

Contraction-Guided Adaptive Partitioning for Reachability Analysis of Neural Network Controlled Systems

Akash Harapanahalli, Saber Jafarpour, and Samuel Coogan

Abstract—In this paper, we present a contraction-guided adaptive partitioning algorithm for improving interval-valued robust reachable set estimates in a nonlinear feedback loop with a neural network controller and disturbances. Based on an estimate of the contraction rate of over-approximated intervals, the algorithm chooses when and where to partition. Then, by leveraging a decoupling of the neural network verification step and reachability partitioning layers, the algorithm can provide accuracy improvements for little computational cost. This approach is applicable with any sufficiently accurate open-loop interval-valued reachability estimation technique and any method for bounding the input-output behavior of a neural network. Using contraction-based robustness analysis, we provide guarantees of the algorithm’s performance with mixed monotone reachability. Finally, we demonstrate the algorithm’s performance through several numerical simulations and compare it with existing methods in the literature. In particular, we report a sizable improvement in the accuracy of reachable set estimation in a fraction of the runtime as compared to state-of-the-art methods.

I. INTRODUCTION

Neural networks have become increasingly popular in control systems in recent years due to their relative ease for in-the-loop computation. These learning-based algorithms are known to be vulnerable to input perturbations—small (possibly adversarial) changes in their input can lead to large variations in their output [1]. As such, runtime verification of the safety and performance of neural network controlled systems is essential in safety-critical applications. This task is generally challenging due to the nonlinear and large-scale structure of the neural networks and their interconnection with nonlinear dynamics [2].

A basic ingredient for verifying control systems is the ability to overapproximate the set of reachable states from a given set of initial conditions, possibly in the presence of disturbances. If, for example, this overapproximation avoids obstacles or reaches a goal region, then the system certifiably satisfies the corresponding safety or performance criteria.

Recently, several promising reachability-based methods have been proposed for verifying stand-alone neural networks or feedback systems with neural networks in the control loop; however, these methods either suffer from

*This work is supported in part by Cisco Systems, the National Science Foundation under Grants #1749357 and #2219755, the Air Force Office of Scientific Research under Grant FA9550-23-1-0303, and NASA under the University Leadership (ULI) Award #80NSSC20M0161 but solely reflects the opinions and conclusions of its authors.

Akash Harapanahalli, Saber Jafarpour, and Samuel Coogan are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, USA, {aharapan, saber, sam.coogan}@gatech.edu




Fig. 1. A snapshot of Algorithm 1 is illustrated for $n = 1$. (top) There are two initial partitions that both compute the neural network verification step (filled circles in the graph representation). Both partitions are integrated from t_{j-1} to $t_{j-1} + \gamma(t_j - t_{j-1})$ to estimate the width of the box at t_j . Since the estimated blue partition width violates the user-defined maximum width ϵ , (bottom) the algorithm returns to t_{j-1} and adds two new partitions to the tree. The maximum neural network verification depth here is 1, so the new partitions do not recompute the neural network verification and instead use the same E^c (14) from the initial blue partition.

large computational complexity, large over-approximation error, or lack of generality. For stand-alone neural networks: Interval Bound Propagation (IBP) [3] is fast, but largely over-conservative; CROWN [4] propagates linear bounds through the network; LipSDP [5] uses semi-definite programming, but is not scalable to large networks; NNV [6] is a set-based verification framework implemented in MATLAB. For neural network closed-loop verification: ReachLP [7] is computationally light, but it can be conservative and only applies to linear systems; ReachSDP [8] can provide tighter bounds, but does not scale well to large networks and only applies to linear systems; JuliaReach [9], CORA [10], NNV [6], POLAR [11], MILP [12], and constrained Zonotope [13] methods can provide very accurate estimates for non-linear systems, but are too expensive for runtime verification. ReachMM [14] combines off the shelf open-loop verifiers with mixed monotonicity to provide efficient but coarse interval over-estimates.

Partitioning of the state space is an effective approach for balancing the trade-off between the accuracy of reachable set estimates and the runtime of the verification strategy. In the literature, there are many existing partitioning algorithms. For standalone ReLU networks, [15] splits the input set into partitions based on the stability of each neuron. In [16], [17], Monte Carlo simulations are used to guide which specific axes to cut. For closed-loop linear systems, ReachLP-Uniform [7] uses a uniform partitioning of the state space, ReachLP-GSG [7] applies the GSG algorithm from [17] to ReachLP, and ReachLipBnB [18] combines a branch and bound algorithm with ℓ_2 -Lipschitz constants obtained from LipSDP.

Contributions: We introduce a contraction-guided adaptive partitioning algorithm for general nonlinear systems with

neural network controllers. In Section VI, we use contraction theory to prove that for the mixed monotone framework ReachMM [14], the expansion rate of interval partitions is directly linked with the contraction rate of the embedding system. In particular, we provide upper bounds on the width of interval reachable set over-estimates—showing that their exponential growth in time accumulates from three main sources: (1) the maximum rate of expansion of the closed loop system in the specific partition, (2) the initial width of the partition, and (3) the approximation error of the neural network verifier. This motivates the creation of Algorithm 1 in Section V, which provides three key features for improving these sources of error. The algorithm is *separative*—it decouples the neural network verification calls and reachability partitioning layers, leading to inexpensive accuracy benefits on (2) without suffering additional computational cost on improving (3); *spatially aware*—it chooses specific regions of the state space to cut based on the contraction rate of the partitions, allocating resources according to where in space (1) is the worst; and *temporally aware*—it chooses to cut on along the trajectory as estimates blow up, allocating resources according to when in time (1) is the worst. By applying the algorithm to the mixed monotone framework [14], ReachMM-CG is applicable in continuous or discrete time, for nonlinear plant models controlled by neural networks with general nonlinear activation functions; this removes several limitations found in most of the above-cited approaches. We find that ReachMM-CG yields reachable set estimates with a 33% accuracy improvement in a quarter of the time compared to state-of-the-art partitioning algorithms on a benchmark example. Finally, while our bounds apply for specifically the strategy proposed in [14], we note that the partitioning algorithm itself can also be coupled with any interval-based verification methods, including many of those discussed above, as well as any future advancements.

II. NOTATION AND MATHEMATICAL PRELIMINARY

We denote the set of extended real numbers by $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. We define the partial order \leq on \mathbb{R}^n as $x \leq y \iff x_i \leq y_i \forall i \in \{1, \dots, n\}$. For every $x \leq y$, we define the interval $[x, y] = \{z : x \leq z \leq y\}$. For $x, \hat{x} \in \mathbb{R}^n$, let $(x, \hat{x}) \in \mathbb{R}^{2n}$ be their concatenation. The southeast partial order \leq_{SE} on \mathbb{R}^{2n} is induced by \leq on \mathbb{R}^n as follows: $(x, \hat{x}) \leq_{\text{SE}} (y, \hat{y}) \iff x \leq y$ and $\hat{y} \leq \hat{x}$. Define the following: $\mathcal{T}_{\geq 0}^{2n} = \{(x, \hat{x}) \in \mathbb{R}^{2n} : x \leq \hat{x}\}$, $\mathcal{T}_{\leq 0}^{2n} = \{(x, \hat{x}) \in \mathbb{R}^{2n} : \hat{x} \leq x\}$, $\mathcal{T}^{2n} = \mathcal{T}_{\geq 0}^{2n} \cup \mathcal{T}_{\leq 0}^{2n}$. We define the vector $x_{[i:\hat{x}]} \in \mathbb{R}^n$ as $(x_{[i:\hat{x}]})_j = \begin{cases} x_j, & j \neq i \\ \hat{x}_j, & j = i \end{cases}$. Define the weighted ℓ_∞ -norm $\|x\|_{\infty, \varepsilon} = \|\text{diag}(\varepsilon)^{-1}x\|_\infty$. In particular, the weighted maximum width of an interval $[x, \bar{x}]$ is $\|\bar{x} - x\|_{\infty, \varepsilon}$. Given a matrix $A \in \mathbb{R}^{n \times n}$, the ℓ_∞ -matrix measure of A is defined by $\mu_\infty(A) = \max_{i \in \{1, \dots, n\}} \{A_{ii} + \sum_{j \neq i} |A_{ij}|\}$. From [19, Table III], we define the weak pairing $\llbracket \cdot, \cdot \rrbracket_\infty : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ associated to the norm $\|\cdot\|_\infty$ as follows:

