

Scalable Forward Reachability Analysis of Multi-Agent Systems with Neural Network Controllers

Oliver Gates, Matthew Newton and Konstantinos Gatsis

Abstract—Neural networks (NNs) have been shown to learn complex control laws successfully, often with performance advantages or decreased computational cost compared to alternative methods. Neural network controllers (NNCs) are, however, highly sensitive to disturbances and uncertainty, meaning that it can be challenging to make satisfactory robustness guarantees for systems with these controllers. This problem is exacerbated when considering multi-agent NN-controlled systems, as existing reachability methods often scale poorly for large systems. This paper addresses the problem of finding overapproximations of forward reachable sets for discrete-time uncertain multi-agent systems with distributed NNC architectures. We first reformulate the dynamics, making the system more amenable to reachability analysis. Next, we take advantage of the distributed architecture to split the overall reachability problem into smaller problems, significantly reducing computation time. We use quadratic constraints, along with a convex representation of uncertainty in each agent’s model, to form semidefinite programs, the solutions of which give overapproximations of forward reachable sets for each agent. Finally, the methodology is tested on two realistic examples: a platoon of vehicles and a power network system.

I. INTRODUCTION

There has been recent interest in the use of neural networks (NNs) for control in closed-loop feedback systems. Neural network controllers (NNCs) can be used to imitate traditional control policies, such as model predictive control (MPC), with reduced computational cost [1], or to implement deep reinforcement learning (RL) policies [2]. Even for simple linear systems, NNCs can be used to implement complex non-linear control laws (which may not be easy to achieve with existing methods). NNs are, however, highly sensitive to input perturbations, so disturbances in the closed-loop system can have adverse effects [3]. This is problematic when NNCs are used in safety-critical systems, and recent work has focused on reachability analysis for systems with NNCs; if we can overapproximate the forward reachable sets, then it can be verified that certain regions of the state space are avoided over a given horizon.

The problem of computing forward reachable sets becomes more challenging when considering a multi-agent system in which each agent is controlled by an NN (or a series of NNs); the effects of small perturbations to the input of one NNC are propagated through the system. Control of multi-agent systems is well-studied, and common goals include consensus, formation control and flocking/swarming [4]. Multi-agent control architectures can be categorised

according to the dependence of each agent’s control input on other agents’ states: (a) centralised control, in which each agent’s control input is a function of all agents’ states; (b) distributed control, in which each agent’s control input is a function of a subset of the other agents’ states; (c) decentralised control, in which each agent’s control input is a function of only its own state [5]. In distributed control, common approaches include state feedback [6] and distributed model predictive control (DMPC) [7].

A number of methods have been proposed for the reachability analysis of systems with NNCs [8]–[16]. In general, the problem of reachability for discrete-time LTI systems is undecidable [17], so methods for safety verification of closed-loop systems often aim to find tight overapproximations of the forward reachable sets. Alternatively, additional restrictions can be placed on the problem to allow the exact reachable sets to be computed. Generally, there is a tradeoff between scalability and tightness of the bounds [3].

In [8], semidefinite programming (SDP) is used to compute overapproximations of the forward reachable sets, and [9] builds on this work by considering parameter-varying systems. In [10] and [11], the input set is partitioned into smaller sets, and a linear programming (LP) approach is used to overapproximate the reachable sets; in [10], the input set partitioning approach is also applied to the method in [8], and a comparison is made between the LP and SDP approaches, demonstrating that the former is faster but the latter results in tighter bounds. The work in [12] restricts the input sets to be constrained zonotopes, allowing for the exact computation of reachable sets, and the work in [13] represents the input sets as hybrid zonotopes, allowing for a class of non-convex input sets. In [14], polynomial zonotopes are used to abstract the closed-loop dynamics, providing tight overapproximations. Other approaches include the use of polynomial optimisation [15] and Bernstein polynomials [16].

Existing reachability methods for NN-controlled systems do not explicitly consider multi-agent systems with distributed control architectures. Similarly to the single-agent case, in which NNCs can be used to learn complex control laws, a distributed NNC architecture can be used to learn complex distributed control laws [18]. An example is DMPC, in which each agent’s controller solves an optimisation problem based on its own state and those of its neighbours. An overview of DMPC is given in [7], and its applications include vehicle platooning [19], frequency regulation in power systems [20] and formation control of UAVs [21].

In this paper, we present a scalable method to compute overapproximations of the forward reachable sets for uncer-

The authors are with the Department of Engineering Science, University of Oxford, UK. Email: oliver.gates@st-hildas.ox.ac.uk, matthew.newton@worc.ox.ac.uk, konstantinos.gatsis@eng.ox.ac.uk

tain multi-agent systems with distributed NNC architectures. To the best of our knowledge, this is the first work which explicitly deals with the multi-agent NNC reachability problem. The main contributions are as follows:

- we reformulate the dynamics, making the system more amenable to reachability analysis;
- we take advantage of the distributed architecture to split the overall reachability problem into smaller problems, using an existing SDP-based approach to overapproximate the forward reachable sets for each agent, and further extend this result to incorporate model uncertainty in the agents' dynamics;
- we demonstrate the effectiveness of this methodology on two realistic multi-agent systems with different structures: a platoon of vehicles and a power network;
- we compare our approach to the approach of treating the multi-agent system as one overall system, and we demonstrate that our approach outperforms the alternative approach.

