Certified Approximate Reachability (CARe): Formal Error Bounds on Deep Learning of Reachable Sets

Prashant Solanki[*1], Nikolaus Vertovec[*2], Yannik Schnitzer[2],
Jasper Van Beers[1], Coen de Visser[1], Alessandro Abate[2]

Abstract— **Recent approaches to leveraging deep learning for computing reachable sets of continuous-time dynamical systems have gained popularity over traditional level-set methods, as they overcome the curse of dimensionality. However, as with level-set methods, considerable care needs to be taken in limiting approximation errors, particularly since no guarantees are provided during training on the accuracy of the learned reachable set. To address this limitation, we introduce an ϵ-approximate Hamilton-Jacobi Partial Differential Equation (HJ-PDE), which establishes a relationship between training loss and accuracy of the *true* reachable set. To formally certify this approximation, we leverage Satisfiability Modulo Theories (SMT) solvers to bound the residual error of the HJ-based loss function across the domain of interest. Leveraging Counter Example Guided Inductive Synthesis (CEGIS), we close the loop around learning and verification, by fine-tuning the neural network on counterexamples found by the SMT solver, thus improving the accuracy of the learned reachable set. To the best of our knowledge, Certified Approximate Reachability (CARe) is the first approach to provide soundness guarantees on learned reachable sets of continuous dynamical systems.**

I. INTRODUCTION

Ensuring the safety of autonomous systems under uncertainty is a fundamental challenge in control and verification. Reachability analysis plays a key role in addressing safety concerns across various domains, including spacecraft trajectory design [35], [37], ground transportation systems [21], [23], air traffic management [20], [32], and flight control [22], [24], [34]. The theoretical foundations of reachability stem from viability theory [3], [4], while computational techniques have been developed for both exact and approximate reachable set computations in hybrid systems [2], [7], [8], [27], [28]. These advances have formulated reachability analysis as an optimal control problem, where reachable, viable, or invariant sets are represented as level sets of a value function satisfying a Hamilton-Jacobi (HJ) partial differential equation (PDE) [5], [9], [10], [22], [25].

Traditionally, level set methods are used to compute the unique viscosity solution of the HJ equation [28]. However, even with high-order numerical schemes such as WENO methods, the non-differentiability of the value function introduces numerical inaccuracies [36]. While refining the grid used in finite element methods such as level set methods improves numerical accuracy, it is hard to predict what

*Authors contributed equally to this article.
[1]Department of Aerospace Engineering, TU Delft, Netherlands {p.solanki, j.j.vanbeers, c.c.devisser}@tudelft.nl
[2]Department of Computer Science, University of Oxford, OX1 3PJ, UK {nikolaus.vertovec, yannik.schnitzer, alessandro.abate}@cs.ox.ac.uk

constitutes a sufficiently fine grid. Moreover, these grid-based approaches suffer from the curse of dimensionality, as the number of required grid points grows exponentially with system dimensionality.

Deepreach [6], a recent approach leveraging neural networks, has enabled solving the HJ equation without relying on finite difference methods. If the loss function enforcing the PDE conditions converges to zero in the region of interest, the learned value function is the unique viscosity solution of the HJ-PDE, ensuring that its zero-level set correctly represents the reachable set. This approach improves scalability to higher-dimensional systems by mitigating the curse of dimensionality.

In this work, we demonstrate an additional advantage of using neural networks to approximate the value function: the ability to provide a formally bounded ϵ-accurate reachable set. While both level-set methods and DeepReach yield empirically accurate solutions, their adherence to the HJ-PDE conditions–and thus the validity of the reachable set– has not been quantified. Follow-up work to DeepReach uses probabilistic methods to recover the reachable set [18], [19], [31]. In contrast, we provide a formal approach to verification that does not rely on recovery of the reachable set, but rather provides a sound over- and under-approximation of the reachable set. Our key contributions are:

1) We establish a formal bound ϵ on the residual error, ensuring that the absolute value of the HJ-based loss function remains within a specified threshold across the domain of interest.
2) We demonstrate that the neural network-based value function, constrained by ϵ, provides a certified approximation of the true value function, allowing for over- and under-approximation of the reachable set.

The remainder of the paper is structured as follows: Section II provides an overview of the reachability problems and introduces the Hamilton-Jacobi (HJ) partial differential equation (PDE) whose solution we aim to learn. Section III presents the deep learning approach for computing the reachable set, while Section IV discusses the formal verification method used to derive and certify an upper bound on the loss across the entire domain of interest. In Section V, we introduce our main contribution, establishing a connection between the residual training loss and an under- and over-approximation of the reachable set. Finally, Section VI provides implementation details and a case study.

II. Problem Setup

In reachability theory, a key objective is to compute the backward reachable set (BRS) of a dynamical system, which consists of all states from which trajectories can reach a given target set at the end of the time horizon. In contrast, the backward reachable tube (BRT) represents the set of states that can reach the target set within a specified time horizon [5]. If the target set represents unsafe states, the BRS/BRT identifies states that may lead to unsafe conditions and should be avoided. To account for adversarial disturbances, safety-critical scenarios are often modeled as a two-player game, where Player 1 represents the control input and Player 2 represents the disturbance input [26].

Mathematically, consider a dynamical system with state $x \in \mathbb{R}^m$ governed by the ordinary differential equation (ODE)

$$\dot{x}(s) \in f(s, x(s), u(s), d(s)), \quad t_0 \leq s \leq T, \qquad (1)$$

where $T \geq t_0$ is the fixed time horizon. The initial state is given by $x(t_0) := x_0$, and the control and disturbance inputs are measurable functions

$$u : [t_0, T] \to \mathcal{U}, \quad d : [t_0, T] \to \mathcal{D},$$

with $\mathcal{U} \subset \mathbb{R}^k$ and $\mathcal{D} \subset \mathbb{R}^l$ being compact sets. The dynamics $f : [t_0, T] \times \mathbb{R}^m \times \mathcal{U} \times \mathcal{D} \to \mathbb{R}^m$ satisfy the following conditions for some constants C_1, C_2:

$$|f(t, x, u, d)| \leq C_1, \qquad (2)$$

$$|f(t, x, u, d) - f(t, \hat{x}, u, d)| \leq C_2 |x - \hat{x}|, \qquad (3)$$

for all $t \in [t_0, T]$, $x, \hat{x} \in \mathbb{R}^m$, $u \in \mathcal{U}$, and $d \in \mathcal{D}$. These assumptions ensure that the ODE (1) admits a unique solution

$$x(t) = \phi(t, t_0, x_0, u(\cdot), d(\cdot)). \qquad (4)$$

Next, we define the sets of control and disturbance policies:

$$\mathcal{M}_{[t,T]} \equiv \{u : [t, T] \to \mathcal{U} \mid u \text{ is measurable}\},$$

$$\mathcal{N}_{[t,T]} \equiv \{d : [t, T] \to \mathcal{D} \mid d \text{ is measurable}\}.$$

A nonanticipative strategy is a mapping $\beta : \mathcal{M}_{[t,T]} \to \mathcal{N}_{[t,T]}$ such that for all $s \in [t, T]$ and for all $u, \hat{u} \in \mathcal{M}_{[t,T]}$, if $u(\tau) = \hat{u}(\tau)$ for almost every $\tau \in [t, s]$, then $\beta[u](\tau) = \beta[\hat{u}](\tau)$ for almost every $\tau \in [t, s]$. The class of such strategies is denoted as $\Delta_{[t,T]}$.

Given a target set $\mathcal{G}$ with a signed distance function $g(x)$ such that $g(x) \leq 0 \iff x \in \mathcal{G}$, the BRT is defined as

$$\text{BRT}_{\mathcal{G}}([t_0, T]) = \{x \mid \exists \beta \in \mathcal{N}_{[t_0,T]}, \forall u \in \mathcal{M}_{[t_0,T]}, \qquad (5)$$
$$\exists s \in [t_0, T] : \phi(s, t_0, x_0, u(\cdot), \beta(\cdot)) \in \mathcal{G}\}.$$

Equivalently, the BRT is given by the subzero level set of the value function $V^*(t, x)$, where

$$V^*(t, x) = \inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))). \qquad (6)$$

Let us introduce the shorthand notation $D_x V(t, x) = \frac{\partial V(t,x)}{\partial x}$ and $D_t V(t, x) = \frac{\partial V(t,x)}{\partial t}$. Then, a standard result of HJ

reachability [5], [12], [26] allows us to reduce the optimization problem in Equation (6) from an infinite-dimensional optimization over policies to an optimization over control and disturbance inputs by formulating the value function as the viscosity solution of the HJ-PDE:

$$D_t V^*(t, x) + \min\{0, \mathcal{H}(t, x, D_x V^*)\} = 0,$$
$$V^*(T, x) = g(x), \qquad (7)$$

with the Hamiltonian defined as

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p].$$

The HJ-PDE can be solved using level-set methods [22], [27], [33] or through curriculum training of a neural network [6].