$$\llbracket x, y \rrbracket_\infty = \max_{i \in I_\infty(y)} y_i x_i,$$

where $I_\infty(x) = \{i \in \{1, \dots, n\} \mid |x_i| = \|x\|_\infty\}$. Consider the following dynamical system

$$\dot{x} = f(x, w) \quad (1)$$

with state vector $x \in \mathbb{R}^n$ and disturbance vector $w \in \mathcal{W} \subset \mathbb{R}^q$. Given a piecewise continuous curve $w(\cdot)$ where $w : [t_0, t] \rightarrow \mathcal{W}$, the trajectory of the system starting from x_0 at time t_0 is given $t \rightarrow \phi_f(t, t_0, x_0, w(\cdot))$. Given an initial set \mathcal{X} , we denote the reachable set of f at some $t \geq t_0$:

$$\mathcal{R}_f(t, t_0, \mathcal{X}, \mathcal{W}) = \left\{ \begin{array}{l} \phi_f(t, t_0, x_0, w(\cdot)), \forall x_0 \in \mathcal{X}, \\ w : \mathbb{R} \rightarrow \mathcal{W} \text{ piecewise cont.} \end{array} \right\} \quad (2)$$

The dynamical system (1) is mixed monotone with respect to the decomposition function $d : \mathcal{T}^{2n} \times \mathcal{T}^{2q} \rightarrow \mathbb{R}^n$ if, for every $i \in \{1, \dots, n\}$,

- i. $d_i(x, x, w, w) = f_i(x, w)$, for every $x \in \mathbb{R}^n$, $w \in \mathbb{R}^q$;
- ii. $d_i(x, \hat{x}, w, \hat{w}) \leq d_i(y, \hat{y}, w, \hat{w})$, for every $x \leq y$ s.t. $x_i = y_i$, and every $\hat{y} \leq \hat{x}$;
- iii. $d_i(x, \hat{x}, w, \hat{w}) \leq d_i(x, \hat{x}, v, \hat{v})$, for every $w \leq v$ and every $\hat{v} \leq \hat{w}$.

With a valid decomposition function d , one can construct an embedding system associated to (1) as follows:

$$\frac{d}{dt} \begin{bmatrix} \underline{x} \\ \bar{x} \end{bmatrix} = \begin{bmatrix} d(\underline{x}, \bar{x}, \underline{w}, \bar{w}) \\ d(\bar{x}, \underline{x}, \bar{w}, \underline{w}) \end{bmatrix} := E(\underline{x}, \bar{x}, \underline{w}, \bar{w}). \quad (3)$$

III. PROBLEM STATEMENT

Consider a nonlinear continuous-time dynamical system of the form

$$\dot{x}(t) = f(x(t), u(t), w(t)), \quad (4)$$

where $x \in \mathbb{R}^n$ is the state of the system, $u \in \mathbb{R}^p$ is the control input to the system, $w \in \mathbb{R}^q$ is a disturbance input to the system, and $f : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}^n$ is a parameterized vector field. We assume that the feedback control policy for the system (4) is given by a k -layer fully connected feed-forward neural network $N : \mathbb{R}^n \rightarrow \mathbb{R}^p$ as follows:

$$\begin{aligned} \xi^{(i)} &= \sigma^{(i-1)} \left(W^{(i-1)} \xi^{(i-1)} + b^{(i-1)} \right), \quad i = 1, \dots, k \\ \xi^{(0)} &= x, \quad N(x) = W^{(k)} \xi^{(k)} + b^{(k)} \end{aligned} \quad (5)$$

where m_i is the number of neurons in the i -th layer, $W^{(i)} \in \mathbb{R}^{m_i \times m_{i-1}}$ is the weight matrix on the i -th layer, $b^{(i)} \in \mathbb{R}^{m_i}$ is the bias vector on the i -th layer, $\xi^{(i)} \in \mathbb{R}^{m_i}$ is the i -th layer hidden variable and σ_i is the activation function for the i -th layer. In practice, $N(x(t))$ cannot be evaluated at each instance of time t , and, instead, the control must be implemented via, *e.g.*, a zero-order hold strategy between sampling instances. We assume that there exists an increasing sequence of control time instances $\{t_0, t_1, t_2, \dots\}$ in which the control input is updated. Thus, the closed-loop system with the neural-network feedback controller is given by:

$$\dot{x}(t) = f(x(t), N(x(t_j)), w(t)) := f^c(j, x(t), w(t)), \quad (6)$$

where $t \in \mathbb{R}_{\geq 0}$ and $j \in \mathbb{Z}_{\geq 0}$ is such that $t \in [t_j, t_{j+1}]$. In our analysis when j is clear from context or does not affect the result, we drop j as an argument of f^c .

The goal of this paper is to verify the behavior of the closed-loop system (6). To verify the safety of a system under uncertainty, one needs to verify the entire reachable set. However, in general, computing the reachable set exactly is not computationally tractable—instead, approaches typically compute an over-approximation $\mathcal{R}_{fc}(t, t_0, \mathcal{X}, \mathcal{W}) \supseteq \mathcal{R}_{fc}(t, t_0, \mathcal{X}, \mathcal{W})$. Therefore, the main challenge addressed in this paper is to develop an approach for providing tight over-approximations of reachable sets while remaining computationally tractable for runtime computation.

IV. INTERVAL REACHABILITY OF NEURAL NETWORK CONTROLLED SYSTEMS

A. General Framework

We assume we have access to an off-the-shelf dynamical system reachability tool that supports *interval analysis*.

Assumption 1 (Open-loop interval reachability). *Given a dynamical system of the form (4), any intervals $\mathcal{X}_0 = [\underline{x}_0, \bar{x}_0] \subseteq \mathbb{R}^n$, $\mathcal{U} = [\underline{u}, \bar{u}] \subseteq \mathbb{R}^p$, and $\mathcal{W} = [\underline{w}, \bar{w}] \subseteq \mathbb{R}^q$, and some initial time t_0 , there exists a reachability algorithm that returns a valid interval approximation $[\underline{x}(t), \bar{x}(t)]$ satisfying*

$$\mathcal{R}_f(t, t_0, \mathcal{X}_0, \mathcal{U}, \mathcal{W}) \subseteq [\underline{x}(t), \bar{x}(t)], \quad \forall t \geq t_0. \quad (7)$$

Reachability analysis of dynamical systems is a classical and well-studied research field with several off-the-shelf toolboxes with this capability, including Flow* [20], Hamilton-Jacobi approach [21], CORA [22], and mixed monotonicity [23].

We also assume access to off-the-shelf neural network verification algorithms that can provide *interval inclusion functions* as follows:

Assumption 2 (Neural network verification). *Given a neural network N of the form (5) and any interval $[y, \bar{y}] \subseteq \mathbb{R}^n$, there exists a neural network verification algorithm that returns a valid inclusion function $\begin{bmatrix} N \\ \bar{N} \end{bmatrix}_{[y, \bar{y}]} : \mathcal{T}_{\geq 0}^{2n} \rightarrow \mathcal{T}_{\geq 0}^{2p}$ satisfying*

$$\underline{N}_{[y, \bar{y}]}(\underline{x}, \bar{x}) \leq N(x) \leq \bar{N}_{[y, \bar{y}]}(\underline{x}, \bar{x}), \quad (8)$$

for any $x \in [\underline{x}, \bar{x}] \subseteq [y, \bar{y}]$.