In Section II, we describe the dynamics, give some remarks about the form of the controller and describe the forward reachability problem. In Section III, we provide a simplification of the dynamics and control input, and in Section IV, we present the reachability method. In Section V, we introduce model uncertainty. In Section VI, we present experiments to demonstrate the method.

A. Notation

The set of real $n \times m$ matrices is denoted by $\mathbb{R}^{n \times m}$, the set of real n -length vectors by \mathbb{R}^n and the set of real numbers by \mathbb{R} . The set of symmetric $n \times n$ matrices is denoted by \mathbb{S}^n . The set of diagonal $n \times n$ matrices is denoted by \mathbb{D}^n . The set of positive integers is denoted by \mathbb{Z}^+ . The cardinality of a set \mathcal{S} is denoted by $|\mathcal{S}|$. The $n \times n$ identity matrix is denoted by I_n . The symbols \geq and \leq apply elementwise to vectors and matrices. $A \preceq 0$ implies that matrix A is negative semidefinite. The number 0 is used to represent the scalar, vector or matrix of appropriate size; the size will be clear from the context.

II. PROBLEM STATEMENT

A. Multi-agent dynamics

We consider a discrete-time multi-agent system of M agents, where each agent i has linear time-invariant dynamics

$$x_{k+1}^{[i]} = A_{ii}x_k^{[i]} + \sum_{j \in \mathcal{N}_i} A_{ij}x_k^{[j]} + B_i u_k^{[i]} + w_k^{[i]}, \quad (1)$$

where for each agent $i \in \{1, \dots, M\} = \mathcal{I}$, $x_k^{[i]} \in \mathbb{R}^{n_x}$ is the local state, $x_k^{[j]} \in \mathbb{R}^{n_x}$ is the j^{th} neighbouring state, $u_k^{[i]} \in \mathbb{R}^{n_u}$ is the control input, $w_k^{[i]} \in \mathbb{R}^{n_x}$ is a known external input, $A_{ii} \in \mathbb{R}^{n_x \times n_x}$ is the state matrix, $A_{ij} \in \mathbb{R}^{n_x \times n_x}$ is the matrix describing the effect of state $x_k^{[j]}$ on agent i , $B_i \in \mathbb{R}^{n_x \times n_u}$ is the input matrix, and \mathcal{N}_i is the set of neighbours. Note that for simplicity, we have assumed that all agents have the same dimensions (although this assumption could be relaxed). We also assume that $|\mathcal{N}_i| > 0 \forall i$.

B. Control input

In a traditional distributed control scheme with proportional feedback, the control input $u_{\text{trad},k}^{[i]}$ might be given by

$$u_{\text{trad},k}^{[i]} = \sum_{j \in \mathcal{N}_i} K_{ij} \left(x_k^{[j]} - x_k^{[i]} \right),$$

where $K_{ij} \in \mathbb{R}^{n_u \times n_x}$ is some gain matrix (which could represent multiple gain matrices) [22]. Similarly, in a DMPC scheme, the i^{th} control input is generated by solving an optimisation problem based on the agent's state $x_k^{[i]}$ and the neighbours' states $x_k^{[j]} \forall j \in \mathcal{N}_i$. In this paper, we consider the extension of the traditional distributed control schemes to NN-based control, in which the i^{th} control input is generated by some non-linear function of the agent's state $x_k^{[i]}$ and the neighbours' states $x_k^{[j]} \forall j \in \mathcal{N}_i$. We also consider the possibility of controller saturation. Hence, the i^{th} control input $u_k^{[i]}$ is given by

$$u_k^{[i]} = \text{sat}_{\mathcal{U}_i} \left[\sum_{j \in \mathcal{N}_i} \pi_{ij} \left(\begin{bmatrix} x_k^{[i]} \\ x_k^{[j]} \end{bmatrix} \right) \right], \quad (2)$$

where $\text{sat}_{\mathcal{U}_i}$ is a projection onto the set $\mathcal{U}_i = \{u \in \mathbb{R}^{n_u} \mid \underline{u}_i \leq u \leq \bar{u}_i\}$, where \underline{u}_i and \bar{u}_i are the lower and upper limits, respectively, for the i^{th} controller, and $\pi_{ij} : \mathbb{R}^{2n_x} \rightarrow \mathbb{R}^{n_u}$ is a function representing the mapping of the input through a multi-layer perceptron (MLP).