III. Deep Learning Approach to HJ Reachability

Inspired by the self-supervised framework of Physics-Informed Neural Networks (PINNs) [29], the solution of the PDE (7) can be effectively learned, as demonstrated in [6].

Let V_θ denote the approximate value function learned by a neural network. To ensure that the zero level set of V_θ correctly represents the BRT, the network must achieve zero loss for the loss function defined as

$$\mathcal{L}(t_i, x_i; \theta) = \mathbb{1}(t_i = T)\mathcal{L}_1(x_i; \theta) + \lambda \mathcal{L}_2(t_i, x_i; \theta), \qquad (8)$$

$$\mathcal{L}_1(x_i; \theta) = \|V_\theta(T, x_i) - g(x_i)\|, \qquad (9)$$

$$\mathcal{L}_2(t_i, x_i; \theta) = \|D_t V_\theta(t_i, x_i) + \min[0, \mathcal{H}(t_i, x_i, D_x V_\theta)]\|, \qquad (10)$$

where $\mathbb{1}(t_i = T)$ is an indicator function ensuring that $\mathcal{L}_1$ is only applied at $t_i = T$, and λ controls the balance between the two loss terms. The nonlinear nature of the neural network might result in areas of the state space with significantly higher loss than empirically observed during training. This acute shortcoming necessitates a formal approach to verifying the approximation errors. To this end, in the subsequent sections, we will

1) establish a verification method to ensure $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x and t within the domain of interest; and
2) derive upper and lower bounds on V_θ in terms of viscosity solutions to the HJ-PDE, thus enabling under- and over-approximations of the BRT/BRS.

IV. Formal Synthesis of Neural Value Functions

Our approach to synthesizing the value function solution follows a two-phase process–a learning phase and a certification phase–that alternates in a Counter Example Guided Inductive Synthesis (CEGIS) loop [1]. The learning phase employs a curriculum training approach, discussed in detail in Section VI, terminating when we empirically satisfy $|\mathcal{L}_1(x_i; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t_i, x_i; \theta)| < \epsilon_2$, for N randomly sampled points $(t_1, x_1), \ldots, (t_n, x_n)$.

To further ensure that $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$ for all x in the domain of interest $\mathcal{X}$ and for all $t \in [t_0, T]$, we design a sound certification phase: this certification is performed via SMT (Satisfiability Modulo Theories) solving,

which symbolically reasons over the continuous domain $[t_0, T] \times \mathcal{X}$, hence generalizing across the sample-based loss error. While computationally intensive, formal verification provides soundness guarantees, allowing us to rigorously validate the learned value function as an approximate solution to the HJ-PDE.

Several approaches exist for certifying the accuracy of the learned value function. To remain general in our choice of dynamical systems, we employ an SMT solver capable of handling quantifier-free nonlinear real arithmetic formulae [13]. This allows us to incorporate arbitrary nonlinear activation functions into the neural network. Furthermore, it enables us to address systems with nonlinear dynamics, which induce nonlinearities in the Hamiltonian. Specifically, we employ `dReal` [15], which supports both polynomial and non-polynomial terms, such as transcendental functions (e.g., trigonometric and exponential functions).

After generating a symbolic representation of the neural network, the SMT solver searches for an assignment of variables (t, x) that satisfies the quantifier-free formula

$$\left(x \in \mathcal{X} \wedge \|V_\theta(T, x) - g(x)\| > \epsilon_1 \right) \vee \tag{11}$$

$$\left(x \in \mathcal{X} \wedge t \in [t_0, T] \wedge \right.$$
$$\left. \|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| > \epsilon_2 \right) \tag{12}$$

The logical disjunction between (11) and (12) allows us to split the SMT call into separate parallel queries. If one query identifies a valid assignment, i.e., a counterexample to either $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ or $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, the remaining SMT calls terminate. A counterexample indicates the necessity to finetune the neural network, as it does not yet approximate a valid value function with sufficient accuracy. We leverage the counterexample for targeted training specifically at inputs where the network is insufficiently accurate. We sample additional training data points for finetuning in close proximity to the counterexample, thus reducing the need for extensive sampling in regions of the input space that are already well approximated. Finetuning and certification are alternately repeated within an inductive synthesis loop (CEGIS), progressively improving the neural value function in targeted regions of the input space, until the certifier determines the Formulae (11) and (12) to be unsatisfiable. This implies that there exists no counterexample to $|\mathcal{L}_1(x; \theta)| < \epsilon_1$ and $|\mathcal{L}_2(t, x; \theta)| < \epsilon_2$, formally proving the validity of the learned value function over the entire domain of interest. In the following section, we derive results establishing how this formal certification can be leveraged to obtain sound over- and under-approximations of the BRT and BRS from the neural value function.

V. ϵ-ACCURATE VALUE FUNCTIONS

The certification of the neural value function in the previous section established that for all $(t, x) \in [t_0, T] \times \mathcal{X}$, there exists some $\hat{\epsilon} \in [0, \epsilon_2]$ such that

$$\|D_t V_\theta(t, x) + \min[0, \mathcal{H}(t, x, D_x V_\theta)]\| = \hat{\epsilon},$$

which can equivalently be rewritten as

$$\left\| D_t V_\theta(t, x) + \min[\pm\hat{\epsilon}, \tilde{\mathcal{H}}(t, x, D_x V_\theta)] \right\| = \hat{\epsilon},$$
$$\tilde{\mathcal{H}}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p \pm \hat{\epsilon}].$$

We consider the worst-case realisation of $\hat{\epsilon}$, i.e., $\hat{\epsilon} = \epsilon_2$. We will now show that, for $\epsilon = (\epsilon_1 + \epsilon_2(T - t))$,

$$V_\theta(t, x) - \epsilon \leq V^*(t, x) \leq V_\theta(t, x) + \epsilon, \tag{13}$$

which allows us to use the neural value function to reason about the true BRT and BRS.

A. Bounding the True Value Function

To derive these bounds, we introduce the modified value functions $\overline{V}$ and $\underline{V}$, which are the unique viscosity solutions under the worst-case realizations of $\hat{\epsilon}$:

$$\overline{V}(t, x) = \inf_{\beta(\cdot) \in \Delta_{[t,T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,T]}}$$
$$\inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon_2 \, ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right], \tag{14}$$

$$\underline{V}(t, x) = \inf_{\beta(\cdot) \in \Delta_{[t,T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,T]}}$$
$$\inf_{\tau \in [t,T]} \left[\int_t^\tau -\epsilon_2 \, ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right]. \tag{15}$$

Since ϵ_2 is a positive constant, it follows that

$$\underline{V}(t, x) \leq V^*(t, x) \leq \overline{V}(t, x). \tag{16}$$

Furthermore, using the inequality $-\epsilon_2(T - t) \leq \int_t^\tau -\epsilon_2 \, ds$, we obtain:

$$\inf_{\tau \in [t,T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \epsilon_2(T - t)$$
$$\leq \inf_{\tau \in [t,T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) - \int_t^\tau \epsilon_2 \, ds \right].$$

Similarly, using $\epsilon_2(T - t) \geq \int_t^\tau \epsilon_2 \, ds$, we obtain:

$$\inf_{\tau \in [t,T]} g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \epsilon_2(T - t)$$
$$\geq \inf_{\tau \in [t,T]} \left[g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) + \int_t^\tau \epsilon_2 \, ds \right].$$

Thus, we derive the key bound:

$$\overline{V}(t, x) - \epsilon_2(T - t) \leq V^*(t, x) \leq \underline{V}(t, x) + \epsilon_2(T - t). \tag{17}$$

B. Relating Bounds to the Neural Value Function

To relate $\overline{V}$ and $\underline{V}$ to the neural value function, we need to establish that these auxiliary value functions are the unique viscosity solutions of the ϵ-modified HJ-PDE. This result is formalized in the following theorem.

Theorem 1: The function $\overline{V}(t, x)$ is the unique viscosity solution of the ϵ-modified HJ-PDE:

$$D_t \overline{V}(t, x) + \min \left\{ \epsilon_2, \mathcal{H}(t, x, D_x \overline{V}) \right\} = 0,$$

with the Hamiltonian defined as

$$\mathcal{H}(t,x,p) = \max_{u\in\mathcal{U}}\min_{d\in\mathcal{D}}\left[f(t,x,u,d)\cdot p + \epsilon_2\right],$$

and the terminal condition

$$\overline{V}(T,x) = g(x).$$

Similarly, $\underline{V}(t,x)$ is the unique viscosity solution of the ϵ-modified HJ-PDE:

$$D_t\underline{V}(t,x) + \min\left\{-\epsilon_2, \mathcal{H}(t,x,D_x\underline{V})\right\} = 0,$$

with the Hamiltonian

$$\mathcal{H}(t,x,p) = \max_{u\in\mathcal{U}}\min_{d\in\mathcal{D}}\left[f(t,x,u,d)\cdot p - \epsilon_2\right],$$

and the terminal condition

$$\underline{V}(T,x) = g(x).$$

The proof of Theorem 1 is given in the appendix.