A large number of the existing neural network verification algorithms can provide bounds of the form (8) for the output of the neural networks, including CROWN (and its subsequent variants) [4], LipSDP [5], and IBP [3].

Combining Assumptions 1 and 2 leads naturally to an algorithm for over-approximating solutions to (6). In particular, starting with $j = 0$ and an initial interval of states $\mathcal{X}_j = [\underline{x}(t_j), \bar{x}(t_j)]$, first obtain $\underline{N}_{[y, \bar{y}]}$ and $\bar{N}_{[y, \bar{y}]}$ for some $[y, \bar{y}] \supseteq \mathcal{X}_j$ from Assumption 2. Next, set $\mathcal{U} = [\underline{N}_{[y, \bar{y}]}(\underline{x}(t_j), \bar{x}(t_j)), \bar{N}_{[y, \bar{y}]}(\underline{x}(t_j), \bar{x}(t_j))]$. Finally, compute the reachable set of the closed-loop system on the interval $t \in (t_j, t_{j+1}]$ as the interval reachable set obtained from the algorithm of Assumption 1, set \mathcal{X}_{j+1} as the reachable set at time t_{j+1} , increment $j \leftarrow j + 1$, and iterate. This iteration serves as the backbone of our proposed partitioning-based algorithm below.

B. Mixed Monotone Framework

One specific framework that satisfies both Assumptions 1 and 2 is developed in [14], where interval reachability of neural network controlled systems is studied through the lens of mixed monotone dynamical systems theory. Suppose that we have access to a decomposition function d for the open-loop system (4) with the open-loop embedding system

$$\frac{d}{dt} \begin{bmatrix} \underline{x} \\ \bar{x} \end{bmatrix} = \begin{bmatrix} d(\underline{x}, \bar{x}, \underline{u}, \bar{u}, \underline{w}, \bar{w}) \\ d(\bar{x}, \underline{x}, \bar{u}, \underline{u}, \bar{w}, \underline{w}) \end{bmatrix} := E(\underline{x}, \bar{x}, \underline{u}, \bar{u}, \underline{w}, \bar{w}). \quad (9)$$

Given $\mathcal{X}_0 = [\underline{x}_0, \bar{x}_0] \subseteq \mathbb{R}^n$, $\mathcal{U} = [\underline{u}, \bar{u}] \subseteq \mathbb{R}^p$, and $\mathcal{W} = [\underline{w}, \bar{w}] \subseteq \mathbb{R}^q$, the trajectory of (9) starting from $(\underline{x}_0, \bar{x}_0)$ provides the inclusion (7) from Assumption 1.

Following the treatment in [14], one can use CROWN [4] to obtain the desired bounds from Assumption 2. Given an interval $[y, \bar{y}]$, the algorithm provides an efficient procedure for finding a tuple $(\underline{C}, \bar{C}, \underline{d}, \bar{d})$ defining linear upper and lower bounds for the output of the neural network

$$\underline{C}(y, \bar{y})x + \underline{d}(y, \bar{y}) \leq N(x) \leq \bar{C}(y, \bar{y})x + \bar{d}(y, \bar{y}), \quad (10)$$

for every $x \in [y, \bar{y}]$. Using these linear bounds, we can construct the inclusion function for any $[\underline{x}, \bar{x}] \subseteq [y, \bar{y}]$:

$$\begin{aligned} \underline{N}_{[y, \bar{y}]}(\underline{x}, \bar{x}) &= [\underline{C}(y, \bar{y})]^+ \underline{x} + [\underline{C}(y, \bar{y})]^- \bar{x} + \underline{d}(y, \bar{y}), \\ \bar{N}_{[y, \bar{y}]}(\underline{x}, \bar{x}) &= [\bar{C}(y, \bar{y})]^+ \bar{x} + [\bar{C}(y, \bar{y})]^- \underline{x} + \bar{d}(y, \bar{y}). \end{aligned} \quad (11)$$

By combining these two tools, one can construct a new closed-loop embedding system to over-approximate the reachable sets of the closed-loop system (6). From [14], with the following definitions,

$$\begin{aligned} \underline{\eta}_j &= \underline{N}_{[y, \bar{y}]}(\underline{x}(t_j), \bar{x}(t_j)_{[i:\underline{x}(t_j)]}), \\ \bar{\eta}_j &= \bar{N}_{[y, \bar{y}]}(\underline{x}(t_j), \bar{x}(t_j)_{[i:\underline{x}(t_j)]}), \\ \underline{\nu}_j &= \underline{N}_{[y, \bar{y}]}(\bar{x}(t_j), \underline{x}(t_j)_{[i:\bar{x}(t_j)]}), \\ \bar{\nu}_j &= \bar{N}_{[y, \bar{y}]}(\bar{x}(t_j), \underline{x}(t_j)_{[i:\bar{x}(t_j)]}), \end{aligned} \quad (12)$$

we use the following “hybrid” function

$$\left(d_{[y, \bar{y}]}^c(j, \underline{x}, \bar{x}, \underline{w}, \bar{w}) \right)_i = \begin{cases} d_i(\underline{x}, \bar{x}, \underline{\eta}_j, \bar{\eta}_j, \underline{w}, \bar{w}), & \underline{x} \leq \bar{x} \\ d_i(\bar{x}, \underline{x}, \bar{\nu}_j, \underline{\nu}_j, \bar{w}, \underline{w}), & \bar{x} \leq \underline{x} \end{cases} \quad (13)$$

for any $[y, \bar{y}] \supseteq [\underline{x}, \bar{x}]$, to create the closed-loop embedding system:

$$\frac{d}{dt} \begin{bmatrix} \underline{x} \\ \bar{x} \end{bmatrix} = \begin{bmatrix} d_{[y, \bar{y}]}^c(j, \underline{x}, \bar{x}, \underline{w}, \bar{w}) \\ d_{[y, \bar{y}]}^c(j, \bar{x}, \underline{x}, \bar{w}, \underline{w}) \end{bmatrix} := E_{[y, \bar{y}]}^c(j, \underline{x}, \bar{x}, \underline{w}, \bar{w}). \quad (14)$$

In our analysis when j is clear from context or does not affect the result, we drop j as an argument of d^c and E^c . For an input set $\mathcal{X}_0 \subseteq [\underline{x}_0, \bar{x}_0]$ and a disturbance set $\mathcal{W} \subseteq [\underline{w}, \bar{w}]$, with $\begin{bmatrix} \underline{x}(t) \\ \bar{x}(t) \end{bmatrix} = \phi_{E_{[y, \bar{y}]}^c}(t, t_0, \begin{bmatrix} \underline{x}_0 \\ \bar{x}_0 \end{bmatrix}, \begin{bmatrix} \underline{w} \\ \bar{w} \end{bmatrix})$ for every $t > t_0$, the following inclusion holds [14, Theorem 1]:

$$\mathcal{R}_{fc}(t, t_0, \mathcal{X}, \mathcal{W}) \subseteq [\underline{x}(t), \bar{x}(t)]. \quad (15)$$

Thus, running one single trajectory of E^c can over-

approximate the reachable set of the closed-loop system (6).

C. Partitioning

While interval analysis techniques can often be computationally inexpensive, as the size of the uncertain sets grow, the reachable set estimates become overly conservative. In a closed-loop, these effects can compound significantly as the error continues to accumulate. This phenomenon is known in the literature as the *wrapping effect*. To mitigate over-conservatism in interval-based techniques, one can split uncertain regions into smaller partitions. A valid over-approximation of the reachable set can be found by taking the union of the over-approximations of the reachable sets for each partition. Due to the locality of the smaller partitions, this technique can drastically reduce the conservatism of interval-based techniques.