C. Multi-layer perceptron

The mapping $s \mapsto \pi_{ij}(s)$ for an L -layer MLP is

$$z_{ij}^0 = s, \quad (3a)$$

$$z_{ij}^{\ell+1} = \sigma^\ell (W_{ij}^\ell z_{ij}^\ell + b_{ij}^\ell), \quad \ell = 0, \dots, L-1, \quad (3b)$$

$$\pi_{ij}(s) = W_{ij}^L z_{ij}^L + b_{ij}^L, \quad (3c)$$

where $z_{ij}^\ell \in \mathbb{R}^{n_\ell}$ is the ℓ^{th} vector of activation values (note that $n_0 = 2n_x$), $W_{ij}^\ell \in \mathbb{R}^{n_{\ell+1} \times n_\ell}$ is the ℓ^{th} weight matrix, $b_{ij}^\ell \in \mathbb{R}^{n_{\ell+1}}$ is the ℓ^{th} bias vector, and $\sigma^\ell : \mathbb{R}^{n_{\ell+1}} \rightarrow \mathbb{R}^{n_{\ell+1}}$ is the ℓ^{th} ReLU activation function, i.e. $\sigma^\ell(y) = \max(y, 0)$, applied elementwise, such that $\sigma^\ell(y) = [\max(y_1, 0) \ \dots \ \max(y_{n_{\ell+1}}, 0)]^\top$, where $y = [y_1 \ \dots \ y_{n_{\ell+1}}]^\top$ is the vector of pre-activation values. For simplicity, we assume that all MLPs have the same structure (size and number of hidden layers) for all i, j .

D. Overapproximation of forward reachable sets

We denote the set of all possible states of the i^{th} agent at time k as $\mathcal{X}_k^{[i]}$, such that $x_k^{[i]} \in \mathcal{X}_k^{[i]}$. Given $\mathcal{X}_k^{[i]}$ and the sets of neighbouring states at time k , i.e. $\mathcal{X}_k^{[j]} \forall j \in \mathcal{N}_i$, we aim to find an overapproximation $\widehat{\mathcal{X}}_{k+1}^{[i]}$ of the reachable set $\mathcal{X}_{k+1}^{[i]}$ at the next time step. Specifically, we aim to find the tightest polytopic overapproximation of the reachable set

$$\min \text{vol} \left(\widehat{\mathcal{X}}_{k+1}^{[i]} \right) \quad (4a)$$

$$\text{subject to } \widehat{\mathcal{X}}_{k+1}^{[i]} \supseteq \mathcal{X}_{k+1}^{[i]}, \quad (4b)$$

$$\widehat{\mathcal{X}}_{k+1}^{[i]} \text{ is a polytope,} \quad (4c)$$

for each agent i , where vol is the n_x -dimensional volume. This could then be applied recursively to overapproximate the next forward reachable sets $\mathcal{X}_{k+2}^{[i]}, \dots, \mathcal{X}_{k+N}^{[i]}$ over a finite horizon N , and used to verify that certain unsafe regions of the state space are avoided over this horizon.

III. REFORMULATION OF DYNAMICS

An obvious approach to the forward reachability problem is to augment the agents' states into one 'overall' state

$$x_k = \begin{bmatrix} x_k^{[1]\top} & \dots & x_k^{[M]\top} \end{bmatrix}^\top,$$

then form a recursion for the overall dynamics and use existing methods to perform reachability analysis on this system. The main issue with this approach is that it ignores the distributed architecture of the system – the dynamics of each agent depend only on its neighbours, not on all other agents. As a result, we reformulate the dynamics given by (1) and (2) to allow reachability analysis to be performed for each agent. In Section VI, we show that there is a significant computational advantage to solving M smaller reachability problems over one large reachability problem for the SDP-based method. We note that the approach of decomposing the reachability problem into smaller problems has been used more generally in work on reachability analysis [23], [24].

A. Simplification of dynamics

Let $g(i, j)$ be a function which returns the j^{th} neighbour of the i^{th} agent. For example, if agent 3 has $\mathcal{N}_3 = \{2, 5, 6\}$, then $g(3, 1) = 2$, $g(3, 2) = 5$ and $g(3, 3) = 6$. Then we can write $\mathcal{N}_i = \{g(i, 1), \dots, g(i, q_i)\}$, where $q_i = |\mathcal{N}_i|$, and let

$$\tilde{A}_i = \begin{bmatrix} A_{ii} & A_{ig(i,1)} & \dots & A_{ig(i,q_i)} \end{bmatrix}, \quad (5)$$

$$\tilde{x}_k^{[i]} = \begin{bmatrix} x_k^{[i]\top} & x_k^{[g(i,1)]\top} & \dots & x_k^{[g(i,q_i)]\top} \end{bmatrix}^\top, \quad (6)$$

such that $\tilde{x}_k^{[i]}$ is the concatenation of the state of agent i and the states of the neighbours of agent i at time k , then (1) can be written as

$$x_{k+1}^{[i]} = \tilde{A}_i \tilde{x}_k^{[i]} + B_i u_k^{[i]} + w_k^{[i]}.$$

B. Simplification of control input

To simplify (2), we aim to represent it as the mapping of $\tilde{x}_k^{[i]}$ through a single MLP $\Pi_i : \mathbb{R}^{(1+q_i)n_x} \rightarrow \mathbb{R}^{n_u}$. Note that the layer sizes can differ from agent to agent, depending on $|\mathcal{N}_i|$. The mapping from $\tilde{x}_k^{[i]}$ to $\Pi_i(\tilde{x}_k^{[i]})$ is then

$$z_{i,k}^0 = \tilde{x}_k^{[i]}, \quad (7a)$$

$$z_{i,k}^{\ell+1} = \sigma_i^\ell (W_i^\ell z_{i,k}^\ell + b_i^\ell), \quad \ell = 0, \dots, L-1, \quad (7b)$$

