



HAL
open science

Input State Stability of Gated Graph Neural Networks

Antonio Marino, Claudio Pacchierotti, Paolo Robuffo Giordano

► **To cite this version:**

Antonio Marino, Claudio Pacchierotti, Paolo Robuffo Giordano. Input State Stability of Gated Graph Neural Networks. IEEE Transactions on Control of Network Systems, 2024, pp.1-12. 10.1109/TCNS.2024.3372710 . hal-04110671v3

HAL Id: hal-04110671

<https://hal.science/hal-04110671v3>

Submitted on 24 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Input State Stability of Gated Graph Neural Networks

Antonio Marino¹, Claudio Pacchierotti², Paolo Robuffo Giordano²

Abstract—In this paper, we aim to find the conditions for input-state stability (ISS) and incremental input-state stability (δ ISS) of Gated Graph Neural Networks (GGNNs). We show that this recurrent version of Graph Neural Networks (GNNs) can be expressed as a dynamical distributed system and, as a consequence, can be analysed using model-based techniques to assess its stability and robustness properties. Then, the stability criteria found can be exploited as constraints during the training process to enforce the internal stability of the neural network. Two distributed control examples, flocking and multi-robot motion control, show that using these conditions increases the performance and robustness of the gated GNNs.

Index Terms—distributed control, graph neural network, stability analysis

I. INTRODUCTION

IN recent years, multi-agent systems have garnered considerable attention [1]. Coordinated multi-agent systems offer enhanced collaborative problem-solving capabilities and flexibility compared to single-agent approaches, making them suitable for various disciplines such as computer science, electrical engineering, and robotics [2]. Effective coordination of multiple agents requires considering group dynamics. In the case of multi-robot applications, generating individual robot motion involves not only local sensing data, but also information about the group state, often acquired through communication with a limited number of neighboring team members [3]. Thus, communication becomes a critical element in realizing distributed solutions for multi-agent systems.

Neural networks have gained prominence in data-driven control applications, given their exceptional approximation capabilities [4]. In distributed control, neural networks are particularly advantageous as they can approximate complex distributed policies without intricate optimizations and designs. The literature showcases various data-driven approaches, especially those rooted in reinforcement learning [5], [6], leveraging input data like images [7]. However, these approaches operate on local sensing data without inter-agent communication. On the other hand, distributed machine learning employs communication to partition the learning process across multiple machines contributing to group knowledge [8].

Leveraging communication, recent trends in data-driven distributed control involve employing Graph Neural Networks (GNNs) to encode distributed control and planning solutions. GNNs excel in predicting and analyzing graphs, proving

valuable in various problems such as text classification, protein interface predictions, and social network decisions [9]. Particularly in the latter cases, GNNs outperform classical neural network architectures [10], offering a fresh perspective on distributed control implementation. In this respect, Gama et al. [11] extend GNNs to flocking control for large teams. Additional examples of the use of GNNs for distributed control are found in space coverage [12], multi-robot path planning [13], and motion planning [14], including obstacle-rich environments [15]. GNNs also enhance multi-agent perception [16] and enable distributed active information acquisition [17], translating multi-robot information gathering into graph representation and formulating GNN-based decision-making.

In the recent literature, a very relevant discussion is about making learning based method robust and stable [18]. In this context, many works applied contraction analysis to demonstrate recurrent neural network stability [19], or directly closed-loop stability in continuous learning [20] and adaptive control [21]. Recently, the stability analyses presented in [22], [23] showcased the concepts of ISS (Input-to-State Stability) and incremental ISS (δ ISS) [24] in the context of LSTMs and GRUs, which are two of the most popular recurrent neural network models. The δ ISS property is a stronger property than plain ISS as it leads to the asymptotically convergence of two state trajectories when their respective input sequences are close, regardless of the initial conditions for the states. In a broader scope, D’Amico et al. [25] has established the criteria for achieving δ ISS in a general class of recurrent neural networks. Additionally, they have formally outlined the Linear Matrix Inequality (LMI) problem that guarantees the stability of the acquired neural network.

Inspired by these last results, this work characterizes the δ ISS properties of the recurrent version of GNN, i.e. Gated Graph Neural Networks (GGNN) [26]. These models use gated mechanisms to deploy distributed recurrent models able to reason on temporal- and spatial-based relationships among the agents. Compared to the recurrent systems in [25], GGNNs present a heightened level of complexity. This is primarily due to the fact that their underlying communication graph is not predetermined and undergoes changes over time, ultimately rendering the entire neural network variant across time.

To the best of our knowledge, this is the first time ISS and δ ISS have been proven for GGNN. Previous works have primarily focused on limited stability results, such as stability to graph perturbations [26], [27]. Instead, this article considers the system internal stability to the input features in a more general dynamical system analysis. The conditions derived in this work are in the form of nonlinear inequalities on GGNN weights: these can be exploited to certify the stability of a

¹ A. Marino is with Univ Rennes, CNRS, Inria, IRISA – Rennes, France. E-mail: antonio.marino@irisa.fr

² C. Pacchierotti and P. Robuffo Giordano are with CNRS, Univ Rennes, Inria, IRISA – Rennes, France. E-mail: {claudio.pacchierotti,prg}@irisa.fr.