V. CONTRACTION-GUIDED ADAPTIVE PARTITIONING

In this section, we introduce a contraction-guided adaptive partitioning algorithm to improve the results of interval reachability methods satisfying Assumptions 1 and 2. This algorithm is motivated by the theory in the next section.

A. Algorithm Description

At a particular control step t_j , define a partition

$$P = ((\underline{x}_0, \bar{x}_0), N, S), \quad (16)$$

where $(\underline{x}_0, \bar{x}_0) \in \mathcal{T}_{>0}^{2n}$ is the initial condition of the partition, $N \in \{\text{TRUE}, \text{FALSE}\}$ is whether or not the partition should compute the neural network verification (8), and S is a set of tuples representing its subpartitions. Algorithm 1 starts by initializing a tree at time t_0 with a single root node $\mathcal{P} = ((\underline{X}_0, \bar{X}_0), \text{TRUE}, \emptyset)$ representing the initial set.

step() Procedure: Along a particular control interval $[t_{j-1}, t_j]$, every partition undergoes the following procedure. If specified, the embedding system E^c is computed using (11, 13, 14). Initially, E^c is integrated to a fraction $\gamma \in (0, 1]$ of the control interval, where the contraction rate is computed by comparing its width to its initial condition. Then, the width of the final box at t_j is estimated. If it violates hyper-parameter ε , then the algorithm returns to time t_{j-1} and divides the initial condition into 2^n sub-partitions. This process is repeated for each of the sub-partitions. Finally, if the estimated width does not violate ε , E^c is fully integrated to t_j . This procedure is visualized in Figure I, and formalized as *step()* in Algorithm 1, which repeats for every control step until some final time T .

Hyper-parameters: There are several user-defined parameters that are important for the algorithm's performance. The choice of the check-contraction factor $\gamma \in (0, 1]$ is purely for computational benefit, and in particular, $\gamma = 1$ checks the true width of the box at t_j instead of an estimate. The maximum partition depth $D_p \in \mathbb{Z}_{\geq 0}$ specifies how deep new partitions can be added in the tree to integrate E^c . The maximum neural network verification depth $D_N \in \mathbb{Z}_{\geq 0}$ specifies how deep neural network verification is performed in the partition tree. D_p and D_N are directly related to the

Algorithm 1: Contraction-Guided Adaptive Partitioning (ReachMM-CG)

Input: initial set: $\mathcal{X}_0 = [\underline{X}_0, \bar{X}_0] \subset \mathbb{R}^n$, control time instances: $\{t_0, t_1, \dots\}$, final time: T
Parameter: desired width $\varepsilon \in \mathbb{R}_{>0}^n$, check contraction factor $\gamma \in (0, 1]$, max partition depth $D_p \in \mathbb{Z}_{\geq 0}$, max neural network verification depth $D_N \in \mathbb{Z}_{\geq 0}$
Output: Over-approximated reachable set trajectory for $t \in [t_0, T]$: $\bar{\mathcal{R}}(t)$

```

1  $\bar{\mathcal{R}}(t_0) = \mathcal{X}_0$ ;
2  $\mathcal{P} = ((\underline{X}_0, \bar{X}_0), \text{TRUE}, \emptyset)$ ;
3  $t_m =$  smallest control time instance  $\geq T$ ;
4 for  $j = 1, \dots, m$  do
5    $(\bar{\mathcal{R}}(\cdot)|_{[t_{j-1}, t_j]}, \mathcal{P}_+) \leftarrow \text{step}(\mathcal{P}, \emptyset, j)$ ;
6    $\mathcal{P} \leftarrow \mathcal{P}_+$ ;
7 end
8 return  $\bar{\mathcal{R}}(t) \forall t \in [t_0, T]$ ;
9 Procedure  $\text{step}((\underline{x}_0, \bar{x}_0), N, S), E^c, j)$ 
10   if  $N$  is TRUE then // NN Verification
11      $E^c \leftarrow E^c$  as (14, 13) using (11) on  $(\underline{x}_0, \bar{x}_0)$ ;
12   if  $S = \emptyset$  then // No Subpartitions
13      $d = \text{get\_depth}()$ ; // Tree Depth
14     if  $d < D_p$  then // Can Partition
15        $t_\gamma \leftarrow t_{j-1} + \gamma(t_j - t_{j-1})$ ;
16        $(\underline{x}, \bar{x})(t_\gamma) \leftarrow \phi_{E^c}(t_\gamma, t_{j-1}, (\underline{x}_0, \bar{x}_0))$ ;
17        $C = \frac{\|\bar{x}(t_\gamma) - \underline{x}(t_\gamma)\|_{\infty, \varepsilon}}{\|\bar{x}_0 - \underline{x}_0\|_{\infty, \varepsilon}}$ ; // Contr Rate
18       if  $C^{1/\gamma} \|\bar{x}_0 - \underline{x}_0\|_{\infty, \varepsilon} > 1$  then
19          $\{\underline{x}_0^k, \bar{x}_0^k\}_{k=1}^{2^n} \leftarrow \text{uni\_div}(\underline{x}_0, \bar{x}_0)$ ;
20          $S \leftarrow \{((\underline{x}_0^k, \bar{x}_0^k), d < D_N, \emptyset)\}_{k=1}^{2^n}$ ;
21          $N \leftarrow N \wedge ((d+1) > D_N)$ ;
22         return  $\text{step}((\underline{x}_0, \bar{x}_0), N, S), E^c, j)$ 
23        $(\underline{x}, \bar{x})(\cdot)|_{[t_{j-1}, t_j]} \leftarrow \phi_{E^c}(\cdot, t_{j-1}, (\underline{x}_0, \bar{x}_0))$ ;
24       return  $([\underline{x}, \bar{x}](\cdot)|_{[t_{j-1}, t_j]}, ((\underline{x}, \bar{x})(t_j), N, \emptyset))$ 
25     else // Iterate Subpartitions
26        $\{(\bar{\mathcal{R}}^k(\cdot), P_+^k) \leftarrow \text{step}(P^k, E^c, j)\}_{P^k \in S}$ ;
27        $(\underline{x}_0, \bar{x}_0)_+ \leftarrow \text{tightest} [\underline{x}, \bar{x}] \supseteq \bigcup_k \bar{\mathcal{R}}^k(t_j)$ ;
28       return  $(\bigcup_k \bar{\mathcal{R}}^k(\cdot), ((\underline{x}_0, \bar{x}_0)_+, N, \{P_+^k\}_{k=1}^{2^n}))$ ;

```

first two terms of the right-hand side of Theorem 1 in the next section—every partitioning layer will improve the first term, while only partitions computing the neural network verification will improve the second term. Finally, the choice of desired width $\varepsilon \in \mathbb{R}_{>0}^n$ specifies how wide partitions can grow before being cut: if chosen too small, the algorithm will partition early in the start ($\varepsilon_i = 0$ for any i implies a uniform initial partitioning); if chosen too large, the algorithm will not partition at all ($\varepsilon_i = \infty$ implies no condition on the i -th component).

B. Discussion

Algorithm 1 has three key features.




Fig. 2. A partition tree structure of Algorithm 1 for a run on the double integrator system (19) for $\varepsilon = 0.1$, $D_p = 10$, $D_N = 2$. The color represents the algorithm's *separation*—only nodes filled with red compute the neural network verification step, and all the integrations are performed only on the leaf nodes. The imbalanced structure is a consequence of the algorithm's *spatial awareness*. As a consequence of the algorithm's *temporal awareness*, the structure of the partition tree deepens as time increases.