As a consequence of Theorem 1, we obtain

$$\underline{V}(t,x) \le V_\theta(t,x) \le \overline{V}(t,x). \tag{18}$$

Finally, incorporating ϵ_1, which perturbs the boundary condition, we have:

$$V_\theta(T,x) = g(x) \pm \epsilon_1. \tag{19}$$

Thus, by combining Equation (17) with Equations (18) and (19), we recover the bound in Equation (13).

VI. CASE STUDIES

Given the page limitations, we prioritize detailing the key steps of the training and verification loop to ensure clarity, and thus omit a full comparison of different reach/avoid examples. The provided toolbox[1] includes additional examples, such as an evader-pursuer scenario, forward and backward reachability problems, and implementations of both the reachable tube and the reachable set.

Let us consider the simple double integrator example, a canonical second-order control system, $\dot{x}_1 = x_2, \dot{x}_2 = u$, where x_1 represents position, x_2 represents velocity, and u is the control input. We consider the problem of learning the forward reachable set, with the initial set defined by

$$g(x) = x_1^2 + x_2^2 - R^2, \tag{20}$$

with $R^2 = 0.5$.

A. Training of $V_\theta(t,x)$

The training procedure, outlined in Algorithm 1 and illustrated in Fig. 1 follows a curriculum learning strategy with three phases:

- **Pretraining phase:** Focuses on boundary conditions at the final time $t = T$.
- **Curriculum phase:** Gradually extends the time horizon from $[T,T]$ to $[t_0,T]$.
- **Finetuning phase:** Further refines the model to minimize loss before certification.

Each training epoch involves generating a random batch of samples and computing the loss $\mathcal{L}(t_i, x_i, \theta)$ for each sample $i \in \{1, \ldots, N\}$. The formal error bounds on the reachable set can be conservative as the results in Theorem 1 assume the worst-case realisation of ϵ_2 along any trajectory. Consequently, even when the mean training loss is low, the existence of regions of the state space with high network loss can lead to conservative over- or under-approximations of the reachable set. To mitigate this, we not only minimize the mean error over the batch but also the maximum error, leading to the training objective of minimizing:

$$\frac{1}{N}\sum_i \mathcal{L}(t_i, x_i, \theta) + \lambda_{\max}\max_i \mathcal{L}(t_i, x_i, \theta).$$

Here, $\lambda_{\max}$ is set to 10% during the curriculum phase and 30% during the finetuning phase. The finetuning phase uses a patience counter, terminating only when the loss consistently stays below a threshold $\lambda_\epsilon\epsilon_2$, with $\lambda_\epsilon = 95\%$. This modified approach to the curriculum training scheme presented in [6] reduces the total training time while ensuring sufficient confidence in the model prior to progressing to the certification phase.

Algorithm 1 Training of $V_\theta(t,x)$

1: **Input:** Thresholds ϵ_1 and ϵ_2
2: **Initialize:** Model V_θ, $t_{\text{current}} \leftarrow T$
3: **while** Training **do**
4:      Generate samples x_i and $t_i \in [t_{\text{current}}, T]$
5:      Compute model output $\hat{y} = V_\theta(t_i, x_i)$
6:      Compute loss $\mathcal{L}_{total} = \mathcal{L}_{mean} + \lambda_{max} \cdot \mathcal{L}_{max}$
7:      Backpropagate $\mathcal{L}_{total}$ and update parameters θ
8:      **if** $\mathcal{P}_{\text{pre-training}}$ **and** $\max(\mathcal{L}_1) < \epsilon_1$ **then**
9:          Progress to $\mathcal{P}_{\text{curriculum}}$
10:     **else if** $\mathcal{P}_{\text{curriculum}}$ **and** $\max(\mathcal{L}_2) < \epsilon_2$ **then**
11:         Expand time horizon: $t_{\text{current}} \leftarrow t_{\text{current}} - \Delta_t$
12:         **if** $t_{\text{current}} = 0$ **then**
13:            Progress to $\mathcal{P}_{\text{finetune}}$
14:            $p \leftarrow 0$
15:         **end if**
16:     **else if** $\mathcal{P}_{\text{finetune}}$ **and** $\max(\mathcal{L}_2) < \lambda_\epsilon\epsilon_2$ **then**
17:         $p \leftarrow p + 1$
18:         **if** $p > 1000$ **then**
19:            **break**   ▷ Stop training
20:         **end if**
21:     **end if**
22: **end while**
23: **Output:** Trained model V_θ

To improve performance and enable smaller networks to approximate the value function effectively, we introduce a polynomial layer as the first layer of the network. This allows subsequent layers to operate not only on the original inputs (t, x) but also on polynomial transformations, such as (t^2, x^2).
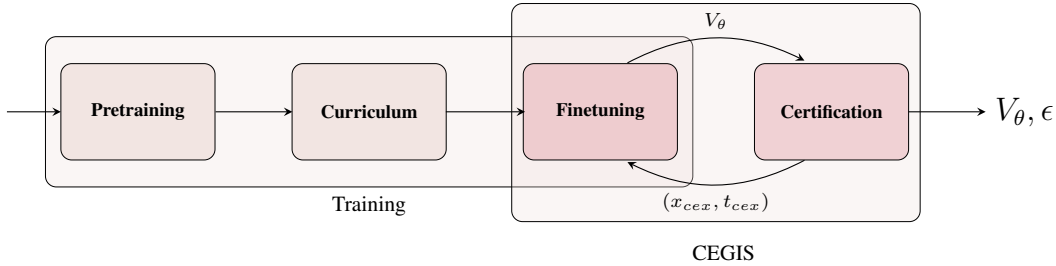
Fig. 1. Flowchart of the training and verification Procedure
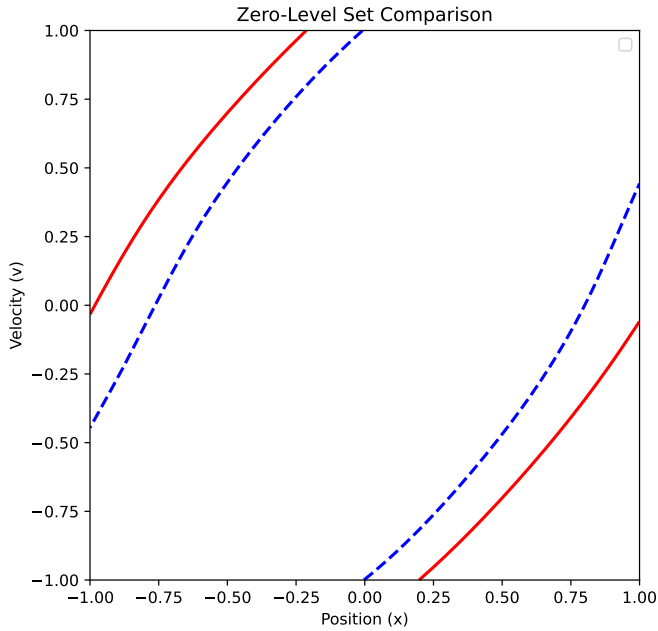

Zero-Level Set Comparison

Fig. 2. The over-approximated reachable set is shown in red, while the learned reachable set is depicted in blue. The initial set at time $t_0 = 0$ is a circle around the centre of radius 0.25, the time horizon is $T = 1$.

B. Formal Verification

To formally verify that the trained network represents an ϵ-accurate value function, we generate a symbolic representation $V_\theta(t, x)$ to instantiate the quantifier-free SMT query described in Section IV. To obtain the corresponding derivative terms $D_t V_\theta(t, x)$ and $\mathcal{H}(t, x, D_x V_\theta)$ we apply symbolic differentiation. To further parallelize the SMT call, we decompose the absolute value expressions of Equations (11)–(12) into separate calls, e.g., we seek x s.t.

$$x \in \mathcal{X} \wedge V_\theta(T, x) - g(x) > \epsilon_1,$$
$$\vee \ x \in \mathcal{X} \wedge V_\theta(T, x) - g(x) < -\epsilon_1.$$

As SMT with nonlinear real arithmetic is undecidable in general, no exact decision procedure can exist for arbitrary HJ-PDEs [17]. However, `dreal` provides a δ-complete decision procedure that allows for user-specified δ deviations in the satisfying assignments [14], [16]. Consequently, while satisfying assignments may be spurious, unsatisfiability carries over to the exact problem. Hence, if the SMT query is unsatisfiable, it formally establishes a global upper bound on the loss function across the entire domain of interest.