This work was supported by the ANR-20-CHIA-0017 project “MULTI-SHARED”

trained neural model, or can be enforced as constraints during the training process to guarantee the stability of the GGNN.

The remaining sections of this paper are organized as follow. Section II reviews the preliminaries about graph neural networks and presents their recurrent versions, RGNNs and GGNNs. Section III and IV present the main stability results for the one-layer and deep GGNN. In Sect. V, we also provide a perspective of these results when considering communication delays. We then show two examples of the conditions found for flocking control, in Sect. VI-A, and multi robot motion control, in Sect. VI-B. Concluding remarks are finally presented in Sect. VII.

II. PRELIMINARIES

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph where $\mathcal{V} = \{v_1, \dots, v_N\}$ is the vertex set (representing the N agents in the group) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set. Each edge $e_k = (i, j) \in \mathcal{E}$ is associated with a weight $w_{ij} \geq 0$ such that $w_{ij} > 0$ if the agent i and j can interact and $w_{ij} = 0$ otherwise. As usual, we denote with $\mathcal{N}_i = \{j \in \mathcal{V} \mid w_{ij} > 0\}$ the set of neighbors of agent i . We also let $\mathbf{A} \in \mathbb{R}^{N \times N}$ be the adjacency matrix with entries given by the weights w_{ij} . Defining the degree matrix $\mathbf{D} = \text{diag}(d_i)$ with $d_i = \sum_{j \in \mathcal{N}_i} w_{ij}$, the Laplacian matrix of the graph is $\mathbf{L} = \mathbf{D} - \mathbf{A}$.

The graph signal $\mathbf{x} \in \mathbb{R}^N$, whose i -component x_i is assigned to agent i , can be processed over the network by the following linear combination rule applied by each agent

$$\mathbf{s}_i \mathbf{x} = \sum_{j \in \mathcal{N}_i} s_{ij} (x_i - x_j), \quad (1)$$

where \mathbf{s}_i is the i -th row of \mathbf{S} . The signal manipulation can be operated by means of any *graph shift operator* \mathbf{S} , e.g., Laplacian, adjacency matrix, weighted Laplacian, which respects the sparsity pattern of the graph. Later, in Sects. VI-A, VI-B, we will use the Laplacian as support matrix, as it is commonly used in distributed control. However, the proposed techniques do not assume the use of a specific support matrix.

Performing k repeated applications of \mathbf{S} on the same signal represents the aggregation of the k -hop neighborhood information. In analogy with traditional signal processing, this property can be used to define a linear graph filtering [28] that processes the multi features signal $\mathbf{x} \in \mathbb{R}^{N \times G}$ with G features:

$$\mathbf{H}_S(\mathbf{x}) = \sum_{k=0}^K \mathbf{S}^k \mathbf{x} \mathbf{H}_k. \quad (2)$$

where the weights $\mathbf{H}_k \in \mathbb{R}^{G \times F}$ define the output of the filter. Note that $\mathbf{S}^k = \mathbf{S}(\mathbf{S}^{k-1})$, so that it can be computed locally with repeated 1-hop communications between a node and its neighbors. Hence, the computation of \mathbf{H}_S is distributed on each node.

A. Graph Neural Network

Although \mathbf{H}_S is simple to evaluate, it can only represent a linear mapping between input and output filters. GNNs increase the expressiveness of the linear graph filters by means of pointwise nonlinearities $\rho : \mathbb{R}^{N \times F_{l-1}} \rightarrow \mathbb{R}^{N \times F_{l-1}}$ following a filter bank. Letting \mathbf{H}_{S_l} be a bank of $F_{l-1} \times F_l$ filters at layer l , the GNN layer is defined as

$$\mathbf{x}_l = \rho(\mathbf{H}_{S_l}(\mathbf{x}_{l-1})), \quad \mathbf{x}_{l-1} \in \mathbb{R}^{N \times F_{l-1}}. \quad (3)$$

Starting by $l = 0$ with F_0 , the signal tensor $\mathbf{x}_{l_n} \in \mathbb{R}^{N \times F_{l_n}}$ is the output of a cascade of l_n GNN layers. This specific type of GNN is commonly referred to as a convolution graph network because each layer utilizes a graph signal convolution (2). By the use of imitation learning or similar techniques, the GNN can learn a distributed policy by finding the optimal filter weights \mathbf{H}_k to propagate information among the agents and generate the desired output. Notably, each agent employs an identical version of the network and exchanges intermediate quantities with the other agents in the network through the application of \mathbf{S} , resulting in an overall distributed neural network. GNNs inherit some interesting properties from graph filters, such as permutational equivariance [27] and their local and distributed nature, showing superior ability to process graph signals [17], [29], [30]. Under dynamic graphs and perturbations of \mathbf{S} , the stability of GNN predictions is related to \mathbf{S} spectral characteristics. In particular, the graph filters composing the GNN have a frequency response, $h(\lambda)$ [31], bounded by the graph support eigenvalues λ if the graph filters are integral Lipschitz:

$$|h(\lambda_j) - h(\lambda_i)| \leq 2C \frac{|\lambda_j - \lambda_i|}{|\lambda_j + \lambda_i|} \quad (4)$$

where $\lambda_j, \lambda_i \in \mathbb{R}$ are any support matrix eigenvalues, and $C > 0$ a proper integral Lipschitz constant. This condition restricts the graph frequency response variability to the midpoint of λ variations.