Separation: In practice, the neural network verification step usually contributes the most computational expense to the algorithm. Introducing the maximum depths D_p and D_N allows us to separate calls to the neural network verification algorithm and the total number of partitions. This allows the algorithm to improve its accuracy without any additional calls to the neural network verifier, leaving it to the user to control the trade-off between computational complexity and accuracy. In Theorem 1, this corresponds with an improvement of the first term separately from the second term.

Spatial awareness: Based on the dynamics of the system, the over-approximation error can be highly sensitive to the spatial location of each partition. By estimating the contraction rate of each partition separately, the algorithm can localize partitioning exactly *where* necessary. In Theorem 1, this curbs exponential growth *where* c_x is the worst.

Temporal awareness: One of the main drawbacks of interval reachability frameworks is that as time evolves, the over-approximation error tends to compound exponentially. As such, by partitioning along trajectories, the algorithm can save computations to improve reachable set estimates exactly *when* necessary. In Theorem 1, this curbs exponential growth *when* c_x is the worst.

Generality of Algorithm 1: The algorithm, as written, specifically uses the Mixed Monotone framework (developed in Section V.B) for the interval reachability of the dynamical system, as well as CROWN (specified by (10)) for the neural network verification. We will exclusively use this setting for the rest of the paper, and refer it as “ReachMM-CG”. However, it is easy to see that one can replace Lines 11, 16, and 23 with any other interval reachability setting satisfying Assumptions 1 and 2.

VI. CONTRACTION-BASED GUARANTEES ON REACHABLE SET OVER-APPROXIMATION

In this section, we use contraction theory to provide rigorous bounds on the accuracy of the mixed-monotone reachable set over-approximations. For the simplicity of analysis in this section, we remove the assumption that the

control is applied in the piecewise constant fashion and instead consider a continuously applied neural network.

Theorem 1 (Accuracy guarantees for mixed-monotone reachability). *Consider the closed-loop system (6) with the neural network controller $u = N(x)$ given by (5) and let $t \mapsto x(t)$ be a trajectory of the closed-loop system (6). Suppose that d is the decomposition function for the open-loop system (4), there is a neural network verification algorithm which provides the bounds of the form (8). Let $t \mapsto \begin{bmatrix} \underline{x}(t) \\ \bar{x}(t) \end{bmatrix}$ be a trajectory of the closed-loop embedding system (14) and $t \mapsto \begin{bmatrix} \underline{y}(t) \\ \bar{y}(t) \end{bmatrix}$ be a curve such that $[x(t), \bar{x}(t)] \subseteq [y(t), \bar{y}(t)]$, for every $t \in \mathbb{R}_{\geq 0}$. Then, with $t_0 = 0$,*

$$\begin{aligned} & \left\| \begin{bmatrix} \underline{x}(t) \\ \bar{x}(t) \end{bmatrix} - \begin{bmatrix} x(t) \\ x(t) \end{bmatrix} \right\|_{\infty} \leq e^{c_x t} \left\| \begin{bmatrix} \underline{x}(0) \\ \bar{x}(0) \end{bmatrix} - \begin{bmatrix} x(0) \\ x(0) \end{bmatrix} \right\|_{\infty} \\ & + \frac{\ell_u^o (e^{c_x t} - 1)}{c_x} \sup_{\substack{\tau \in [0, t], \\ z \in [\underline{y}(\tau), \bar{y}(\tau)]}} \left\| \begin{bmatrix} \underline{N}_{[\underline{y}, \bar{y}]}(z, z) \\ \bar{N}_{[\underline{y}, \bar{y}]}(z, z) \end{bmatrix} - \begin{bmatrix} N(z) \\ N(z) \end{bmatrix} \right\|_{\infty} \\ & + \frac{\ell_w^o (e^{c_x t} - 1)}{c_x} \sup_{\tau \in [0, t]} \left\| \begin{bmatrix} \underline{w}(\tau) \\ \bar{w}(\tau) \end{bmatrix} - \begin{bmatrix} w(\tau) \\ w(\tau) \end{bmatrix} \right\|_{\infty} \quad (17) \end{aligned}$$

where

$$\begin{aligned} c_x &= \sup_{z, \bar{z} \in \Omega_t} \mu_{\infty}(D_{[\underline{z}, \bar{z}]}^c E_{[\underline{y}(t), \bar{y}(t)]}^c(z, \bar{z}, \underline{w}, \bar{w})), \\ \ell_u^o &= \sup_{\substack{x \in \Omega_t, w \in [\underline{w}, \bar{w}], \\ \underline{u}, \bar{u} \in [\underline{N}(x, x), \bar{N}(x, x)]}} \|D_{[\underline{u}, \bar{u}]}^c E^o(x, x, \underline{u}, \bar{u}, w, w)\|_{\infty}, \\ \ell_w^o &= \sup_{\substack{x \in \Omega_t, z, \bar{z} \in [\underline{w}, \bar{w}], \\ \underline{u}, \bar{u} \in [\underline{N}(x, x), \bar{N}(x, x)]}} \|D_{[\underline{z}, \bar{z}]}^c E^o(x, x, \underline{u}, \bar{u}, z, \bar{z})\|_{\infty}, \end{aligned}$$

and the set Ω_t is defined by $\Omega_t = \bigcup_{\tau \in [0, t]} [y(\tau), \bar{y}(\tau)] \subseteq \mathbb{R}^n$.

Proof: We first make the following conventions:

$$\begin{aligned} r(t) &= \begin{bmatrix} \underline{x}(t) \\ \bar{x}(t) \end{bmatrix} - \begin{bmatrix} x(t) \\ x(t) \end{bmatrix}, \\ \textcircled{1} &= E_{[\underline{y}, \bar{y}]}^c(\underline{x}, \bar{x}, \underline{w}, \bar{w}) - E_{[\underline{y}, \bar{y}]}^c(x, x, \underline{w}, \bar{w}), \\ \textcircled{2} &= E_{[\underline{y}, \bar{y}]}^c(x, x, \underline{w}, \bar{w}) - E_{[\underline{y}, \bar{y}]}^c(x, x, w, w), \\ \textcircled{3} &= E_{[\underline{y}, \bar{y}]}^c(x, x, w, w) - \begin{bmatrix} f^c(x, w) \\ f^c(x, w) \end{bmatrix}. \end{aligned}$$

Therefore, we can compute

$$\begin{aligned} \frac{1}{2} D^+ \|r(t)\|_{\infty}^2 &= \left[E_{[\underline{y}, \bar{y}]}^c(\underline{x}, \bar{x}, \underline{w}, \bar{w}) - \begin{bmatrix} f^c(x, w) \\ f^c(x, w) \end{bmatrix}, r(t) \right]_{\infty} \\ &\leq \llbracket \textcircled{1}, r(t) \rrbracket_{\infty} + \llbracket \textcircled{2}, r(t) \rrbracket_{\infty} + \llbracket \textcircled{3}, r(t) \rrbracket_{\infty}, \end{aligned}$$

where D^+ is the upper Dini Derivative with respect to time, the first equality holds by the curve-norm derivative property of the weak pairing $\llbracket \cdot, \cdot \rrbracket_{\infty}$ [19, Theorem 25(ii)] and the second inequality holds by the subadditivity of weak pairing $\llbracket \cdot, \cdot \rrbracket_{\infty}$ [19, Definition 15, property (i)]. By [19, Theorem 18], the first term on the RHS of the above equation can be estimated $\llbracket \textcircled{1}, r(t) \rrbracket_{\infty} \leq c_x \|r(t)\|_{\infty}^2$. We also note that, for the closed-loop system, we have

$$\begin{aligned} E_{[\underline{y}, \bar{y}]}^c(x, x, \underline{w}, \bar{w}) &= E^o(x, x, \underline{N}_{[\underline{y}, \bar{y}]}(x, x), \bar{N}_{[\underline{y}, \bar{y}]}(x, x), \underline{w}, \bar{w}), \\ E_{[\underline{y}, \bar{y}]}^c(x, x, w, w) &= E^o(x, x, \underline{N}_{[\underline{y}, \bar{y}]}(x, x), \bar{N}_{[\underline{y}, \bar{y}]}(x, x), w, w). \end{aligned}$$