$$\Pi_i(\tilde{x}_k^{[i]}) = W_i^L z_{i,k}^L + b_i^L, \quad (7c)$$

where $z_{i,k}^\ell \in \mathbb{R}^{\tilde{n}_i^\ell}$ is the ℓ^{th} vector of activation values (note that $\tilde{n}_0^i = (1 + q_i)n_x$), $\sigma_i^\ell : \mathbb{R}^{\tilde{n}_i^{\ell+1}} \rightarrow \mathbb{R}^{\tilde{n}_i^{\ell+1}}$ is the ℓ^{th} ReLU activation function, and the weight matrices and bias vectors

are given by $W_i^0 = \mathfrak{W}_i^0 \Lambda_i$, $b_i^0 = \mathfrak{b}_i^0$, $W_i^\ell = \mathfrak{W}_i^\ell$, $b_i^\ell = \mathfrak{b}_i^\ell$ for $\ell = 1, \dots, L-1$, and $W_i^L = \Omega_i \mathfrak{W}_i^L$, $b_i^L = \Omega_i \mathfrak{b}_i^L$, where

$$\mathfrak{W}_i^\ell = \text{blkdiag} \left(W_{ig(i,1)}^\ell, \dots, W_{ig(i,q_i)}^\ell \right),$$

$$\mathfrak{b}_i^\ell = \begin{bmatrix} b_{ig(i,1)}^\ell \top & \dots & b_{ig(i,q_i)}^\ell \top \end{bmatrix}^\top,$$

for $\ell = 0, \dots, L$, where $\Lambda_i \in \mathbb{R}^{2q_i n_x \times \tilde{n}_0^i}$ and $\Omega_i \in \mathbb{R}^{n_u \times q_i n_u}$ are given by

$$\Lambda_i = \begin{bmatrix} I_{n_x} & 0 & 0 & 0 & \dots & 0 \\ 0 & I_{n_x} & 0 & 0 & \dots & 0 \\ I_{n_x} & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & I_{n_x} & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ I_{n_x} & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & \dots & I_{n_x} \end{bmatrix}, \quad \Omega_i = \begin{bmatrix} I_{n_u} \\ \vdots \\ I_{n_u} \end{bmatrix}^\top.$$

Finally, we can write

$$\sum_{j \in \mathcal{N}_i} \pi_{ij} \left(\begin{bmatrix} x_k^{[i]} \\ x_k^{[j]} \end{bmatrix} \right) \equiv \Pi_i(\tilde{x}_k^{[i]}).$$

C. Summary of reformulation

Lemma 1: The dynamics given by (1) and (2) can be written as

$$x_{k+1}^{[i]} = \tilde{A}_i \tilde{x}_k^{[i]} + B_i \text{sat}_{\mathcal{U}_i} \left[\Pi_i(\tilde{x}_k^{[i]}) \right] + w_k^{[i]}, \quad (8)$$

where \tilde{A}_i and $\tilde{x}_k^{[i]}$ are defined in (5) and (6), respectively, and $\Pi_i(\tilde{x}_k^{[i]})$ is given by (7a)–(7c).

IV. FORWARD REACHABILITY ANALYSIS

To overapproximate the forward reachable sets, we extend the reachability method outlined in [3] and [8] for the multi-agent case. The method in [8], called Reach-SDP, uses quadratic constraints (QCs) and semidefinite programming; the input set, NN and reachable set are abstracted using QCs. This involves forming quadratic inequalities for sets and functions by pre- and post-multiplying a matrix by a 'basis vector'. Using 'change-of-basis' matrices allows the same basis vector to be used for all inequalities, which allows linear matrix inequalities (LMIs) to be formed for forward reachability analysis of the closed-loop system.

A. Incorporation of control limits into MLP

We first add two layers to the MLP in (7a)–(7c) to account for the control limits [8]. The mapping from $\tilde{x}_k^{[i]}$ to $u_k^{[i]}$ is now defined by (7a), (7b) and

$$z_{i,k}^{L+1} = \sigma_i^L (W_i^L z_{i,k}^L + b_i^L - \underline{u}_i), \quad (9a)$$

$$z_{i,k}^{L+2} = \sigma_i^{L+1} (-z_{i,k}^{L+1} + \bar{u}_i - \underline{u}_i), \quad (9b)$$

$$u_k^{[i]} = -z_{i,k}^{L+2} + \bar{u}_i. \quad (9c)$$

We then define the overall basis vector for the QCs as $[z_{i,k}^\top \ 1]^\top$, where $z_{i,k}^\top = [z_{i,k}^0 \top \ z_{i,k}^1 \top \ \dots \ z_{i,k}^{L+2} \top]^\top$.

B. QCs for input set, MLP and reachable set

We present the following Lemmas for the input set, MLP and reachable set. These are slightly modified from those originally presented in [3] and [8] for the multi-agent case, and the proofs are given in [25].

Lemma 2: Consider the state of agent i and the states of the neighbours of agent i at time k . If these are bounded by a hyper-rectangular input set, i.e.

$$\tilde{\mathcal{X}}_k^{[i]} = \left\{ x \in \mathbb{R}^{\bar{n}_0^i} \mid \underline{\hat{x}}_k^{[i]} \leq x \leq \bar{x}_k^{[i]} \right\}, \quad (10)$$

then

$$\begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix}^\top \Delta_k^i(\Gamma) \begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix} \geq 0,$$

where $\Delta_k^i(\Gamma)$ is defined in [25].