B. Gated Graph Neural Network

Recurrent models of GNNs can solve time-dependent problems. These models, similarly to recurrent neural networks (RNNs), are known as graph recurrent neural networks (GRNNs). GRNNs utilize memory to learn patterns in data sequences, where the data is spatially encoded within graphs, regardless of the team size of the agents [32]. However, traditional GRNNs encounter challenges such as vanishing gradients, which are also found in RNNs. Additionally, they face difficulties in handling long sequences in space, where certain nodes or paths within the graph might be assigned more importance than others in long-range exchanges, causing imbalances in the graph's informational encoding.

Forgetting factors can be applied to mitigate this problem, reducing the influence of past or new signal on the state. A Gated Graph Neural Network (GGNN) [26] is a recurrent Graph Neural Network that uses a gating mechanism to control how the past information influences the update of the GNN states. We can add a state and an input gate, $\hat{\mathbf{q}}, \tilde{\mathbf{q}} \in Q \subseteq [0, 1]^{N \times F}$, that are multiplied via the Hadamard product \circ by the state and the inputs of the network, respectively. These two gates regulate how much the past information or the input are used to update the network internal state. GGNNs admit the following state-space representation [33],

$$\begin{cases} \tilde{\mathbf{q}} = \sigma(\tilde{\mathbf{A}}_S(\mathbf{x}) + \tilde{\mathbf{B}}_S(\mathbf{u}) + \tilde{\mathbf{b}}) \\ \hat{\mathbf{q}} = \sigma(\hat{\mathbf{A}}_S(\mathbf{x}) + \hat{\mathbf{B}}_S(\mathbf{u}) + \hat{\mathbf{b}}) \\ \mathbf{x}^+ = \sigma_c(\hat{\mathbf{q}} \circ \mathbf{A}_S(\mathbf{x}) + \tilde{\mathbf{q}} \circ \mathbf{B}_S(\mathbf{u}) + \mathbf{b}) \end{cases} \quad (5)$$

with $\sigma(x) = \frac{1}{1+e^{-x}}$ being the logistic function, and $\sigma_c(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$ being the hyperbolic tangent. $\hat{\mathbf{A}}_S, \hat{\mathbf{B}}_S$ are graph filters

of the forgetting gate, \tilde{A}_S, \tilde{B}_S are graph filters (2) of the input gate, and A_S and B_S are the state graph filters (2). $\hat{\mathbf{b}}, \tilde{\mathbf{b}}, \mathbf{b} \in \mathbb{R}^{N \times F}$ are respectively the biases of the gates and the state built as $\mathbf{1}_N \otimes \mathbf{b}$ with the same bias for every agents. The weights and biases of the graph filters can be acquired using imitation learning techniques [34]

We note that, in the literature, there exist several Gated GNNs structures [9], [26], [33], sharing gating mechanisms of different kinds (temporal, attention, and so on). It is easy to check that the system in (5) generalizes the systems above, thereby offering a broader framework for analyzing their dynamic stability. We want to characterize the dynamic stability properties of such neural models that combine gating mechanisms with a distributed computation. Therefore, we have to analyze the state \mathbf{x} evolution and its convergence as a nonlinear dynamical system.

C. Incremental Input State Stability

In this section, we recall the definition of ISS and δ ISS that will be used throughout the paper. The ISS property guarantees that, regardless of the starting conditions, limited inputs or disturbances lead to limited system states. This attribute is advantageous in control systems, as they require stability and robustness. However, this principle does not hold for systems that are inherently unstable. Recalling the definitions of \mathcal{KL} , \mathcal{K}_∞ functions [35], and the infinite norm $\|\cdot\|_\infty$, the following definition of ISS is given

Definition II.1 (ISS). System (5) is called *input-to-state stable* if there exist functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}_\infty$ such that, for any $t \in \mathbb{Z} \geq 0$, any initial state $\mathbf{x}(0) \in \mathcal{X}$ any input sequence $\mathbf{u} \in \mathcal{U}$ it holds that:

$$\|\mathbf{x}(t)\|_\infty \leq \beta(\|\mathbf{x}(0)\|_\infty, t) + \gamma_u(\|\mathbf{u}\|_\infty) + \gamma_b(\|\mathbf{b}\|_\infty) \quad (6)$$

A further desirable property is incremental ISS (δ ISS) [36]. The δ ISS property ensures that any pair of state trajectories converge towards each other even if they start from different initial conditions. Moreover, their difference is bounded only by the differences of their inputs (e.g., an ideal control corrupted by an additive noise), thus enhancing the system robustness [37].