TABLE I

TRAINING AND VERIFICATION RESULTS OF THE CEGIS LOOP

Iter.	ϵ	Training (s)	Verif. (s)	Result
1	0.30	836.92	1367	Certified
2	0.27	-	22	Counterexample found
3	0.27	181.27	2735	Certified
4	0.243	-	5537	Certified
5	0.2187	-	237	Counterexample found

Checking SMT queries with nonlinear real arithmetic can be computationally expensive, especially as unsatisfiability constitutes a full proof for validity. In practice, if a counterexample exists, the SMT query typically terminates within a reasonable timeframe.

For illustration, we consider training and verifying the double integrator system using a single hidden layer with 16 neurons, employing sine activation functions. After training with $\epsilon := (\epsilon_1 + \epsilon_2(T - t_0))$, and $T = 1, t_0 = 0, \epsilon_2 = 0.95\epsilon = 0.285, \epsilon_1 = 0.05\epsilon = 0.015$ we progress to certification, which takes 1367 seconds resulting in no counterexample being found. Subsequently, we attempt to reduce ϵ by 10%. Within 22 seconds, one of the parallelized SMT queries finds a counterexample. Thus, at this stage, our trained neural network satisfies $\mathcal{L}(t, x) < 0.3$ but not $\mathcal{L}(t, x) < 0.27$. Following a Counterexample-Guided Inductive Synthesis (CEGIS) approach, we leverage the identified counterexample to refine our model (Fig. 1). Specifically, we re-enter the finetuning phase of training with the reduced ϵ this time ensuring that 10% of our samples are drawn near the counterexample. This refinement improves the model such that a subsequent verification attempt fails to find a new counterexample, thereby establishing a tighter bound on the loss. We continue the CEGIS loop for a total of 5 iterations with the results shown in Table I. The fine-tuning phase of the final iteration is unable to obtain an empirical loss below 0.2187, thus, the best certifiable result with $\epsilon = 0.243$ is used to produce the final over-approximated reachable set, shown in Fig. 2.

Our approach is applicable to arbitrary continuous-time reachability problems. However, the use of `dreal` practically limits the network size, as increasing the depth and number of neurons adds complexity to the formulas (11)-(12). If an empirical result is satisfactory, Theorem 1 can provide an over- or under-approximation of the reachable set, even when only an empirical bound on the loss is available.

VII. CONCLUSIONS

The primary goal of this paper is to establish a relation between the training loss and the induced level of approximation error on the final BRS/BRT. To this end, we introduce a modified HJ-PDE allowing us to relate a bound on the loss function used during training, to the true solutions of the reachability problem. To obtain the bound on the loss function, we could use an empirical data-driven approach to estimating the bound on the training loss, though this would fail to constitute a formal certificate. Thus we introduce an approach to verification using a Satisfiability Modulo Theory solver, enabling us to provide a formal certificate for approximate reachability.

APPENDIX

For the proofs in the remaining sections, we ease notation by omitting the subscript of ϵ, i.e. $\epsilon_2 = \epsilon$. we introduce the payoff function:

$$
\begin{aligned}
P(u, d) &= P_{x,t}(u(\cdot), \beta(\cdot)) \\
&= \int_t^T \epsilon \, ds + g(\phi(T, t, x, u(\cdot), \beta(\cdot))),
\end{aligned} \quad (21)
$$

where $g : \mathbb{R}^m \to \mathbb{R}$ satisfies the following conditions:

$$
|g(x)| \le C_2, \quad (22)
$$

$$
|g(x) - g(\hat{x})| \le C_2 |x - \hat{x}|, \quad (23)
$$

for some constant C_2 and for all $t_0 \le t \le T$, $x, \hat{x} \in \mathbb{R}^m$, $u \in \mathcal{U}$, and $d \in \mathcal{D}$. In the differential game setup considered in this paper, the control u is chosen to maximize $P(u, d)$, while the disturbance d is chosen to minimize $P(u, d)$. The term ϵ represents a positive perturbation ($\epsilon \in \mathbb{R}^+$). Since the proofs for $\underline{V}$ follow directly from those of $\overline{V}$, we focus only on the case of $\overline{V}$ and ease notation by using V to denote $\overline{V}$. We ease notation be introducing $\mathbf{x}_{u,\beta}(\tau) := \phi(\tau, t, x, u(\cdot), \beta(\cdot))$.

A. Bellman Optimality

Theorem 2: The value function defined in Equation (14) satisfies the Bellman optimality principle:

$$
\begin{aligned}
\overline{V}(t, x) = \inf_{\beta(\cdot) \in \Delta_{[t,t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \Bigg\{ \\
\overline{V}(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\
\inf_{\tau \in [t,t+\sigma]} \Big[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \Big] \Bigg\}.
\end{aligned}
$$

Proof: We define

$$
\begin{aligned}
W(t, x) = \inf_{\beta(\cdot) \in \Delta_{[t,t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \Bigg\{ \\
V(t + \sigma, \mathbf{x}_{u,\beta}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\
\inf_{\tau \in [t,t+\sigma]} \Big[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \Big] \Bigg\}.
\end{aligned}
$$

Part I: Lower bound of $W(t, x)$:

Fix $\gamma > 0$, and choose $\beta_1 \in \Delta_{[t,t+\sigma]}$ such that:

$$
\begin{aligned}
W(t, x) \ge \sup_{u_1(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \Bigg\{ \\
V(t + \sigma, \mathbf{x}_{u_1,\beta_1}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\
\inf_{\tau \in [t,t+\sigma]} \Big[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \Big] \Bigg\} - \gamma.
\end{aligned}
$$

Thus, for all $u_1(\cdot) \in \mathcal{M}_{[t,t+\sigma]}$, there exists a $\beta_1 \in \Delta_{[t,t+\sigma]}$ such that:

$$
\begin{aligned}
W(t, x) \ge \min \Bigg\{ V(t + \sigma, \mathbf{x}_{u_1,\beta_1}(t + \sigma)) + \int_t^{t+\sigma} \epsilon \, ds, \\
\inf_{\tau \in [t,t+\sigma]} \Big[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \Big] \Bigg\} - \gamma. \quad (24)
\end{aligned}
$$

Now, consider the term $V(t + \sigma, x)$:

$$
\begin{aligned}
V(t + \sigma, x) = \inf_{\beta(\cdot) \in \Delta_{[t+\sigma,T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma,T]}} \inf_{\tau \in [t+\sigma,T]} \Bigg[\int_{t+\sigma}^\tau \epsilon \, ds \\
+ g(\phi(\tau, t + \sigma, x, u(\cdot), \beta(\cdot))) \Bigg].
\end{aligned}
$$

Similarly, we can find a $\beta_2(\cdot) \in \Delta_{[t+\sigma,T]}$ and $u_2(\cdot) \in \mathcal{M}_{[t+\sigma,T]}$ such that:

$$
\begin{aligned}
V(t + \sigma, x) \ge \inf_{\tau \in [t+\sigma,T]} \Bigg[\int_{t+\sigma}^\tau \epsilon \, ds \\
+ g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \Bigg] - \gamma.
\end{aligned}
$$

Now, define the combined control policy $u(\cdot) \in \mathcal{M}_{[t,T]}$ and disturbance strategy $\beta(\cdot) \in \Delta_{[t,T]}$ as:

$$
u(\tau) = \begin{cases} u_1(\tau) & \text{if } t \le \tau < t + \sigma, \\ u_2(\tau) & \text{if } t + \sigma \le \tau \le T. \end{cases}
$$

$$
\beta[u](\tau) = \begin{cases} \beta_1[u_1](\tau) & \text{if } t \le \tau < t + \sigma, \\ \beta_2[u_2](\tau) & \text{if } t + \sigma \le \tau \le T. \end{cases}
$$

Subsequently, we obtain:

$$W(t,x) \geq \min \left\{ \int_t^{t+\sigma} \epsilon \, ds + \right.$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right],$$

$$\left. \inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u_1(\cdot), \beta_1(\cdot))) \right] \right\} - 2\gamma,$$

$$\geq \min \Big\{$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right]$$

$$\left. \inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u_1(\cdot), \beta_1(\cdot))) \right] \right\} - 2\gamma,$$

$$= \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t, x, u(\cdot), \beta(\cdot))) \right] - 2\gamma.$$

This holds true for all $u(\cdot) \in \mathcal{M}_{[t,T]}$, and hence we conclude:

$$W(t,x) \geq \sup_{u(\cdot) \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] - 2\gamma.$$

Finally, this gives:

$$W(t,x) + 2\gamma \geq V(t,x).$$

Part II: Upper Bound $W(t,x)$
For all $\beta_1(\cdot) \in \Delta_{[t,t+\sigma]}$, we have:

$$W(t,x) \leq \sup_{u(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \Big\{$$

$$V(t+\sigma, \mathbf{x}_{u,\beta_1}(t+\sigma)) + \int_t^{t+\sigma} \epsilon \, ds,$$

$$\inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta_1}(\tau)) \right] \Big\}.$$