Lemma 3: First, we define $\hat{z}_{i,k}^{\ell+1} = W_i^\ell z_{i,k}^\ell + b_i^\ell$ for $\ell = 0, \dots, L-1$, $\hat{z}_{i,k}^{L+1} = W_i^L z_{i,k}^L + b_i^L - \underline{u}_i$, and $\hat{z}_{i,k}^{L+2} = -z_{i,k}^{L+1} + \bar{u}_i - \underline{u}_i$. Then, we can write $z_{i,k}^\ell = \sigma_i^{\ell-1}(\hat{z}_{i,k}^\ell)$ for $\ell = 1, \dots, L+2$. If the activation values for the extended NN satisfy $z_{i,k}^\ell = \sigma_i^{\ell-1}(\hat{z}_{i,k}^\ell)$ for $\ell = 1, \dots, L+2$ then

$$\begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix}^\top \Theta^i(\lambda, \nu, \eta) \begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix} \geq 0,$$

where $\Theta^i(\lambda, \nu, \eta)$ is defined in [25].

Lemma 4: If

$$\begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix}^\top \Psi_k^i(h_p) \begin{bmatrix} \mathbf{z}_{i,k} \\ 1 \end{bmatrix} \leq 0,$$

for $p = 1, \dots, m$, where $\Psi_k^i(h_p)$ is defined in [25], then $\hat{\mathcal{X}}_{k+1}^{[i]}$ is a polytope defined by H_1, \dots, H_m and h_1, \dots, h_m , as in

$$\hat{\mathcal{X}}_{k+1}^{[i]} = \left\{ x \in \mathbb{R}^{n_x} \mid H_1^\top x \leq h_1, \dots, H_m^\top x \leq h_m \right\}.$$

C. Solution to reachability problem

By combining the results in Lemmas 1–4, we arrive at the following result for the overapproximation of the forward reachable set of agent i .

Theorem 1: Consider a discrete-time LTI system of the form in (1) and (2), where the structure of each MLP is given by (3a)–(3c). At time k , let the state of agent i and the states of its neighbours be bounded by a hyper-rectangular input set, as in (10). If there exists a solution to

$$\begin{aligned} & \min_{\Gamma, \lambda, \nu, \eta, h_p} h_p \\ & \text{subject to } \Delta_k^i(\Gamma) + \Theta^i(\lambda, \nu, \eta) + \Psi_k^i(h_p) \preceq 0, \end{aligned}$$

for $p = 1, \dots, m$, where H_1, \dots, H_m are specified by the user, then the resulting polytope is the solution to (4a)–(4c).

Proof: See [25].

V. MODEL UNCERTAINTY

In this section, we introduce model uncertainty into the dynamics and extend Theorem 1 to account for this. Consider the dynamics in (1). Instead of assuming direct knowledge of the state and input matrices, we now assume that they lie in the convex hull of a finite number of matrices, i.e.

$$A_{ii} \in \text{co} \left\{ A_{ii}^1, \dots, A_{ii}^{C_i} \right\}, \quad B_i \in \text{co} \left\{ B_i^1, \dots, B_i^{D_i} \right\}, \quad (11)$$

where $C_i, D_i \in \mathbb{Z}^+$, for $i = 1, \dots, M$. Note that Ψ_k^i depends on A_{ii} and B_i , so we can write $\Psi_k^i(h_p; A_{ii}, B_i)$.

Theorem 2: Consider the formulation in Theorem 1, but with the addition of model uncertainty in the state and input matrices described by (11). If there exists a solution to

$$\begin{aligned} & \min_{\Gamma, \lambda, \nu, \eta, h_p} h_p \\ & \text{subject to } \Delta_k^i(\Gamma) + \Theta^i(\lambda, \nu, \eta) + \Psi_k^i(h_p; A_{ii}^c, B_i^d) \preceq 0, \\ & \quad \forall c \in \{1, \dots, C_i\}, d \in \{1, \dots, D_i\}, \end{aligned}$$

for $p = 1, \dots, m$, where H_1, \dots, H_m are specified by the user, then the resulting polytope is the solution to (4a)–(4c).

Proof: See [25].

VI. EXPERIMENTS

In this section, we use two realistic examples of multi-agent systems to demonstrate our results. We also give a comparison to the approach proposed at the start of Section III (performing reachability analysis on the overall system). We then introduce model uncertainty into one of the systems and analyse this case. Simulations were performed on MATLAB, and CVX with MOSEK was used to solve the semidefinite programs. The NNs were trained to approximate distributed MPC schemes with a given horizon.