Definition II.2 (δ ISS). System (5) is called *incrementally input-to-state stable* [24] if there exist functions $\beta_\delta \in \mathcal{KL}$ and $\gamma_\delta \in \mathcal{K}_\infty$ such that, for any $t \in \mathbb{Z} \geq 0$, any initial states $\mathbf{x}(0)_1, \mathbf{x}(0)_2 \in \mathcal{X}$ any input sequences $\mathbf{u}_1, \mathbf{u}_2 \in \mathcal{U}$ it holds that:

$$\|\mathbf{x}(t)_1 - \mathbf{x}(t)_2\|_\infty \leq \beta_\delta(\|\mathbf{x}(0)_1 - \mathbf{x}(0)_2\|_\infty, t) + \gamma_\delta(\|\mathbf{u}_1 - \mathbf{u}_2\|_\infty) \quad (7)$$

Remark 1. In the neural network context, the δ ISS property ensures that any difference in the initial conditions will be eventually discarded, and thus the same outputs will correspond to the same observations. Moreover, since the stability is valid for $t > 0$, for a training with a finite time sequence dataset it is guaranteed that all the NN state trajectories converge to a unique solution.

III. ONE-LAYER GGNN STABILITY

In this section we will discuss the stability properties of a single layer GGNN. For the rest of the paper the following assumption will be made

Assumption 1. The input \mathbf{u} is unity-bounded: $\mathbf{u} \in \mathcal{U} \subseteq [-1, 1]^{N \times G}$, i.e. $\|\mathbf{u}\|_\infty \leq 1$.

This is a quite mild assumption since the input signal is usually normalized or it is the result of others network layers with unitary output activation functions.

Before stating the sufficient conditions for the ISS of GGNN, we will first introduce the notation for the following quantities

$$\begin{aligned} S_K &\triangleq [I, S, \dots, S^K] \\ A &\triangleq [A_0, \dots, A_K]^T & B &\triangleq [B_0, \dots, B_K]^T \\ \tilde{A} &\triangleq [\tilde{A}_0, \dots, \tilde{A}_K]^T & \hat{A} &\triangleq [\hat{A}_0, \dots, \hat{A}_K]^T \\ \tilde{B} &\triangleq [\tilde{B}_0, \dots, \tilde{B}_K]^T & \hat{B} &\triangleq [\hat{B}_0, \dots, \hat{B}_K]^T \end{aligned} \quad (8)$$

where K is the filters length. Then, in light of assumption (1) and knowing that $\|\mathbf{x}\| \leq 1$, each gate feature q_i satisfies:

$$\begin{aligned} |\hat{q}_i| &\leq \|\hat{\mathbf{q}}\|_\infty \leq \max_{\mathbf{u} \in \mathcal{U}, \mathbf{x} \in \mathcal{X}} \|\sigma(\hat{A}_S(\mathbf{x}) + \hat{B}_S(\mathbf{u}) + \hat{\mathbf{b}})\|_\infty \\ &\leq \max_{\mathbf{u} \in \mathcal{U}, \mathbf{x} \in \mathcal{X}} \|\sigma(\hat{A}_S(\mathbf{x}) + \hat{B}_S(\mathbf{u}) + \hat{\mathbf{b}})\|_\infty \\ &\leq \sigma(\max_{\mathbf{u} \in \mathcal{U}, \mathbf{x} \in \mathcal{X}} \|\hat{A}_S(\mathbf{x}) + \hat{B}_S(\mathbf{u}) + \hat{\mathbf{b}}\|_\infty) \\ &\leq \sigma(\|\hat{A}_S(\mathbf{x}_{max}) + \hat{B}_S(\mathbf{u}_{max}) + \hat{\mathbf{b}}\|_\infty) \end{aligned} \quad (9)$$

where \mathbf{x}_{max} and \mathbf{u}_{max} are the maximum values of \mathbf{x} and \mathbf{u} . Recalling that a graph filter in equation (2) can be written as $\hat{H}_S(\mathbf{x}) = [I, S, \dots, S^K](I_K \otimes \mathbf{x})[\hat{H}_0, \hat{H}_1, \dots, \hat{H}_K]^T$ and using notation (8), equation (9) becomes

$$\begin{aligned} |\hat{q}_i| &\leq \sigma(\|S_K\|_\infty(\|\hat{A}\|_\infty\|\mathbf{x}\|_\infty + \|\hat{B}\|_\infty\|\mathbf{u}\|_\infty) + \|\hat{\mathbf{b}}\|_\infty) \\ &= \sigma(\|S_K\|_\infty(\|\hat{A}\|_\infty + \|\hat{B}\|_\infty) + \|\hat{\mathbf{b}}\|_\infty) \triangleq \sigma_{\hat{q}}. \end{aligned} \quad (10)$$

We identify the induced ∞ -norm as $\|\cdot\|_\infty$. Similarly for \tilde{q}_i

$$|\tilde{q}_i| \leq \sigma(\|S_K\|_\infty(\|\tilde{A}\|_\infty + \|\tilde{B}\|_\infty) + \|\tilde{\mathbf{b}}\|_\infty) \triangleq \sigma_{\tilde{q}}. \quad (11)$$

Theorem 1. A sufficient condition for the ISS of a single-layer GGNN network is that $\mathcal{A} \leq 1$, where

$$\mathcal{A} \triangleq \sigma_{\hat{q}}\|S_K\|_\infty\|A\|_\infty. \quad (12)$$

The proof based on the results for GRU and LSTM [22], [23] is provided in the Appendix I-A.