Then for a fixed $\gamma > 0$, there exists a $u_1 \in \mathcal{M}_{[t,t+\sigma]}$ such that:

$$W(t,x) \leq \min \Big\{ V(t+\sigma, \mathbf{x}_{u_1,\beta_1}(t+\sigma)) + \int_t^{t+\sigma} \epsilon \, ds,$$

$$\inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \right] \Big\} + \gamma.$$

Consider the term

$$V(t+\sigma, \mathbf{x}_{u,\beta}(t+\sigma)) = \inf_{\beta \in \Delta_{[t+\sigma,T]}} \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma,T]}}$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u(\cdot), \beta(\cdot))) \right].$$

Then for all $\beta_2 \in \Delta_{[t+\sigma,T]}$, we have:

$$V(t+\sigma, x) \leq \sup_{u(\cdot) \in \mathcal{M}_{[t+\sigma,T]}}$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u(\cdot), \beta_2(\cdot))) \right].$$

For a fixed $\gamma > 0$, there exists a $u_2 \in \mathcal{M}_{[t+\sigma,T]}$ such that:

$$V(t+\sigma, x) \leq \inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds \right.$$

$$\left. + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] + \gamma.$$

We define the combined control policy $u(\cdot) \in \mathcal{M}_{[t,T]}$ and disturbance strategy $\beta(\cdot) \in \Delta_{[t,T]}$ as before, such that:

$$W(t,x) \leq \min \Big\{$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right],$$

$$\inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \right] \Big\} + 2\gamma.$$

$$\leq \min \Big\{$$

$$\inf_{\tau \in [t+\sigma,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\phi(\tau, t+\sigma, x, u_2(\cdot), \beta_2(\cdot))) \right],$$

$$\inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta_1}(\tau)) \right] \Big\} + 2\gamma.$$

$$\leq \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] + 2\gamma. \tag{25}$$

It follows from Equation (14) that:

$$V(t,x) = \inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right].$$

Thus, there exists $\beta \in \Delta_{[t,T]}$ such that for a given $\gamma > 0$:

$$V(t,x) \geq \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] - \gamma.$$

Thus for all $u \in \mathcal{M}_{[t,T]}$, we get:

$$V(t,x) + \gamma \geq \inf_{\tau \in [t,T]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right]. \tag{26}$$

Combining this with Equation (25), we get:

$$W(t,x) \leq V(t,x) + 3\gamma. \tag{27}$$

Finally, since γ is arbitrary, we can let $\gamma \to 0$ and obtain:

$$V(t,x) = \inf_{\beta(\cdot) \in \Delta_{[t,t+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t,t+\sigma]}} \min \Big\{$$

$$V(t+\sigma, \mathbf{x}_{u,\beta}(t+\sigma)) + \int_t^{t+\sigma} \epsilon \, ds,$$

$$\inf_{\tau \in [t,t+\sigma]} \left[\int_t^{\tau} \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right] \Big\}.$$

■

The proof of Theorem 1 relies on several further properties of the value function, which will be provided in the form of Lemmas 1–4. As in the previous section, we will use V for ease of notation. We begin with the monotonicity of the value function.

B. Monotonicity of the value function

Lemma 1: For all $(t, x) \in [0, T] \times \mathbb{R}^m$

$$V(t, x) \leq V(t + \sigma, x) + \int_t^{t+\sigma} \epsilon \, ds$$

and $V(T, x) = g(x)$

Proof: $V(T, x) = g(x)$ is trivial and can be directly observed from the definition of the value function, Equation (14). Consider

$$V(t, x) = \inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right]$$

Let us assume for the sake of contradiction that

$$V(t, x) > V(t + \sigma, x) + \int_t^{t+\sigma} \epsilon \, ds,$$

which can be equivalently stated as

$$\inf_{\beta_1 \in \Delta_{[t,T]}} \sup_{u_1 \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \right]$$

$$> \inf_{\beta_2 \in \Delta_{[t+\sigma,T]}} \sup_{u_2 \in \mathcal{M}_{[t+\sigma,T]}} \inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds \right.$$

$$\left. + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] + \int_t^{t+\sigma} \epsilon \, ds.$$

Thus, for all $\beta_1 \in \Delta_{[t,T]}$, there exists a $\beta_2 \in \Delta_{[t+\sigma,T]}$ such that

$$\sup_{u_1 \in \mathcal{M}_{[t,T]}} \inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \right]$$

$$> \sup_{u_2 \in \mathcal{M}_{[t+\sigma,T]}} \inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds \right.$$

$$\left. + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right] + \int_t^{t+\sigma} \epsilon \, ds.$$

Furthermore, for all $u_2 \in \mathcal{M}_{[t+\sigma,T]}$, there exists $u_1 \in \mathcal{M}_{[t,T]}$ such that

$$\inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u_1,\beta_1}(\tau)) \right]$$

$$> \inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u_2(\cdot), \beta_2(\cdot))) \right]$$

$$+ \int_t^{t+\sigma} \epsilon \, ds.$$

Now let us define:

$$u(\tau) \equiv \begin{cases} u_1(\tau) & \text{if } t \leq \tau < t + \sigma, \\ u_1(\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

and

$$\beta[u](\tau) \equiv \begin{cases} \beta_1[u](\tau) & \text{if } t \leq \tau < t + \sigma, \\ \beta_2[u](\tau) & \text{if } t + \sigma \leq \tau \leq T. \end{cases}$$

Using the uniqueness of the solution of the ODE (1), it follows that

$$\inf_{\tau \in [t,T]} \left[\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right]$$

$$> \inf_{\tau \in [t+\sigma,T]} \left[\int_{t+\sigma}^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u(\cdot), \beta(\cdot))) \right]$$

$$> \inf_{\tau \in [t+\sigma,T]} \left[\int_t^\tau \epsilon \, ds + g(\phi(\tau, t + \sigma, x, u(\cdot), \beta(\cdot))) \right],$$

which is a contradiction, as the infimum over a larger set is always less than or equal to the infimum over a smaller set. ∎

As an immediate corollary of Lemma 1, we have

Corollary 1: For all $(t, x) \in [0, T] \times \mathbb{R}^m$

$$V(\tau, x(\tau)) + \inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \int_t^\tau \epsilon \, ds \leq$$

$$\inf_{\beta \in \Delta_{[t,T]}} \sup_{u \in \mathcal{M}_{[t,T]}} \left(\int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau)) \right)$$

C. Existence and Uniqueness

We provide 2 lemmas that show that the value function is bounded and Lipschitz continuous. This is required for the existence and uniqueness guarantees.

Lemma 2: The value function, defined in the Equation (14), is bounded i.e. $|V(t, x)| \leq C_4$ Where C_4 is a constant

Proof:

$$P(u(\cdot), \beta(\cdot)) = \int_t^T \epsilon \, ds + g(\phi(T, t, x, u(\cdot), \beta(\cdot)))$$

Using Equation (22) and the fact that ϵ is a constant, it follows that $|P(u(\cdot), \beta(\cdot))| \leq (T - t)\epsilon + C_2$. This holds for all $u(\cdot) \in \mathcal{M}_{[t,T]}$ and $\beta(\cdot) \in \Delta_{[t,T]}$. Thus implying $|V(t, x)| \leq C_4$. ∎

Lemma 3: The value function, defined in Equation (14), is Lipschitz continuous i.e. $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$ for all $0 \leq t \leq T$, $0 \leq \hat{t} \leq T$ and $x, \hat{x} \in \mathbb{R}^m$.

