Variational Inequalities”. *SIAM Journal on Control and Optimization* 51.3 (2013), pp. 2102–2131.
- [11] R. J. Elliott and N. J. Kalton. *The existence of value in differential games*. Vol. 126. American Mathematical Soc., 1972.
- [12] L. C. Evans and P. E. Souganidis. “Differential games and representation formulas for solutions of Hamilton-Jacobi-Isaacs equations”. *Indiana University mathematics journal* 33.5 (1984), pp. 773–797.
- [13] I. M. Mitchell. *A toolbox of level set methods*. 2004. URL: <http://www.cs.ubc.ca/~mitchell/ToolboxLS>.
- [14] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. *IEEE Transactions on Automatic Control* 50.7 (2005), pp. 947–957.
- [15] I. Mitchell. “Application of Level Set Methods to Control and Reachability Problems in Continuous and Hybrid Systems”. PhD thesis. Stanford University, 2002.
- [16] I. M. Mitchell. “Application of Level Set Methods to Control and Reachability Problems in Continuous and Hybrid Systems”. PhD thesis. Stanford University, 2002.
- [17] S. Osher and R. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer Verlag, 2003.
- [18] S. Osher and C.-W. Shu. “High-Order Essentially Nonoscillatory Schemes for Hamilton-Jacobi Equations”. *SIAM Journal on Numerical Analysis* 28.4 (1991), pp. 907–922.
- [19] B. Peng, B. Merriman, et al. “A PDE-based fast local level set method”. *Journal of computational physics* 155 (1999), pp. 410–438.
- [20] M. Quincampoix and O.-S. Serea. “A viability approach for optimal control with infimum cost”. *Annals. Stiint. Univ. Al. I. Cuza Iasi, sI a, Mat* 1 (2002), pp. 1–20.
- [21] A. Rapaport. “Characterization of Barriers of Differential Games”. *Journal of optimization theory and applications* 97.1 (1998), pp. 151–179.
- [22] E. Roxin. “Axiomatic approach in differential games”. *Journal of Optimization Theory and Applications* 3.3 (1969), pp. 153–163.
- [23] J. A. Sethian. “A fast marching level set method for monotonically advancing fronts”. *Proceedings of the National Academy of Sciences* 93.4 (1996), pp. 1591–1595.
- [24] C.-W. Shu and S. Osher. “Efficient implementation of essentially non-oscillatory shock-capturing schemes”. *Journal of Computational Physics* 77.2 (1988), pp. 439–471.
- [25] P. Varaiya. “On the existence of solutions to a differential game”. *SIAM Journal on Control* 5.1 (1967), pp. 153–162.