The δ ISS stability requires to analyze the evolution of the maximum distance of two states trajectories $\mathbf{x}_1, \mathbf{x}_2$, starting from two different initial conditions $\mathbf{x}_1(0), \mathbf{x}_2(0)$ and having two different inputs $\mathbf{u}_1, \mathbf{u}_2$. We also need another assumption

Assumption 2. Given two support matrices $\|S_1(t)\|_\infty, \|S_2(t)\|_\infty, \forall t \in \mathbb{Z}^+$ associated with two different graphs, they are bounded by the same $\|\bar{S}\|_\infty$.

Theorem 2. Under the assumption 2, a sufficient condition for the system (5) to be δ ISS is $\mathcal{A}_\delta \leq 1$; where

$$\begin{aligned} \mathcal{A}_\delta &\triangleq \sigma_{\hat{q}}\|\bar{S}\|_\infty\|A\|_\infty + \frac{1}{4}\|\bar{S}\|_\infty^2\|\hat{A}\|_\infty\|A\|_\infty \\ &\quad + \frac{1}{4}\|\bar{S}\|_\infty^2\|\tilde{A}\|_\infty\|B\|_\infty. \end{aligned} \quad (13)$$

The proof to this theorem is reported in Appendix I-B. Assumption 2 is reasonable for multi-agent systems, since graphs always have a finite number of agents with finite number of links between each other. When we deal with scalability and dynamic graphs, the assumption may be restrictive based on the choice of the support matrix. For examples, by using the adjacency matrix the upper bound of its norm would be the maximum number of the links for one agent, i.e. N (the team size). Therefore, training the GGNN using the adjacency matrix and being $\|\tilde{S}\|_\infty = N$, the stability condition would be respected for teams with a maximum number of links for each agent up to N . For group with $N' > N$ we can guarantee stability if the agents have a number of neighbors less than N .

Remark 2. The condition presented in Theorem 2 also offers an approach for verifying the δ ISS characteristic of a GGNN layer, in contrast of relying on a statistical analysis of state convergence. This condition, while sufficient, provides a computationally inexpensive tool for assessment.

Remark 3. In practice we can solve this issue by constraining, at runtime, the cardinality of $\#\mathcal{N}_i < N$ for every agent i in the team. For a GGNN that uses normalized support matrices, e.g., normalized Laplacian, the assumption is met without any further restrictions on the graph topology. Moreover the use of normalized support reduces the regularizing term in the loss function at training time. In the experiments we will use normalized Laplacian to compare stable and unstable neural network.

Remark 4. Note that we did not make any assumptions on the shape of S ; the graph can be made of connected or disconnected sub-graphs without affecting the stability.

While the identified conditions are deemed sufficient, it is important to acknowledge the possibility that certain GGNN layers may exhibit δ ISS or ISS characteristics without necessarily conforming to the conditions outlined in Theorem 2 or Theorem 1. In the upcoming sections, we will demonstrate that satisfying these conditions, although potentially restrictive, leads to enhanced performance.

IV. DEEP GGNN STABILITY

With the word ‘‘deep’’ we refer to the stack of multiple layers of the Neural Networks. It is natural to ask whether the stability properties extend from the single layer to a multi-layer structure, commonly used in practice. A deep GGNN, whose representation is reported in the supplementary material, is made by interconnections among layers by feeding the future state of one layer to the next one,

$$\begin{cases} \tilde{\mathbf{q}}^i = \sigma(\tilde{A}_S(\mathbf{x}^i + \tilde{B}_S(\mathbf{u}^i)) + \tilde{\mathbf{b}}) \\ \hat{\mathbf{q}}^i = \sigma(\hat{A}_S(\mathbf{x}^i + \hat{B}_S(\mathbf{u}^i)) + \hat{\mathbf{b}}) \\ \mathbf{x}^{i+} = \sigma_c(\hat{\mathbf{q}}^i \circ A_S(\mathbf{x}^i) + \tilde{\mathbf{q}}^i \circ B_S(\mathbf{u}^i) + \mathbf{b}) \\ \mathbf{u}^i = \mathbf{x}^{i-1+}, \quad \mathbf{u}^1 = \mathbf{u} \end{cases} \quad (14a)$$

for all the layers $i \in \{1, \dots, M\}$. The output of the network results from a graph output filter

$$\mathbf{y} = Y_S(\mathbf{x}^M) + \mathbf{b}_y. \quad (14b)$$

The following then holds

Theorem 3. The GGNN network is ISS if $\mathcal{A}^i \leq 1$ for every layer $i \in \{1, \dots, M\}$, where \mathcal{A}^i are defined like in Theorem 1

Proof. The deep GGNN (14) can be considered as a cascade of subsystems, so it is ISS if every subsystem is ISS. \square

The δ ISS condition is more complex since there does not exist, to our knowledge, a general study in the literature for a cascade of δ ISS systems. However, in this case we can state

Theorem 4. The deep GGNN is δ ISS stable if $\mathcal{A}_\delta^i \leq 1$ for every layer $i \in \{1, \dots, M\}$, where \mathcal{A}_δ^i are defined like in Theorem 2.