Proof: We introduce the following notation

$$P_{t,x}(\tau, u, \beta[u]) := \int_t^\tau \epsilon \, ds + g(\mathbf{x}_{u,\beta}(\tau))$$

Part I

We will show that $V(t_1, x_1) - V(t_2, x_2) \leq \overline{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma$, where $\overline{C}$ is some constant. Without loss of generality we can assume that $t_1 \leq t_2$. Furthermore, for any β_1

$$V(t_1, x_1) = \inf_{\beta \in \Delta_{[t_1,T]}} \sup_{u \in \mathcal{M}_{[t_1,T]}} \inf_{\tau \in [t_1,T]} P_{t_1,x_1}(\tau, u, \beta[u])$$

$$\leq \sup_{u \in \mathcal{M}_{[t_1,T]}} \inf_{\tau \in [t_1,T]} P_{t_1,x_1}(\tau, u, \beta_1[u])$$

Subsequently, for a fixed $\gamma > 0$, $\exists u_1 \in \mathcal{M}_{[t_1,T]}$ such that for all $\tau \in [t_1, T]$ we have

$$V(t_1, x_1) \leq P_{t_1,x_1}(\tau, u_1, \beta_1[u_1]) + \gamma \tag{28}$$

Similarly, we can show that for a fixed $\gamma > 0$, $\exists \beta_2 \in \Delta_{[t_2,T]}$ such that

$$V(t_2, x_2) \geq \sup_{u \in \mathcal{M}_{[t_2,T]}} \inf_{\tau \in [t_2,T]} P_{t_2,x_2}(\tau, u, \beta_2[u]) - \gamma \quad (29)$$

For an arbitrary $u_1 \in \mathcal{M}_{[t_2,T]}$ we have that

$$V(t_2, x_2) \geq \inf_{\tau \in [t_1,T]} P_{t_2,x_2}(\tau, u_1, \tilde{\beta}[u_1]) - \gamma$$

where $\tilde{\beta}[u_1]$ is defined as

$$\tilde{\beta}[u_1](\tau) \equiv \begin{cases} \beta_1[u_1](\tau) & \text{if} \quad t_1 \leq \tau < t_2 \\ \beta_2[u_1](\tau) & \text{if} \quad t_2 \leq \tau \leq T \end{cases}$$

Using the definition of infimum $\exists \tau \in [t_1, T]$ such that

$$V(t_2, x_2) \geq P_{t_2,x_2}(\tau, u_1, \tilde{\beta}[u_1]) - 2\gamma$$
$$\Rightarrow -V(t_2, x_2) \leq -P_{t_2,x_2}(\tau, u_1, \tilde{\beta}[u_1]) + 2\gamma \quad (30)$$

Now let $x_1(\cdot)$ be the solution for the time horizon ($t_1 \leq s \leq T$) of the following ODE

$$\begin{cases} \frac{dx_1}{ds} = f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) \\ x_1(t_1) = x_1 \end{cases}$$

Let $x_2(\cdot)$ be the solution for the time horizon ($t_2 \leq s \leq T$) of the ODE

$$\begin{cases} \frac{dx_2}{ds} = f(s, x_2(s), u_1(s), \beta_2[u_1](s)) \\ x_2(t_2) = x_2 \end{cases}$$

Then it follows from equation (28) and (30), that

$$V(t_1, x_1) - V(t_2, x_2) \leq$$
$$P_{t_1,x_1}(\tau, u_1, \beta_1[u_1]) - P_{t_2,x_2}(\tau, u_1, \tilde{\beta}[u_1]) + 3\gamma$$
$$= \int_{t_1}^{\tau} \epsilon ds + g(x_1(\tau)) - \int_{t_2}^{\tau} \epsilon ds - g(x_2(\tau)) + 3\gamma$$
$$= g(x_1(\tau)) - g(x_2(\tau)) + 3\gamma + (t_2 - t_1)\epsilon \quad (31)$$

Using Equation (23), it follows that

$$\left| g(x_1(\tau)) - g(x_2(\tau)) \right| \leq C_2 \left| x_1(\tau) - x_2(\tau) \right| \quad (32)$$

Furthermore, since $\tilde{\beta}[u_1](s) = \beta_2[u_1](s)$ for all $s \in$

$[t_2, T]$, we have that

$$|x_1(\tau) - x_2(\tau)| = \left| x_1 + \int_{t_1}^{\tau} f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) ds \right.$$
$$\left. - x_2 - \int_{t_2}^{\tau} f(s, x_2(s), u_1(s), \beta_2[u_1](s)) ds \right|$$
$$= \left| x_1 + \int_{t_1}^{t_2} f(s, x_1(s), u_1(s), \tilde{\beta}[u_1](s)) ds - x_2 \right.$$
$$+ \int_{t_2}^{\tau} f(s, x_1(s), u_1(s), \beta_2[u_1](s))$$
$$\left. - f(s, x_2(s), u_1(s), \beta_2[u_1](s)) ds \right|$$
$$\leq \left| x_1(t_2) - x_2(t_2) \right|$$
$$+ \int_{t_2}^{\tau} \left| f(s, x_1(s), u_1(s), \beta_2[u_1](s)) \right.$$
$$\left. - f(s, x_2(s), u_1(s), \beta_2[u_1](s)) \right| ds$$
$$\leq \left| x_1(t_2) - x_2(t_2) \right| + C_1 \int_{t_2}^{\tau} |x_1(s) - x_2(s)| ds$$
$$\leq \left| x_1(t_2) - x_2(t_2) \right| e^{(\tau - t_2)C_1}$$

where the second inequality uses Equations (2) and (3) and the last inequality is due to the Bellman-Gronwall Lemma [30].

Using Equation (2), the following holds for any $u \in \mathcal{U}, d \in \mathcal{D}$

$$|x_1(t_2) - x_1| \leq C_3|t_1 - t_2|, \quad (33)$$

$$|x_1(t_2) - x_1(t_1)| = \left| \int_t^{t_2} f(s, x_2(s), u(s), d(s)) \, ds \right.$$
$$\left. - \int_t^{t_1} f(s, x(s), u(s), d(s)) \, ds \right|$$
$$\leq |C_1(t_2 - t) - C_1(t_1 - t)|$$
$$\leq C_1|t_1 - t_2|.$$

Thus we obtain

$$\left| g(x_1(\tau)) - g(x_2(\tau)) \right| \leq C_2 \left| x_1(\tau) - x_2(\tau) \right|$$
$$\leq C_2 e^{(\tau - t_2)C_1} |x_1(t_2) - x_2(t_2)|$$
$$\leq C_2 e^{(\tau - t_2)C_1} |x_1(t_2) - x_1|$$
$$+ C_2 e^{(\tau - t_2)C_1} |x_1 - x_2|$$
$$\leq C_3 e^{(\tau - t_2)C_1} |t_1 - t_2| + e^{(\tau - t_2)C_1} |x_1 - x_2|$$

Subsequently, there exists some constant $\overline{C}$, such that

$$V(t_1, x_1) - V(t_2, x_2) \leq \overline{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma \quad (34)$$

Part II
Now we will show that $V(t_2, x_2) - V(t_1, x_1) \leq \overline{C}(|t_1 - t_2| + |x_1 - x_2|) + 3\gamma$. We have that

$$V(t_2, x_2) = \inf_{\beta \in \Delta_{[t_2,T]}} \sup_{u \in \mathcal{M}_{[t_2,T]}} \inf_{\tau \in [t_2,T]} P_{t_2,x_2}(\tau, u, \beta[u])$$

Thus for all $\beta \in \Delta_{[t_2, T]}$.

$$V(t_2, x_2) \leq \sup_{u \in \mathcal{M}(t_2)} \inf_{\tau \in [t_2, T]} P_{t_2, x_2}(\tau, u, \beta[u])$$

Subsequently, for a fixed $\gamma > 0$, $\exists u_2 \in \mathcal{M}_{[t_2, T]}$ such that for all $\tau \in [t_2, T]$ we have

$$V(t_2, x_2) \leq P_{t_2, x_2}(\tau, u_2, \beta[u_1]) + \gamma \qquad (35)$$

Similarly, we can show that for a fixed $\gamma > 0$, $\exists \beta_1 \in \Delta_{[t_1, T]}$ such that

$$V(t_1, x_1) \geq \sup_{u \in \mathcal{M}_{[t_1, T]}} \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u, \beta_1[u]) - \gamma \quad (36)$$

For an arbitrary $u_1 \in \mathcal{M}_{[t_2, T]}$ we have that

$$V(t_1, x_1) \geq \inf_{\tau \in [t_1, T]} P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) - \gamma$$

Using the definition of infimum, $\exists \tau \in [t_1, T]$ such that

$$V(t_1, x_1) \geq P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) - 2\gamma$$
$$\Rightarrow -V(t_1, x_1) \leq -P_{t_1, x_1}(\tau, u_1, \beta_1[u_1]) + 2\gamma \qquad (37)$$

Let us define the extended policy

$$\tilde{u}(\tau) \equiv \begin{cases} u_1(\tau) & \text{if} \quad t_1 \leq \tau < t_2 \\ u_2(\tau) & \text{if} \quad t_2 \leq \tau \leq T \end{cases}$$

Now by using Equation (35) and (37), it follows that

$$V(t_2, x_2) - V(t_1, x_1) \leq$$
$$P_{t_2, x_2}(\tau, \tilde{u}, \beta_1[\tilde{u}]) - P_{t_1, x_1}(\tau, \tilde{u}, \beta_1[\tilde{u}]) + 3\gamma \qquad (38)$$

Following similar arguments as in Part 1 and since we can choose γ to be arbitrarily small, we obtain $|V(t, x) - V(\hat{t}, \hat{x})| \leq C_4(|t - \hat{t}| + |x - \hat{x}|)$. ∎

Next, let us recall the definition of the modified Hamiltonian

$$\mathcal{H}(t, x, p) = \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} [f(t, x, u, d) \cdot p + \epsilon]. \qquad (39)$$