Therefore, by [19, Definition 15, property (iv)], the second term in the RHS can be estimated as

$$\| \textcircled{2}, r(t) \|_\infty \leq \ell_w^\circ \left\| \begin{bmatrix} w(t) \\ \bar{w}(t) \end{bmatrix} - \begin{bmatrix} w(t) \\ w(t) \end{bmatrix} \right\|_\infty \|r(t)\|_\infty.$$

Finally, by the definition of E^c , we have

$$\begin{aligned} \textcircled{3} &= E^o(x, x, \underline{N}_{[y, \bar{y}]}(x, x), \bar{N}_{[y, \bar{y}]}(x, x), w, w) \\ &\quad - E^o(x, x, N(x), N(x), w, w). \end{aligned}$$

Thus, by [19, Definition 15, property (iv)], the third term in the RHS can be estimated as

$$\| \textcircled{3}, r(t) \|_\infty \leq \ell_u^\circ \left\| \begin{bmatrix} \underline{N}_{[y, \bar{y}]}(x, x) \\ \bar{N}_{[y, \bar{y}]}(x, x) \end{bmatrix} - \begin{bmatrix} N(x) \\ N(x) \end{bmatrix} \right\|_\infty \|r(t)\|_\infty.$$

Therefore, gathering all the terms and dividing by $\|r(t)\|_\infty$,

$$\begin{aligned} D^+ \|r(t)\|_\infty &\leq c_x \|r(t)\|_\infty + \ell_u^\circ \left\| \begin{bmatrix} \underline{N}_{[y, \bar{y}]}(x, x) \\ \bar{N}_{[y, \bar{y}]}(x, x) \end{bmatrix} - \begin{bmatrix} N(x) \\ N(x) \end{bmatrix} \right\|_\infty \\ &\quad + \ell_w^\circ \left\| \begin{bmatrix} w(t) \\ \bar{w}(t) \end{bmatrix} - \begin{bmatrix} w(t) \\ w(t) \end{bmatrix} \right\|_\infty. \end{aligned}$$

The result then follows from the Grönwall–Bellman inequality [19, Lemma 11].

Remark 1.

- 1) *Theorem 1 explains the three key features of Algorithm 1. (i) separation allows the improvement of term 1 independent of term 2 on the RHS of (17); (ii) spatial awareness allows for cutting the partitions with the worst contraction rates c_x and lipschitz bounds $\ell_u^\circ, \ell_w^\circ$, since each partition defines a different set Ω_t ; (iii) temporal awareness allows for improving the exponential growth present in $e^{c_x t}$. The hyper-parameters D_p, D_N of Algorithm 1 control the first and second terms respectively of (17).*
- 2) *The constant c_x in Theorem 1 is the contraction rate of the closed-loop embedding system (14). Theorem 1 shows that c_x is also the rate of contraction of the box over-approximation of reachable sets using mixed monotonicity. For $c_x < 0$, the upper bound shows that the error associated with uncertainties in initial conditions converges to 0 as $t \rightarrow \infty$.*

Finding the value of the contraction parameter c_x in Theorem 1 is usually complicated due to the large number of parameters of the neural network and its interconnection with the nonlinear dynamics of the system. In the next theorem, we provide a simple method for over-approximating c_x using the input-output Lipschitz bounds of the neural network.

Theorem 2 (Upper bound on closed-loop contraction rate). *Consider the closed-loop embedding system (14) with the neural network controller $u = N(x)$ given by (5). Suppose that, for the neural network verification in (8), we have*

$$\sup_{x \in \Omega_t} \left\| D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} \begin{bmatrix} \underline{N}_{[y, \bar{y}]} \\ \bar{N}_{[y, \bar{y}]} \end{bmatrix} \right\|_\infty = \text{Lip}_\infty$$

and, for the open-loop embedding system (9), we have

$$\begin{aligned} c_x^\circ &= \sup_{\substack{x, \bar{x} \in \Omega_t, w \in [\underline{w}, \bar{w}], \\ \underline{u}, \bar{u} \in [\underline{N}(x, \bar{x}), \bar{N}(x, \bar{x})]}} \left(\mu_\infty \left(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E^o(x, \bar{x}, \underline{u}, \bar{u}, w, \bar{w}) \right) \right) \\ \ell_u^\circ &= \sup_{\substack{x, \bar{x} \in \Omega_t, w \in [\underline{w}, \bar{w}], \\ \underline{u}, \bar{u} \in [\underline{N}(x, \bar{x}), \bar{N}(x, \bar{x})]}} \left\| D_{\begin{bmatrix} \underline{u} \\ \bar{u} \end{bmatrix}} E^o(x, \bar{x}, \underline{u}, \bar{u}, w, \bar{w}) \right\|_\infty \end{aligned}$$

where the set Ω_t is defined by $\Omega_t = \bigcup_{\tau \in [0, t]} [y(\tau), \bar{y}(\tau)] \subseteq \mathbb{R}^n$. Then $\mu_\infty(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E_{[y, \bar{y}]}^c) \leq c_x^\circ + \ell_u^\circ \text{Lip}_\infty$.

Proof: Note that, for every $\underline{x} \leq \bar{x}$ and every $\underline{w} \leq \bar{w}$,

$$E_{[y, \bar{y}]}^c(x, \bar{x}, \underline{w}, \bar{w}) = E^o(x, \bar{x}, \underline{N}_{[y, \bar{y}]}(x, \bar{x}), \bar{N}_{[y, \bar{y}]}(x, \bar{x}), \underline{w}, \bar{w}).$$

Therefore, using the chain rule, we have

$$\begin{aligned} \mu_\infty \left(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E_{[y, \bar{y}]}^c \right) &= \mu_\infty \left(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E^o + D_{\begin{bmatrix} \underline{u} \\ \bar{u} \end{bmatrix}} E^o D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} \begin{bmatrix} \underline{N}_{[y, \bar{y}]} \\ \bar{N}_{[y, \bar{y}]} \end{bmatrix} \right) \\ &\leq \mu_\infty \left(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E^o \right) + \mu_\infty \left(D_{\begin{bmatrix} \underline{u} \\ \bar{u} \end{bmatrix}} E^o \ D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} \begin{bmatrix} \underline{N}_{[y, \bar{y}]} \\ \bar{N}_{[y, \bar{y}]} \end{bmatrix} \right) \\ &\leq \mu_\infty \left(D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} E^o \right) + \left\| D_{\begin{bmatrix} \underline{u} \\ \bar{u} \end{bmatrix}} E^o \right\|_\infty \left\| D_{\begin{bmatrix} x \\ \bar{x} \end{bmatrix}} \begin{bmatrix} \underline{N}_{[y, \bar{y}]} \\ \bar{N}_{[y, \bar{y}]} \end{bmatrix} \right\|_\infty \\ &\leq c_x^\circ + \ell_u^\circ \text{Lip}_\infty. \end{aligned}$$

where the second inequality holds by the subadditive property of matrix measure $\mu_\infty(\cdot)$ [24, Lemma 2.10] and the third inequality holds using the fact that $\mu_\infty(AB) \leq \|AB\|_\infty \leq \|A\|_\infty \|B\|_\infty$, for every $A \in \mathbb{R}^{2n \times 2p}$ and $B \in \mathbb{R}^{2p \times 2n}$.