The proof is reported in Appendix I-C. δ ISS is a preferable property for deep GGNN as well as for single layer networks, as this property guarantees that any two state trajectories will converge under the same inputs.

V. GGNN STABILITY UNDER COMMUNICATION DELAY

The previous scheme of GGNN can represent any dynamics on a graph. However, it does not take into account the communication steps and the possible delay in the application of the support matrix S . In each communication step, performed with sampling time T , we considered that the nodes communicate the data used in the graph filters and, at each T , we obtain the output of a GGNN layer computed on the data $[\mathbf{x}(t-K), \mathbf{x}(t-(K-1)), \dots, \mathbf{x}(t)]$. We express the delayed system in (5) as:

$$\begin{cases} \tilde{\mathbf{q}} = \sigma(\tilde{A}_{St}(\mathbf{x}(t-K)) + \tilde{B}_{St}(\mathbf{u}(t-K)) + \tilde{\mathbf{b}}) \\ \hat{\mathbf{q}} = \sigma(\hat{A}_{St}(\mathbf{x}(t-K)) + \hat{B}_{St}(\mathbf{u}(t-K)) + \hat{\mathbf{b}}) \\ \mathbf{x}(t+1) = \sigma_c(\hat{\mathbf{q}} \circ A_{St}(\mathbf{x}(t-K)) + \tilde{\mathbf{q}} \circ B_{St}(\mathbf{u}(t-K)) + \mathbf{b}) \\ \mathbf{y} = Y_{St}(\mathbf{x}(t-K)) + \mathbf{b}_y \end{cases} \quad (15)$$

where the delay of K communication steps can be represented by a delay between the input, old of K steps, and the updating rule of the state. Moreover, we must consider graph filters with dynamic support matrices, i.e. matrices that change between each communication step. Thus, the graph filters considered so far become

$$H_{St}(\mathbf{x}(t)) = \begin{bmatrix} I_N \\ \vdots \\ \prod_{\tau=t-(K-1)}^t S(\tau) \end{bmatrix}^T \text{diag} \left(\begin{bmatrix} \mathbf{x}(t) \\ \vdots \\ \mathbf{x}(t-K) \end{bmatrix} \right) \begin{bmatrix} H_0 \\ \vdots \\ H_K \end{bmatrix}$$

This filter expression is called unit-delayed filter [38]. Under Assumption 2, the infinite norm of the support matrices in time is upper bounded. Therefore, the formulation (15) allows us to conclude that the system remains δ ISS under the same conditions of the theorem 2.

Remark 5. As noted in [38], the trajectories and the underlying graph observed at training time and the one observed at deployment are different, causing an increase of the training error. This is not an issue since, at a reasonable sampling time, the sequence of support matrices will not present drastic changes as their spectral characteristics are

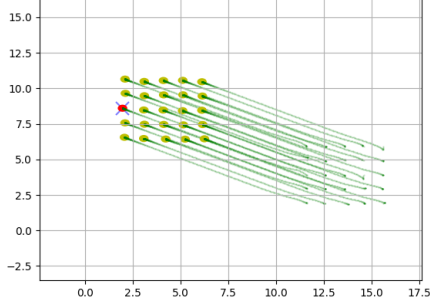


Fig. 1: Flocking control: a group of agents (yellow dots) move in order to reach the same velocity and to avoid collision. The leader (red dot) moves in order to reach the target (blue cross) and avoids the collision with the other agents.

similar and, under Theorem 2, the graph filters satisfy (4) with $C = (\|H\|_\infty - 1)\|\bar{S}\|_\infty/2$, and thus are stable to graph perturbations.

VI. VALIDATION EXAMPLES

We experimentally confirmed the stability condition on two distributed control scenarios: flocking control and multi-robot motion control. Despite their simplicity, these scenarios have been previously studied [39], [40] using different approaches such as GNN and LSTM, allowing for a direct comparison with stable GGNN sGGNN). We also explicitly compare a sGGNN which satisfies the condition in Theorem 2 and a GGNN that does not satisfy either Theorem 1 or 2. The stability condition is imposed by the following regularization in addition to the training loss:

$$\Pi = \sum_{i=0}^L \rho_- \min(0, \delta \mathcal{A}_i - 1 - \epsilon) + \rho_+ \max(0, \delta \mathcal{A}_i - 1 - \epsilon). \quad (16)$$

This regularization applies the condition in the Theorem 2 for each layer of GGNN with $0 < \rho_- \ll \rho_+$. By choosing ρ_- , Π enforces fast convergence pushing $\delta \mathcal{A}_i$ toward zero at the cost of less accuracy, since the condition in Theorem 2 is only a sufficient condition. For this reason, the choice of ρ_- must be designed in the training process to achieve a desirable estimation error. In our implementation, we used $\rho^+ = 1$ and $\rho^- = 0.01$.

A. Flocking Control Example

In the following, we show an application of the sGGNN on a case study involving flocking control (Fig. 1) with a leader. In the problem of flocking, the agents are initialized to follow random velocities while the goal is to have them all fly at the same velocity while avoiding collisions with each other. Moreover, one of the agents takes the leader role, conducting the team toward a target unknown to the other agents. Flocking is a canonical problem in decentralized robotics [41], [42].