A. Vehicle platooning

In the first example, we consider the example of control of a platoon of vehicles. Control of a vehicular platoon has a number of benefits, including improved safety, higher road capacity, lower emissions, and/or reduced congestion [26]–[29]. In this example, we consider the adaptive cruise control problem, in which each vehicle aims to maintain a fixed distance from the vehicle in front, whilst travelling at a given velocity. The continuous-time longitudinal dynamics of each vehicle are given by [30]

$$\dot{x}^{[i]}(t) = \bar{A}_{ii}x^{[i]}(t) + \bar{A}_{ii-1}x^{[i-1]}(t) + \bar{B}_i u^{[i]}(t),$$

(note that $\mathcal{N}_i = \{i-1\}$) where the state vector is

$$x^{[i]}(t) = [e^{[i]}(t) \quad v^{[i]}(t) \quad a^{[i]}(t)]^\top,$$

where $e^{[i]}(t)$ is the distance error between vehicle i and vehicle $i-1$, $v^{[i]}(t)$ is the velocity of the i^{th} vehicle, $a^{[i]}(t)$ is the acceleration of the i^{th} vehicle, and $u^{[i]}(t)$ is the i^{th} acceleration input. The matrices \bar{A}_{ii} , \bar{A}_{ii-1} and \bar{B}_i are defined in [25]. Note that the lead vehicle ($i=1$) has no physical neighbour, but this can be resolved by imagining a virtual vehicle [30] with state $x^{[0]} = [0 \quad \bar{v} \quad 0]^\top$, where \bar{v} is the reference velocity. The dynamics are discretised

assuming zero-order hold (ZOH) with a sample period $T = 0.1$ s, treating $x^{[i-1]}(t)$ and $u^{[i]}(t)$ as exogenous inputs, such that the discrete-time dynamics are in the form in (1), where $w_k^{[i]} = 0 \forall i \in \mathcal{I}$. The NNs have 2 hidden layers, both with 15 neurons. The i^{th} control input is

$$u_k^{[i]} = \text{sat}_{\mathcal{U}_i} \left[\pi_{i-1} \left(\begin{bmatrix} e_k^{[i]} \\ v_k^{[i-1]} - v_k^{[i]} \\ a_k^{[i-1]} - a_k^{[i]} \end{bmatrix} \right) \right],$$

where $e_k^{[i]} \equiv e^{[i]}(kT)$, $v_k^{[i]} \equiv v^{[i]}(kT)$ and $a_k^{[i]} \equiv a^{[i]}(kT)$. The forward reachable sets were computed for 5 time steps, $M = 9$ agents, initial conditions given by $\mathcal{X}_0^{[i]} = \{x \in \mathbb{R}^3 \mid \underline{x} \leq x \leq \bar{x}\} \forall i \in \mathcal{I}$, where $\underline{x}^\top = [-0.1 \ 19.95 \ -0.01]$ and $\bar{x}^\top = [0.1 \ 20.05 \ 0.01]$, and controller limits given by $\bar{u}_i = -\underline{u}_i = 5 \forall i \in \mathcal{I}$. A step change of -2 was applied to the reference velocity at $k = 0$ (from $\bar{v} = 20$ to $\bar{v} = 18$). The results are shown in Figure 1 for the first three vehicles. A comparison between the




Fig. 1. Plots of the reachable sets in red (solid) for distance error and velocity, and simulated trajectories (blue) for the vehicle platooning example; the initial set is shown in red (dashed)

computation time for this approach (Reach-SDP-MA) and the existing method (Reach-SDP) is shown in Table 1 for different values of M .

Table 1. Comparison of methods (times in s)

M	1	2	3	4	5
Reach-SDP-MA	3.85	7.04	11.60	14.17	20.82
Reach-SDP [8]	4.01	75.39	562.94	3636.62	41025.05

Next, we extend the vehicle platooning example to account for model uncertainty. Consider the case in which $A_{ii} \in \text{co}\{(1-\delta)A^0, (1+\delta)A^0\}$ and $B_i \in \text{co}\{(1-\delta)B^0, (1+\delta)B^0\}$, where $A^0 \in \mathbb{R}^{3 \times 3}$ and $B^0 \in \mathbb{R}^3$ are the nominal values of A_{ii} and B_i , respectively, and $\delta = 0.01$. The results are shown in Figure 2. Because of the uncertainty, the size of the exact reachable sets increases, so the overapproximations are larger (compared to Figure 1).

B. Power network system

For the second example, we consider automatic generation control of a power network system. Unlike the previous example, the dynamics are not identical across agents, and some agents have more than one neighbour. There is also an




Fig. 2. Plots of the reachable sets in red (solid) for distance error and velocity, and simulated trajectories (blue) for the vehicle platooning example with model uncertainty; the initial set is shown in red (dashed)

exogenous input term. This system consists of M generation areas, and the aim is to reduce the frequency deviation in each area, in spite of load changes. Common approaches to this control problem include decentralised and distributed MPC schemes [20], [31], [32]. The continuous-time dynamics of each area are given by [31], [33]

$$\dot{x}^{[i]}(t) = \bar{A}_{ii}x^{[i]}(t) + \sum_{j \in \mathcal{N}_i} \bar{A}_{ij}x^{[j]}(t) + \bar{B}_i u^{[i]}(t) + \bar{L}_i \Delta P_L^{[i]}(t),$$

where the state vector for area i is

$$x^{[i]}(t) = \begin{bmatrix} \Delta\theta^{[i]}(t) & \Delta\omega^{[i]}(t) & \Delta P_m^{[i]}(t) & \Delta P_v^{[i]}(t) \end{bmatrix}^\top,$$