VII. NUMERICAL EXPERIMENTS

A. Nonlinear vehicle

Consider the nonlinear dynamics of a vehicle from [25]:

$$\begin{aligned} \dot{p}_x &= v \cos(\phi + \beta(u_2)) & \dot{\phi} &= \frac{v}{\ell_r} \sin(\beta(u_2)) \\ \dot{p}_y &= v \sin(\phi + \beta(u_2)) & \dot{v} &= u_1 \end{aligned} \quad (18)$$

where $[p_x, p_y]^\top \in \mathbb{R}^2$ is the displacement of the center of mass, $\phi \in [-\pi, \pi)$ is the heading angle in the plane, and $v \in \mathbb{R}_{\geq 0}$ is the speed of the center of mass. Control input u_1 is the applied force, input u_2 is the angle of the front wheels, and $\beta(u_2) = \arctan\left(\frac{\ell_f}{\ell_f + \ell_r} \tan(u_2)\right)$ is the slip slide angle. Let $x = [p_x, p_y, \phi, v]^\top$ and $u = [u_1, u_2]^\top$. We use the open-loop decomposition function and the neural network controller ($4 \times 100 \times 100 \times 2$ ReLU) defined in [14]. The neural network controller is applied at evenly spaced intervals of 0.25 seconds apart. This neural network is trained to mimic an MPC that stabilizes the vehicle to the origin while avoiding a circular obstacle centered at (4, 4) with a radius of 2.

ReachMM-CG (presented in Algorithm 1) is used to provide over-approximations of the reachable sets for $t \in [0, 1.25]$ as pictured in Figure 3. We consider the initial set of $[7.9, 8.1]^2 \times [-\frac{2\pi}{3} - 0.01, -\frac{2\pi}{3} + 0.01] \times [1.99, 2.01]$. All integrations are performed using Euler integration with a step size of 0.01, and CROWN is computed using

All code is available at https://github.com/gtfacts/ReachMM_CDC2023




Fig. 3. The over-approximated reachable sets of the closed-loop nonlinear vehicle model (18) in the (p_x, p_y) coordinates are shown in blue for the initial set $[7.9, 8.1]^2 \times [-\frac{2\pi}{3} - 0.01, -\frac{2\pi}{3} + 0.01] \times [1.99, 2.01]$ over the time interval $[0, 1.25]$. They are computed using Algorithm 1 with $\varepsilon = [0.2, 0.2, \infty, \infty]^T$, $D_p = 2$, and $D_N = 1$. The average runtime across 100 runs with standard deviation is reported, as well as the volume of the over-bounding box at the final time $T = 1.25$. 200 true trajectories of the system are shown in the time-varying yellow line.

auto-LiRPA [26]. In Table I, the performance of ReachMM-CG with different selections of hyper-parameters (all with check contraction factor $\gamma = 0.1$) is compared to ReachMM, the non-adaptive purely initial set partitioning strategy from [14]. The runtimes are averaged over 100 runs, with mean and standard deviation reported. The volumes for the tightest over-bounding box for the approximated reachable set at the final time step is reported. The adaptive strategies see an improved run-time over the non-adaptive counterparts, and in some cases see a tighter approximation.

ε	D_p, D_N	Runtime (s)	Volume
non-adaptive	(2, 1)	1.851 ± 0.010	1.988
$[0.2, 0.2, \infty, \infty]$	(2, 1)	1.583 ± 0.010	1.689
$[0.25, 0.25, \infty, \infty]$	(2, 1)	1.243 ± 0.008	1.846
non-adaptive	(2, 2)	4.274 ± 0.023	0.803
$[0.2, 0.2, \infty, \infty]$	(2, 2)	3.332 ± 0.012	0.787
$[0.25, 0.25, \infty, \infty]$	(2, 2)	2.636 ± 0.008	0.986

TABLE I

THE PERFORMANCE OF REACHMM-CG ON THE VEHICLE MODEL.

B. Linear Discrete-Time Double Integrator

Consider a zero-order hold discretization of the classical double integrator with a step-size of 1 (from [7]):

$$x_{t+1} = \underbrace{\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}}_A x_t + \underbrace{\begin{bmatrix} 0.5 \\ 1 \end{bmatrix}}_B u_t \quad (19)$$

Here, we consider a fixed actuation step size of 1—the same as the integration step. In this special case (discrete-time LTI systems), [14, Corollary 4] shows that the following is a valid closed-loop embedding system that provides tighter bounds than (14):

$$\begin{bmatrix} \underline{x}_{t+1} \\ \bar{x}_{t+1} \end{bmatrix} = \begin{bmatrix} M^+ & M^- \\ M^- & M^+ \end{bmatrix} \begin{bmatrix} x_t \\ \bar{x}_t \end{bmatrix} + \begin{bmatrix} B^+ & B^- \\ B^- & B^+ \end{bmatrix} \begin{bmatrix} d_t \\ \bar{d}_t \end{bmatrix}$$

with $\underline{C}, \bar{C}, d, \bar{d}$ taken from CROWN as (10) on (x_t, \bar{x}_t) , and $M = A + B^+C + B^-C$ and $\bar{M} = A + B^+C + B^-C$.

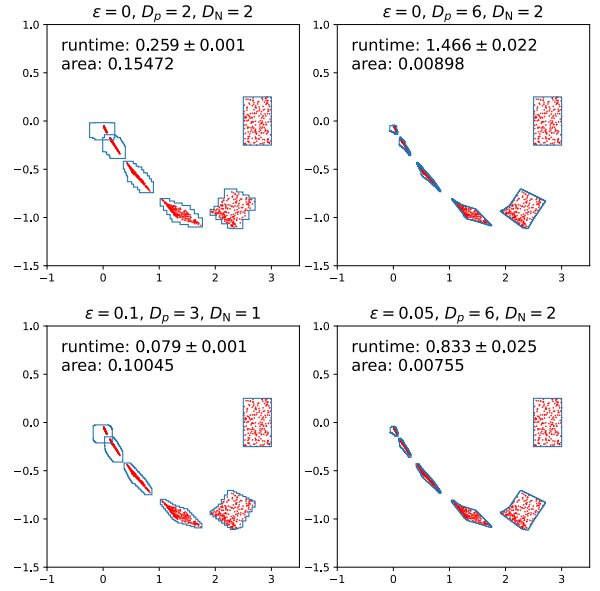


Fig. 4. The over-approximated reachable sets of the closed-loop double integrator model (19) are shown in blue for the initial set $[2.5, 3] \times [-0.25, 0.25]$ and final time $T = 5$. They are computed using Algorithm 1 with the specified parameters in the title of each plot. The average runtime and standard deviation across 100 runs is reported, as well as the true area of the final reachable set. 200 true trajectories are shown in red. The horizontal axis is x_1 and the vertical axis is x_2 .

We use the neural network from [7] ($2 \times 10 \times 5 \times 1$, ReLU). The performance of Algorithm 1 for various choices of ε , D_p , and D_N is shown in Figure 4, for the initial set $[2.5, 3] \times [-0.25, 0.25]$ and a final time of $T = 5$.

Comparing with the literature: Additionally, we compare the proposed ReachMM-CG to state-of-the-art partitioning algorithms for linear discrete-time systems: ReachMM [14] (uniform initial partitioning), ReachLP-Uniform [7] (ReachLP with uniform initial partitioning), ReachLP-GSG (ReachLP with greedy sim-guided partitioning) [7], and ReachLipBnB [18] (branch-and-bound using LipSDP [5]).