Dynamics and Expert Controller

We considered N agents described by the position $\mathbf{r}(t) \in \mathbb{R}^{N \times 2}$ and the velocity $\mathbf{v}(t) \in \mathbb{R}^{N \times 2}$ with a double integrator dynamics

$$\mathbf{r}(t+1) = \mathbf{r}(t) + T\mathbf{v}(t); \quad \mathbf{v}(t+1) = \mathbf{v}(t) + T\mathbf{u}(t);$$

with the discrete acceleration $\mathbf{u}(t) \in \mathbb{R}^{N \times 2}$ taken as system input. Note that the agent dynamics is used for building the dataset and for simulation purpose, but it is not provided to the learning algorithm.

Since the objective is to make all the agents reach the same velocity, the control must be tuned in order to minimize the following cost function

$$J(\mathbf{v}(t)) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{v}_i(t) - \frac{1}{N} \sum_{j=1}^N \mathbf{v}_j(t)\|_2^2 \quad (17)$$

where $\mathbf{v}_i, \mathbf{v}_j$ are the velocities for the agent i and j , respectively. The cost function measures the distance of the agent velocities from the average velocities of the team. The cost $J(\mathbf{v}(t))$ under the control $\mathbf{u}(t)$ can be analysed on the time horizon T such that $t \in [0, T]$ to evaluate the convergence rate of the system. Moreover, the leader, randomly picked among the agents, does not follow the objective in (17) but rather it minimizes its distance from the target (\mathbf{d}). Therefore, the expert controller [39] for the followers is given by

$$\mathbf{u}_f(t) = -\mathbf{L}(t)\mathbf{v}(t) - \nabla_{\mathbf{r}} CA(\mathbf{r}(t), \mathbf{r}_j(t))|_{j=1 \dots N} \quad (18a)$$

and for leader it is

$$\mathbf{u}_l(t) = -W_p(\mathbf{r}_l(t) - \mathbf{d}(t)) - \nabla_{\mathbf{r}_l} CA(\mathbf{r}_l(t), \mathbf{r}_j(t))|_{j=1 \dots N} \quad (18b)$$

where W_p is a gain, $\mathbf{r}_l \in \mathbb{R}^2$ is the leader position, $\nabla_{\mathbf{r}} CA(\mathbf{r}(t), \mathbf{r}_j(t))/\nabla_{\mathbf{r}_l} CA(\mathbf{r}_l(t), \mathbf{r}_j(t))$ are the gradient of the collision avoidance potential with respect to the position of the agents/leader \mathbf{r}/\mathbf{r}_l , evaluated at the position $\mathbf{r}(t)/\mathbf{r}_l(t)$ and the position of every other agent $\mathbf{r}_j(t)$ at time t . The i -element of $\nabla_{\mathbf{r}} CA$ for each robot i with respect to robot j is given by [43]

$$\nabla_{\mathbf{r}_i} CA(\mathbf{r}_{ij}) = \begin{cases} -\frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^4} - \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^2} & \text{if } \|\mathbf{r}_{ij}\|_2 \leq R_{CA} \\ \mathbf{0} & \text{otherwise} \end{cases} \quad (19)$$

with $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ and $R_{CA} > 0$ indicating the minimum acceptable distance between agents. This potential function is a non negative, non smooth function that goes to infinity when the distance reduces and grows when the distance exceeds R_{CA} , in order to avoid the team losing the connectivity [43]. $\mathbf{u}_f(t), \mathbf{u}_l(t)$ are a centralized controller since computing them requires agent i to have instantaneous evaluation of $\mathbf{L}(t)\mathbf{v}(t)$ and $\mathbf{r}_j(t)$ of every other agent j in the team. R_{CA} and W_p are tunable parameters of the controllers.

Neural Network Architecture

We assume that the agents form a communication graph when they are in a sphere of radius R between each others and that exchanges occur at the sampling time T , so that the action clock and the communication clock coincide.

The input features vector $\mathbf{w}_i \in \mathbb{R}^{10}$ of the robot i for the designed neural network is

$$\mathbf{w}_i = \left[\mathbf{v}_i, \sum_{j \in \mathcal{N}_{S_i}} \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^4}, \sum_{j \in \mathcal{N}_{S_i}} \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^2}, \{\mathbf{0}_2, \mathbf{r}_l - \mathbf{d}\}, \{[0, 1], [1, 0]\} \right] \quad (20)$$

where $\mathcal{N}\mathcal{S}_i$ is the set of the sensing agents within a sphere of radius R_{CA} centred in the robot i . Moreover, the vector contains the zero vector $\mathbf{0}_2 \in \mathbb{R}^{1 \times 2}$ and the one-hot encoding $[0, 1]$, if the agent is a follower while $(r_l - d)$ and $[1, 0]$, if the agent is a leader. We chose the one-hot encoding instead of the binary one, because it allows differentiating the neural network weights between the leader and the follower. Note that all the information in the vector w_i are locally available at the sampling/control time T .