Then the final Lemma required for Theorem 1 is

Lemma 4: Let $\varphi \in C^1([t_0, T] \times \mathbb{R}^m)$ and $\theta > 0$. Then if φ satisfies

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \leq -\theta \leq 0. \qquad (40)$$

then for a small enough $\sigma > 0$, there exists $\beta \in \Delta_{[t_0, t_0 + \sigma]}$ such that for all $u \in \mathcal{M}_{[t_0, t_0 + \sigma]}$

$$\varphi((t_0 + \sigma), \phi(t_0 + \sigma, t_0, x_0, u(\cdot), \beta(\cdot))) - \varphi(t_0, x_0)$$
$$+ \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq -\frac{\theta \sigma}{2}$$

Conversely, if φ satisfies

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \geq \theta \geq 0 \qquad (41)$$

then for a small enough $\sigma > 0$, there exists $u \in \mathcal{M}_{[t_0, t_0 + \sigma]}$ such that for all $\beta \in \Delta_{[t_0, t_0 + \sigma]}$

$$\varphi((t_0 + \sigma), \phi(t_0 + \sigma, t_0, x_0, u(\cdot), \beta(\cdot))) - \varphi(t_0, x_0)$$
$$+ \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta \sigma}{2}.$$

Proof: **Part I** we define

$$\Lambda(t, x, u, d) = D_t \varphi(t, x) + f(t, x, u, d) \cdot D_x \varphi(t, x) + \epsilon.$$

Then if $\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u, d) \leq -\theta < 0$ then for each $u \in \mathcal{U}$ there exists $d = d(u) \in \mathcal{D}$ such that $\Lambda(t_0, x_0, u, d) \leq -\theta$. Since Λ is uniformly continuous we have

$$\Lambda(t_0, x_0, \tilde{u}, d) \leq \frac{-3\theta}{4}$$

for all $\tilde{u} \in \mathcal{B}(u, r) \cap \mathcal{U}$ and some $r = r(u) > 0$. Since $\mathcal{U}$ is compact there exist finitely many distinct points $u_1, u_2, ... u_n \in \mathcal{U}$, $d_1, d_2, ... d_n \in \mathcal{D}$ and $r_1, r_2, ... r_n > 0$ such that

$$\mathcal{U} \subset \bigcup_{i=1}^{n} B(u_i, r_i)$$

and

$$\Lambda(t_0, x_0, \tilde{u}, d_i) \leq \frac{-3\theta}{4} \quad \text{for} \quad \tilde{u} \in \mathcal{B}(u_i, r_i)$$

We define $\beta_1 : \mathcal{U} \to \mathcal{D}$, setting

$$\beta_1(u) = d_k \text{ if } u \in \mathcal{B}(u_k, r_k) \setminus \bigcup_{i=1}^{k-1} \mathcal{B}(u_i, r_i) \quad (k = 1, \ldots, n).$$

Thus $\Lambda(t_0, x_0, u, \beta_1(u)) \leq \frac{-3\theta}{4}$ for all $u \in \mathcal{U}$. Since Λ is uniformly continuous we therefore have for each sufficiently small $\delta > 0$

$$\Lambda(s, x(s), u, \beta_1(u)) \leq \frac{-\theta}{2}$$

for all $u \in \mathcal{U}, t_0 \leq s \leq t_0 + \delta$ and any solution $x(\cdot)$ of Equation (1) on $(t_0, t_0 + \delta)$ for any $d(\cdot)$, $u(\cdot)$ with initial condition $x(t_0) = x_0$.

Finally we define $\beta \in \Delta_{[t_0, T]}$ in the following way:

$$\beta[u](s) = \beta_1(u(s))$$

for each $u \in \mathcal{M}_{[t_0, T]}$. It then follows that

$$\Lambda(s, x(s), u(s), \beta[u](s)) \leq \frac{-\theta}{2} \quad (t_0 \leq s \leq t_0 + \sigma),$$

for each $u \in \mathcal{M}_{[t_0, T]}$. Notice that

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) = \varphi(t_0, x_0)$$
$$+ \int_{t_0}^{t_0 + \sigma} f(s, x(s), u(s), \beta[u](s)) \cdot D_x \varphi(s, x(s))$$
$$+ D_t \varphi(s, x(s)) ds, \qquad (42)$$

such that integrating over Λ yields

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq \frac{-\theta \sigma}{2}$$

Part II Set

$$\Lambda(t, x, u, d) = D_t \varphi(t, x) + f(t, x, u, d) \cdot D_x \varphi(t, x) + \epsilon.$$

Then if

$$\max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u, d) \geq \theta > 0,$$

there exists a $u^* \in \mathcal{U}$ such that

$$\min_{d \in \mathcal{D}} \Lambda(t_0, x_0, u^*, d). \geq \theta$$

Since Λ is uniformly continuous, we have that

$$\Lambda(s, x(s), u^*, d) \geq \frac{\theta}{2}$$

provided $t_0 \leq s \leq t_0 + \delta$ (for any small $\delta > 0$) and $x(\cdot)$ solves ODE on $(t_0, t_0 + \delta)$ for any $u(\cdot)$, $d(\cdot)$ with initial candidates $x(t_0) = x_0$. Hence for $u(\cdot) \equiv u^*$ and any and $\beta \in \Delta_{[t_0, T]}$

$$D_t \varphi(s, x(s)) + f(s, x(s), u(s), \beta[u](s)) \cdot D_x \varphi(s, x(s))$$
$$+ \epsilon \geq \frac{\theta}{2}$$

Integrating this expression form t_0 to $t_0 + \sigma$ and subtracting $\varphi(t_0, x_0)$ we obtain

$$\varphi((t_0 + \sigma), x(t_0 + \delta)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta \sigma}{2}$$

■

We are now in a position to provide the proof of Theorem 1

D. Proof of Theorem 1

Proof: To prove this it is sufficient to prove the following ([11], [27], [32])

1) For $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ such that $V - \varphi$ attains a local maximum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$ then

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \geq 0$$

2) For $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ such that $V - \varphi$ has a local minimum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$ then

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \leq 0$$

Part 1
Let $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ and suppose that $V - \varphi$ attains a local maximum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$. Let us assume, for the sake of contradiction, that $\exists \theta > 0$ such that

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \leq -\theta \leq 0$$

Case I: $\mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) < \epsilon$.

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0)) \leq -\theta$$

According to the Lemma 4, this implies

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq \frac{-\theta \sigma}{2}$$
$$(43)$$

Since we know that $V - \varphi$ has a maximum at (t_0, x_0), it follows that $V(t_0, x_0) - \varphi(t_0, x_0) \geq V(t_0 + \sigma, x_0(t_0 + \sigma)) -$

$\varphi(t_0 + \sigma, x_0(t_0 + \sigma))$, such that, using the result of Equation (43),

$$V(t_0 + \sigma, x_0(t_0 + \sigma)) + \int_{t_0}^{t_0 + \sigma} \epsilon ds + \frac{\theta \sigma}{2} \leq V(t_0, x_0)$$

But this leads to a contradiction, since we know from Theorem 2 that

$$V(t, x) \leq V(t + \sigma, x(t + \sigma)) + \int_t^{t + \sigma} \epsilon ds.$$

Case II: $\epsilon \leq \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))$

$$D_t \varphi(t_0, x_0) + \epsilon \leq -\theta$$

Integrating the above from t_0 to $t_0 + \sigma$, it follows that

$$\varphi(t_0 + \sigma, x_0) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \leq -\theta \sigma \quad (44)$$

Since we know that $V - \varphi$ has a maximum, following similar argumentation as before,

$$V(t_0 + \sigma, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds + \theta \sigma \leq V(t_0, x_0)$$

But from Lemma 1 we know that

$$V(t, x) \leq V(t + \sigma, x) + \int_{t_0}^{t_0 + \sigma} \epsilon ds.$$

This contradiction implies

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x \varphi(t_0, x_0))\} \geq 0$$

Part 2
Let $\varphi(t, x) \in C^1([0, T] \times \mathbb{R}^m)$ and suppose that $V - \varphi$ attains a local minimum at $(t_0, x_0) \in [0, T] \times \mathbb{R}^m$. We assume, for the sake of contradiction, that $\exists \theta > 0$ such that

$$D_t \varphi(t_0, x_0) + \min\{\epsilon, \mathcal{H}(t_0, x_0, D_x V)\} \geq \theta. \quad (45)$$

This implies

$$D_t \varphi(t_0, x_0) + \epsilon \geq \theta$$

such that

$$\varphi(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \theta \sigma \geq \frac{\theta \sigma}{2}.$$

Furthermore, Equation (45) implies

$$D_t \varphi(t_0, x_0) + \mathcal{H}(t_0, x_0, D_x V) \geq \theta.$$

such that by Lemma 4, we obtain

$$\varphi((t_0 + \sigma), x(t_0 + \sigma)) - \varphi(t_0, x_0) + \int_{t_0}^{t_0 + \sigma} \epsilon ds \geq \frac{\theta \sigma}{2}$$