where $\Delta\theta^{[i]}(t)$, $\Delta\omega^{[i]}(t)$, $\Delta P_m^{[i]}(t)$ and $\Delta P_v^{[i]}(t)$ are the deviations in rotor angle, frequency, mechanical power and steam valve position, respectively, from the nominal values [34], $u^{[i]}(t)$ is the reference power, $\Delta P_L^{[i]}(t)$ is the local power load, and \bar{A}_{ii} , \bar{A}_{ij} , \bar{B}_i and \bar{L}_i are defined in [25]. In this example, we consider $M = 4$ generation areas (Scenario 1 in [34]), where $\mathcal{N}_1 = \{2\}$, $\mathcal{N}_2 = \{1, 3\}$, $\mathcal{N}_3 = \{2, 4\}$ and $\mathcal{N}_4 = \{3\}$. The dynamics are discretised assuming ZOH with a sample period $T = 1$ s, treating $x^{[j]}(t) \forall j \in \mathcal{N}_i$, $u^{[i]}(t)$ and $\Delta P_L^{[i]}(t)$ as exogenous inputs, such that the discrete-time dynamics are in the form in (1). The NNs have 2 hidden layers, both with 10 neurons. The i^{th} control input is

$$u_k^{[i]} = \text{sat}_{\mathcal{U}_i} \left[\sum_{j \in \mathcal{N}_i} \pi_{ij} \left(\begin{bmatrix} x_k^{[i]} \\ x_k^{[j]} \end{bmatrix} - \begin{bmatrix} x_{\text{ref},k}^{[i]} \\ x_{\text{ref},k}^{[j]} \end{bmatrix} \right) \right],$$

where $x_{\text{ref},k}^{[i]}$ and $x_{\text{ref},k}^{[j]}$ are the state and neighbour reference values, respectively. These are incorporated into the reachability analysis by modifying the first weight matrices and bias vectors. The forward reachable sets were computed for 3 time steps with initial conditions given by $\mathcal{X}_0^{[i]} = \{x \in \mathbb{R}^4 \mid \underline{x} \leq x \leq \bar{x}\} \forall i \in \mathcal{I}$, where $\bar{x}^\top = -\underline{x}^\top = [10^{-4} \ 10^{-7} \ 10^{-3} \ 10^{-3}]$, and controller limits given in [34]. A step change of -0.15 was applied to $\Delta P_{L,k}^{[i]} \equiv \Delta P_L^{[i]}(kT) \forall i \in \mathcal{I}$, and the results are shown in Figure 3.

VII. CONCLUSION AND FUTURE WORK

In this paper, we presented a scalable method to overapproximate the forward reachable sets of multi-agent systems with distributed NNC architectures. The effectiveness of this

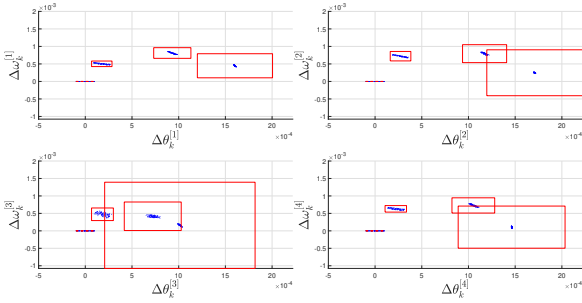


Fig. 3. Plots of the reachable sets in red (solid) for angle and frequency deviation and simulated trajectories (blue) for the power network example; the initial set is shown in red (dashed); here, $\Delta\theta_k^{[i]} \equiv \Delta\theta^{[i]}(kT) \forall i \in \mathcal{I}$ and $\Delta\omega_k^{[i]} \equiv \Delta\omega^{[i]}(kT) \forall i \in \mathcal{I}$

method was demonstrated on realistic examples, and it was shown to be significantly faster than the overall reachability method whilst producing the same bounds. Opportunities for future work include using the general framework presented in this paper to improve other reachability methods, such as LP-based methods. Also, the multi-agent reachability method could be used in robust NNC training.

REFERENCES

- [1] B. Karg and S. Lucia, "Efficient representation and approximation of model predictive control laws via deep learning," *IEEE Transactions on Cybernetics*, vol. 50, no. 9, pp. 3866–3878, 2020.
- [2] H. Li, T. Wei, A. Ren, Q. Zhu, and Y. Wang, "Deep reinforcement learning: Framework, applications, and embedded implementations: Invited paper," in *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2017, pp. 847–854.
- [3] 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.
- [4] P. Shi and B. Yan, "A survey on intelligent control for multiagent systems," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 51, no. 1, pp. 161–175, 2021.
- [5] M. Andreasson, "Control of multi-agent systems with applications to distributed frequency control power systems," Ph.D. dissertation, KTH Royal Institute of Technology, 2013.
- [6] M. Andreasson, D. V. Dimarogonas, H. Sandberg, and K. H. Johansson, "Distributed control of networked dynamical systems: Static feedback, integral action and consensus," *IEEE Transactions on Automatic Control*, vol. 59, no. 7, pp. 1750–1764, 2014.
- [7] R. Negenborn and J. Maestre, "Distributed model predictive control: An overview and roadmap of future research opportunities," *IEEE Control Systems Magazine*, vol. 34, no. 4, pp. 87–97, 2014.
- [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 *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 5929–5934.
- [9] A. Sadeghzadeh and P.-L. Garoche, "Reachability analysis of linear parameter-varying systems with neural network controllers," in *2022 IEEE Conference on Control Technology and Applications (CCTA)*, 2022, pp. 1372–1377.
- [10] M. Everett, G. Habibi, C. Sun, and J. P. How, "Reachability analysis of neural feedback loops," *IEEE Access*, vol. 9, pp. 163 938–163 953, 2021.
- [11] W. Xiang and T. T. Johnson, "Reachability analysis and safety verification for neural network control systems," *arXiv preprint arXiv:1805.09944*, 2018.
- [12] Y. Zhang and X. Xu, "Reachability analysis and safety verification of neural feedback systems via hybrid zonotopes," *arXiv preprint arXiv:2210.03244*, 2022.
- [13] —, "Safety verification of neural feedback systems based on constrained zonotopes," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2737–2744.