Each algorithm is run with two different sets of hyper-parameters, aiming to compare their performances across various regimes. The setup for ReachMM-CG is (ε, D_p, D_N) ; ReachMM is (D_p, D_N) ; ReachLP-Uniform is $\#$ initial partitions, ReachLP-GSG is $\#$ of total propagator calls, ReachLipBnB is ε . Their runtimes are averaged over 100 runs, with mean and standard deviation reported. The true areas of the reachable sets are computed using Python packages (Shapely, polytope). The performance is outlined in Table II, and notable reachable sets are displayed in Figure 5. ReachMM-CG outperforms SOTA across the board: for both setups, ReachMM-CG is significantly faster than the other methods while returning a tighter reachable set.

VIII. CONCLUSIONS

In this paper, we propose an adaptive partitioning approach for interval reachability analysis of neural network controlled systems. The algorithm uses *separation* of the neural network verifier and the dynamical system reachability tool, is *spatially aware* in choosing the right locations of the state space

Method	Setup	Runtime (s)	Area
ReachMM-CG (our method)	(0.1, 3, 1) (0.05, 6, 2)	0.079 ± 0.001 0.833 ± 0.025	1.0 · 10⁻¹ 7.5 · 10⁻³
ReachMM	(2, 2)	0.259 ± 0.001	1.5 · 10 ⁻¹
	(6, 2)	1.466 ± 0.022	9.0 · 10 ⁻³
ReachLP-Unif	4	0.212 ± 0.002	1.5 · 10 ⁻¹
	16	3.149 ± 0.004	1.0 · 10 ⁻²
ReachLP-GSG	55	0.913 ± 0.031	5.3 · 10 ⁻¹
	205	2.164 ± 0.042	8.8 · 10 ⁻²
ReachLipBnB	0.1	0.956 ± 0.067	5.4 · 10 ⁻¹
	0.001	3.681 ± 0.100	1.2 · 10 ⁻²

TABLE II

THE PERFORMANCE OF REACHMM-CG ON THE LTI DISCRETE-TIME DOUBLE INTEGRATOR MODEL, COMPARED TO OTHER STATE OF THE ART PARTITIONING ALGORITHMS.




Fig. 5. The over-approximated reachable sets of the closed-loop double integrator model (19) are compared for three different algorithms on two different runtime regimes, for the initial set $[2.5, 3] \times [-0.25, 0.25]$ and final time $T = 5$. The experiment setup and performances are reported in Table II. 200 true trajectories are shown in red. The horizontal axis is x_1 and the vertical axis is x_2 .

to partition, and is *temporally aware* to partition along trajectories rather than merely on the initial set. Using contraction theory for mixed monotone reachability analysis, we provide formal guarantees for the algorithm’s performance. Finally, using simulations, we demonstrate significant improvement in runtime, accuracy, and generality as compared to existing partitioning approaches.

REFERENCES

- [1] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” in *International Conference on Learning Representations*, 2014.
- [2] S. Dutta, X. Chen, and S. Sankaranarayanan, “Reachability analysis for neural feedback systems using regressive polynomial rule inference,” in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, p. 157–168, 2019.
- [3] S. Gowal, K. Dvijotham, R. Stanforth, R. Bunel, C. Qin, J. Uesato, R. Arandjelovic, T. A. Mann, and P. Kohli, “Scalable verified training for provably robust image classification,” in *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4841–4850, 2019.
- [4] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, “Efficient neural network robustness certification with general activation functions,” in *Advances in Neural Information Processing Systems*, vol. 31, p. 4944–4953, 2018.
- [5] M. Fazlyab, M. Morari, and G. J. Pappas, “Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming,” *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 1–15, 2022.
- [6] H.-D. Tran, X. Yang, D. Manzanos Lopez, P. Musau, L. V. Nguyen, W. Xiang, S. Bak, and T. T. Johnson, “NNV: The neural network

- verification tool for deep neural networks and learning-enabled cyber-physical systems,” in *Computer Aided Verification*, pp. 3–17, Springer International Publishing, 2020.
- [7] M. Everett, G. Habibi, C. Sun, and J. P. How, “Reachability analysis of neural feedback loops,” *IEEE Access*, vol. 9, pp. 163938–163953, 2021.
- [8] H. Hu, M. Fazlyab, M. Morari, and G. J. Pappas, “Reach-SDP: Reachability analysis of closed-loop systems with neural network controllers via semidefinite programming,” in *59th IEEE Conference on Decision and Control*, pp. 5929–5934, 2020.
- [9] C. Schilling, M. Forets, and S. Guadalupe, “Verification of neural-network control systems by integrating Taylor models and zonotopes,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 8169–8177, 2022.
- [10] N. Kochdumper, C. Schilling, M. Althoff, and S. Bak, “Open- and closed-loop neural network verification using polynomial zonotopes,” in *NASA Formal Methods* (K. Y. Rozier and S. Chaudhuri, eds.), (Cham), pp. 16–36, Springer Nature Switzerland, 2023.
- [11] C. Huang, J. Fan, X. Chen, W. Li, and Q. Zhu, “POLAR: A polynomial arithmetic framework for verifying neural-network controlled systems,” in *Automated Technology for Verification and Analysis: 20th International Symposium (ATVA)*, pp. 414–430, Springer, 2022.
- [12] C. Sidrane, A. Maleki, A. Irfan, and M. J. Kochenderfer, “OVERT: An algorithm for safety verification of neural network control policies for nonlinear systems,” *Journal of Machine Learning Research*, vol. 23, no. 117, pp. 1–45, 2022.
- [13] Y. Zhang and X. Xu, “Safety verification of neural feedback systems based on constrained zonotopes,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 2737–2744, 2022.
- [14] S. Jafarpour, A. Harapanahalli, and S. Coogan, “Interval reachability of nonlinear dynamical systems with neural network controllers,” in *Learning for Dynamics and Control Conference*, pp. 12–25, PMLR, 2023.
- [15] V. Rubies-Royo, R. Calandra, D. M. Stipanovic, and C. Tomlin, “Fast neural network verification via shadow prices,” *arXiv preprint arXiv:1902.07247*, 2019.
- [16] M. Everett, G. Habibi, and J. P. How, “Robustness analysis of neural networks via efficient partitioning with applications in control systems,” *IEEE Control Systems Letters*, vol. 5, no. 6, pp. 2114–2119, 2021.
- [17] W. Xiang, H.-D. Tran, X. Yang, and T. T. Johnson, “Reachable set estimation for neural network control systems: A simulation-guided approach,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 5, pp. 1821–1830, 2021.
- [18] T. Entesari, S. Sharifi, and M. Fazlyab, “ReachLipBnB: A branch-and-bound method for reachability analysis of neural autonomous systems using lipschitz bounds,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1003–1010, 2023.
- [19] A. Davydov, S. Jafarpour, and F. Bullo, “Non-Euclidean contraction theory for robust nonlinear stability,” *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6667–6681, 2022.
- [20] X. Chen, E. Ábrahám, and S. Sankaranarayanan, “Flow*: An analyzer for non-linear hybrid systems,” in *Computer Aided Verification: 25th International Conference (CAV)*, pp. 258–263, Springer, 2013.
- [21] 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)*, pp. 2242–2253, 2017.
- [22] M. Althoff, “An introduction to CORA 2015,” in *Proceedings of the Workshop on Applied Verification for Continuous and Hybrid Systems*, p. 120–151, 2015.
- [23] S. Coogan, “Mixed monotonicity for reachability and safety in dynamical systems,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 5074–5085, 2020.
- [24] F. Bullo, *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.0 ed., 2022.
- [25] P. Polack, F. Alché, B. d’Andréa Novel, and A. de La Fortelle, “The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 812–818, 2017.
- [26] K. Xu, Z. Shi, H. Zhang, Y. Wang, K.-W. Chang, M. Huang, B. Kailkhura, X. Lin, and C.-J. Hsieh, “Automatic perturbation analysis for scalable certified robustness and beyond,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 1129–1141, 2020.