Since $V - \varphi$ obtains a local minimum at (t_0, x_0), i.e.

$$V(t_0, x_0) - \varphi(t_0, x_0) \leq$$
$$V(t_0 + \sigma, x(t_0 + \sigma)) - \varphi(t_0 + \sigma, x(t_0 + \sigma)).$$

it follows that

$$V(t_0 + \sigma, x(t_0 + \sigma)) - V(t_0, x_0) + \int_{t_0}^{t_0+\sigma} \epsilon ds \geq \frac{\theta \sigma}{2}.$$
(46)

Following the same logic and we also obtain

$$V(t_0 + \sigma, x_0) - V(t_0, x_0) + \int_{t_0}^{t_0+\sigma} \epsilon ds \geq \frac{\theta \sigma}{2}.$$
(47)

Recall the definition of $V(t, x)$ obtained from Theorem 2. We consider two cases

Case I

$$V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0+\sigma} \epsilon ds$$

$$\leq \inf_{\tau \in [t_0, t_0+\sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right]$$

then

$$V(t_0, x_0) = V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0+\sigma} \epsilon ds$$

substituting this in equation (46) leads to a contradiction.

Case II

$$\inf_{\tau \in [t_0, t_0+\sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right] \leq$$

$$V(t_0 + \sigma, x(t_0 + \sigma)) + \int_{t_0}^{t_0+\sigma} \epsilon ds$$

Then

$$V(t_0, x_0) = \inf_{\beta(\cdot) \in \Delta_{[t_0, t_0+\sigma]}} \sup_{\nu(\cdot) \in \mathcal{M}_{[t_0, t_0+\sigma]}}$$

$$\inf_{\tau \in [t_0, t_0+\sigma]} \left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right]$$

The minima in τ must occur at $\tau = t_0$ and the minimizer is unique. If this was not true then there would exist some $\tau \in [t_0, t_0 + \sigma]$ such that

$$V(t_0, x_0) = \inf_{\beta(\cdot) \in \Delta_{[t_0, t_0+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t_0, t_0+\sigma]}}$$

$$\left[\int_{t_0}^{\tau} \epsilon ds + g(x(\tau)) \right]$$

Thus using corollary 1 we get that

$$V(\tau, x(\tau)) + \int_{t_0}^{\tau} \epsilon ds \leq V(t_0, x_0), \quad \tau \in [t_0, t_0 + \sigma]$$

This will lead to contradiction with equation (46). Furthermore, we know that for all $\tau \in [t_0, t_0 + \sigma]$ using Lemma 1

$$V(t_0, x_0) \leq V(t_0 + \sigma, x_0) +$$

$$\inf_{\beta(\cdot) \in \Delta_{[t_0, t_0+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t_0, t_0+\sigma]}} \int_{t_0}^{t_0+\sigma} \epsilon ds \leq g(x_0)$$

Therefore we have

$$V(t_0, x_0) = V(t_0 + \sigma, x_0) +$$

$$\inf_{\beta(\cdot) \in \Delta_{[t_0, t_0+\sigma]}} \sup_{u(\cdot) \in \mathcal{M}_{[t_0, t_0+\sigma]}} \int_{t_0}^{t_0+\sigma} \epsilon ds = g(x_0)$$

This along with equation (47) leads to contradiction. ∎

REFERENCES

[1] A. Abate, C. David, P. Kesseli, D. Kroening, and E. Polgreen. *Counterexample Guided Inductive Synthesis Modulo Theories*, page 270–288. Springer International Publishing, 2018.

[2] E. Asarin, O. Bournez, T. Dang, O. Maler, and A. Pnueli. Effective synthesis of switching controllers for linear systems. *Proceedings of the IEEE*, 88(7):1011–1025, 2000.

[3] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011.

[4] J.-P. Aubin and H. Frankowska. Viability kernel of control systems. In *Nonlinear synthesis*, pages 12–33. Springer, 1991.

[5] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.

[6] S. Bansal and C. J. Tomlin. Deepreach: A deep learning approach to high-dimensional reachability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1817–1824. IEEE, 2021.

[7] P. Cardaliaguet. A differential game with two players and one target. *SIAM Journal on Control and Optimization*, 34(4):1441–1460, 1996.

[8] P. Cardaliaguet, M. Quincampoix, and P. Saint-Pierre. Differential games through viability theory: Old and recent results. In *Advances in dynamic game theory*, pages 3–35. Springer, 2007.

[9] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin. Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking. *IEEE Transactions on Automatic Control*, 66(12):5861–5876, 2021.

[10] 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*, 63(11):3675–3688, 2018.

[11] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert. Robust control barrier–value functions for safety-critical control. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 6814–6821. IEEE, 2021.

[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):773–797, 1984.

[13] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure1. *Journal on Satisfiability, Boolean Modeling and Computation*, 1(3–4):209–236, May 2007.

[14] S. Gao, J. Avigad, and E. M. Clarke. δ-complete decision procedures for satisfiability over the reals. In *IJCAR*, volume 7364 of *Lecture Notes in Computer Science*, pages 286–300. Springer, 2012.

[15] S. Gao, S. Kong, and E. M. Clarke. *dReal: An SMT Solver for Nonlinear Theories over the Reals*, page 208–214. Springer Berlin Heidelberg, 2013.

[16] S. Gao, S. Kong, and E. M. Clarke. dreal: An smt solver for nonlinear theories over the reals. In *International conference on automated deduction*, pages 208–214. Springer, 2013.

[17] D. Kroening and O. Strichman. *Decision Procedures - An Algorithmic Point of View*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2008.

[18] A. Lin and S. Bansal. Generating formal safety assurances for high-dimensional reachability. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, page 10525–10531. IEEE, May 2023.

[19] A. Lin and S. Bansal. Verification of neural reachable tubes via scenario optimization and conformal prediction. In *Proceedings of the 6th Annual Learning for Dynamics and Control Conference*, volume 242 of *Proceedings of Machine Learning Research*, pages 719–731. PMLR, 15–17 Jul 2024.

[20] C. Livadas, J. Lygeros, and N. A. Lynch. High-level modeling and analysis of the traffic alert and collision avoidance system (tcas). *Proceedings of the IEEE*, 88(7):926–948, 2000.

[21] C. Livadas and N. A. Lynch. Formal verification of safety-critical hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 253–272. Springer, 1998.

[22] J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

[23] J. Lygeros, D. N. Godbole, and S. Sastry. Verified hybrid controllers for automated vehicles. *IEEE transactions on automatic control*, 43(4):522–539, 1998.

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

[25] J. N. Maidens, S. Kaynama, I. M. Mitchell, M. M. Oishi, and G. A. Dumont. Lagrangian methods for approximating the viability kernel in high-dimensional systems. *Automatica*, 49(7):2017–2029, 2013.

[26] K. Margellos and J. Lygeros. Hamilton–jacobi formulation for reach–avoid differential games. *IEEE Transactions on automatic control*, 56(8):1849–1861, 2011.

[27] I. Mitchell, A. M. Bayen, and C. J. Tomlin. Validating a hamilton-jacobi approximation to hybrid system reachable sets. In *International workshop on hybrid systems: Computation and control*, pages 418–432. Springer, 2001.

[28] I. M. Mitchell et al. A toolbox of level set methods. *UBC Department of Computer Science Technical Report TR-2007-11*, 1:6, 2007.

[29] M. Raissi, P. Perdikaris, and G. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, Feb. 2019.

[30] S. Sastry. *Nonlinear Systems*. Springer New York, 1999.

[31] M. Tayal, A. Singh, S. Kolathaya, and S. Bansal. A physics-informed machine learning framework for safe and optimal control of autonomous systems. *arXiv preprint: 2502.11057*, 2025.

[32] C. Tomlin, I. Mitchell, and R. Ghosh. Safety verification of conflict resolution manoeuvres. *IEEE Transactions on Intelligent Transportation Systems*, 2(2):110–120, 2001.

[33] C. J. Tomlin, J. Lygeros, and S. S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88(7):949–970, 2000.

[34] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Safety-aware hybrid control of airborne wind energy systems. *Journal of Guidance, Control, and Dynamics*, 47(2):326–338, 2024.

[35] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Multi-objective minimum time optimal control for low-thrust trajectory design. In *2021 European Control Conference (ECC)*, pages 1975–1980, 2021.

[36] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Verification of safety critical control policies using kernel methods. In *2022 European Control Conference (ECC)*, pages 1870–1875, 2022.

[37] N. Vertovec, S. Ober-Blöbaum, and K. Margellos. Multi-objective low-thrust spacecraft trajectory design using reachability analysis. *European Journal of Control*, 69:100758, 2023.