- [14] N. Kochdumper, C. Schilling, M. Althoff, and S. Bak, "Open-and closed-loop neural network verification using polynomial zonotopes," *arXiv preprint arXiv:2207.02715*, 2022.
- [15] M. Newton and A. Papachristodoulou, "Reachability analysis of neural feedback loops using sparse polynomial optimisation," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 2745–2750.
- [16] J. Fan, C. Huang, X. Chen, W. Li, and Q. Zhu, "ReachNN*: A tool for reachability analysis of neural-network controlled systems," *Automated Technology for Verification and Analysis*, p. 537.
- [17] N. Fijalkow, J. Ouaknine, A. Pouly, J. Sousa-Pinto, and J. Worrell, "On the decidability of reachability in linear time-invariant systems," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*. ACM, 2019.
- [18] K. Muandet, K. Fukumizu, B. Sriperumbudur, and B. Schölkopf, "Kernel mean embedding of distributions: A review and beyond," *Foundations and Trends® in Machine Learning*, vol. 10, no. 1-2, pp. 1–141, 2017.
- [19] P. Wang, H. Deng, J. Zhang, L. Wang, M. Zhang, and Y. Li, "Model predictive control for connected vehicle platoon under switching communication topology," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 7817–7830, 2022.
- [20] H. Hu, K. Gatsis, M. Morari, and G. J. Pappas, "Non-cooperative distributed mpc with iterative learning," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 5225–5232, 2020, 21st IFAC World Congress.
- [21] Q. Yuan and X. Li, "Distributed model predictive formation control for a group of uavs with newton-euler dynamics," in *2020 European Control Conference (ECC)*, 2020, pp. 1484–1489.
- [22] F. Lewis, H. Zhang, K. Hengster-Movric, and A. Das, *Cooperative Control of Multi-Agent Systems: Optimal and Adaptive Design Approaches*, ser. Communications and Control Engineering. Springer London, 2014.
- [23] X. Chen and S. Sankaranarayanan, "Decomposed reachability analysis for nonlinear systems," in *2016 IEEE Real-Time Systems Symposium (RTSS)*, 2016, pp. 13–24.
- [24] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin, "Decomposition of reachable sets and tubes for a class of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 11, pp. 3675–3688, 2018.
- [25] O. Gates, M. Newton, and K. Gatsis, "Scalable forward reachability analysis of multi-agent systems with neural network controllers," *arXiv preprint arXiv:2309.03846*, 2023.
- [26] J. C. Zegers, E. Semsar-Kazerooni, J. Ploeg, N. van de Wouw, and H. Nijmeijer, "Consensus control for vehicular platooning with velocity constraints," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 5, pp. 1592–1605, 2018.
- [27] Y. Zheng, S. Eben Li, J. Wang, D. Cao, and K. Li, "Stability and scalability of homogeneous vehicular platoon: Study on the influence of information flow topologies," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 1, pp. 14–26, 2016.
- [28] Y. Zheng, S. E. Li, K. Li, and L.-Y. Wang, "Stability margin improvement of vehicular platoon considering undirected topology and asymmetric control," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 4, pp. 1253–1265, 2016.
- [29] J. Ploeg, B. T. M. Scheepers, E. van Nunen, N. van de Wouw, and H. Nijmeijer, "Design and experimental evaluation of cooperative adaptive cruise control," in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2011, pp. 260–265.
- [30] S. Öncü, J. Ploeg, N. van de Wouw, and H. Nijmeijer, "Cooperative adaptive cruise control: Network-aware analysis of string stability," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 4, pp. 1527–1537, 2014.
- [31] S. Rivero, M. Farina, and G. Ferrari-Trecate, "Plug-and-play decentralized model predictive control for linear systems," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2608–2614, 2013.
- [32] M. Zeilinger, Y. Pu, S. Rivero, G. Ferrari-Trecate, and C. Jones, "Plug and play distributed model predictive control based on distributed invariance and optimization," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 5770–5776.
- [33] H. Saadat, *Power System Analysis*, ser. Electrical and Computer Engineering Series. WCB/McGraw-Hill, 1999.
- [34] S. Rivero and G. Ferrari-Trecate, "Hycon2 benchmark: Power network system," *arXiv preprint arXiv:1207.2000*, 2012.