The core of the neural network for the flocking control is a layer of GGNN with $F = 50$ features in the hidden state and filters length $K = 2$. Note that the choice of K affects the complexity of the stability condition imposed since it will constraint more parameters. The input features are first processed by a cascade of two fully connected layers of 128 nodes before feeding the graph neural network. A readout of two layers with 128 nodes combines the F -features GGNN hidden state to get the bidimensional control u saturated to the maximum admissible control. The input layers and the readout that encapsulate the GGNN shape a more realistic setting to test the stability of the graph neural network that is usually used in combination with other kinds of neural models. Following the Remark 3, we used normalized Laplacian as a support matrix

Training

We collected a dataset by recording 120 trajectories, further separated into three subsets of training, validation and test set using the proportion 70% – 10% – 20%, respectively. Each trajectory is generated by randomly positioning the agents in a square such that their inter-distance is between 0.6 m and 1.0 m and their initial velocities are picked at random from the interval $[-2, 2]$ m/s in each direction. The leader is randomly selected among the agents and the target position is randomly located within a square of length 20 m centered at the location of the leader. Regardless of the target location, the trajectories have a duration of 2.5 s and input saturation at 5 m/s^2 . Moreover, the 120 trajectories are recorded with a random number of agents among $N = [4, 6, 10, 12, 15]$. We fixed the communication range to $R = 4$ m and the sensing to $R_{CA} = 1$ m. We trained the models for 120 epochs and executed the DAGGER algorithm [34] every 20 epochs. The algorithm evaluates the expert controller in (18) on the enrolled state trajectories applying the learned control and adding them to the training set. Note that, thanks to the use of DAGGER, we do not need a large dataset. We solve the imitation learning problem using the ADAM algorithm [44] with a learning rate $1e - 3$ and forgetting factors 0.9 and 0.999. The loss function used for imitation learning is the mean squared error between the output of the model and the optimal control action.

Results

In Fig. 2, we show a comparison between the stable GGNN (sGGNN) and non-stable GGNN controller for the flocking controller case. We evaluate the two controllers on 3 sets of experiments with 40 trajectories each. In the experiments, we varied team size, communication range and network delay to test the robustness of the controllers. Figures 2b, 2d and 2f report the leader position error evaluated after a fixed time of 2.5 s with respect to the leader starting location, i.e. e_f/e_s

with e_f, e_s respectively being the final and the initial square distance of the leader from the target. Figures 2a, 2c and 2e show the average flocking error (17) in the interval $0 - 2.5$ s in logarithmic scale. Moreover, we consider a failure when the control leads to an agent-agent collision, the leader-target distance diverges, or any agent-agent distance diverges (i.e. the team splits). In the first experiment, controllers’ scalability is evaluated for team sizes $N = [4, 10, 25, 50]$ while maintaining a fixed communication range of 4 m and no communication delay. Figs. 2a and 2b show sGGNN’s improvement in achieving the flocking state, as seen by the average error, which is generally closer to the expert controller than GGNN. As the number of agents increases, inter-agent collision avoidance leads to low flocking error due to motion constraints for distributed and centralized controllers. However, in the case of 50 agents, the leader struggles to drive the team toward the target due to the group cohesion, resulting in higher flocking and leader errors for all controllers. In general, sGGNN and GGNN show similar leader errors, even if, on average, this error is 7%-12% smaller for the non-stable neural network, at the cost of a higher flocking error.

In the second set of experiments, sGGNN demonstrates enhanced robustness when communication range differs from the training range $R = 4$ m, as can be seen in Figs. 2c, 2d, where the range varies between $[2, 4, 8, 10]$ and the team size is set at $N = 25$ under instantaneous communication. Specifically, when the range is $R = 2$ m, GGNN causes a flocking error of 0.40 m/s with outliers up to 0.68 m/s, greater than the average of 0.13 m/s of the sGGNN. Moreover, with GGNN, we experienced more group division and a consequent splitting of the team, which explains the increase in the flocking error. However, in this condition, we also report the leader divergence and a consequent worse leader error with respect to the stable control of 27%. While the stable controller never results in a failure situation, the non-stable GGNN causes 12% failures with a range of $R = 2$ m. As expected, when the communication range increases, the flocking errors decrease for both controllers, since they are able to communicate with more agents at the same time. However, in this case, the leader error increases, since the leader is often encapsulated by the other team agents and it is thus forced to follow them to not collide, causing a slower convergence to the target. Note that this behavior also affects the expert controller.

In the last set of experiments, we consider not instantaneous communication (delay of $T = 0.01$ s) and evaluate again the transference at scale. With large team sizes, the non-stable controller fails 20%-35% of the experimented trajectories, while the stable one succeeded 100%. Figures 2e, 2f show the consistency of the results with respect to the non-delay case, confirming that the controller remains stable even with not instantaneous communication. In some cases, sGGNN shows a high error variance. This behavior can be seen, for example, in Fig 2a with $N=10$, and in Fig 2c with $R=4m$, where the sGGNN errors boxes are large around the mean value. However, sGGNN error is overall lower and remains consistently closer to the expert controller one.

In Fig. 3, we also reported the performances of the sGGNN

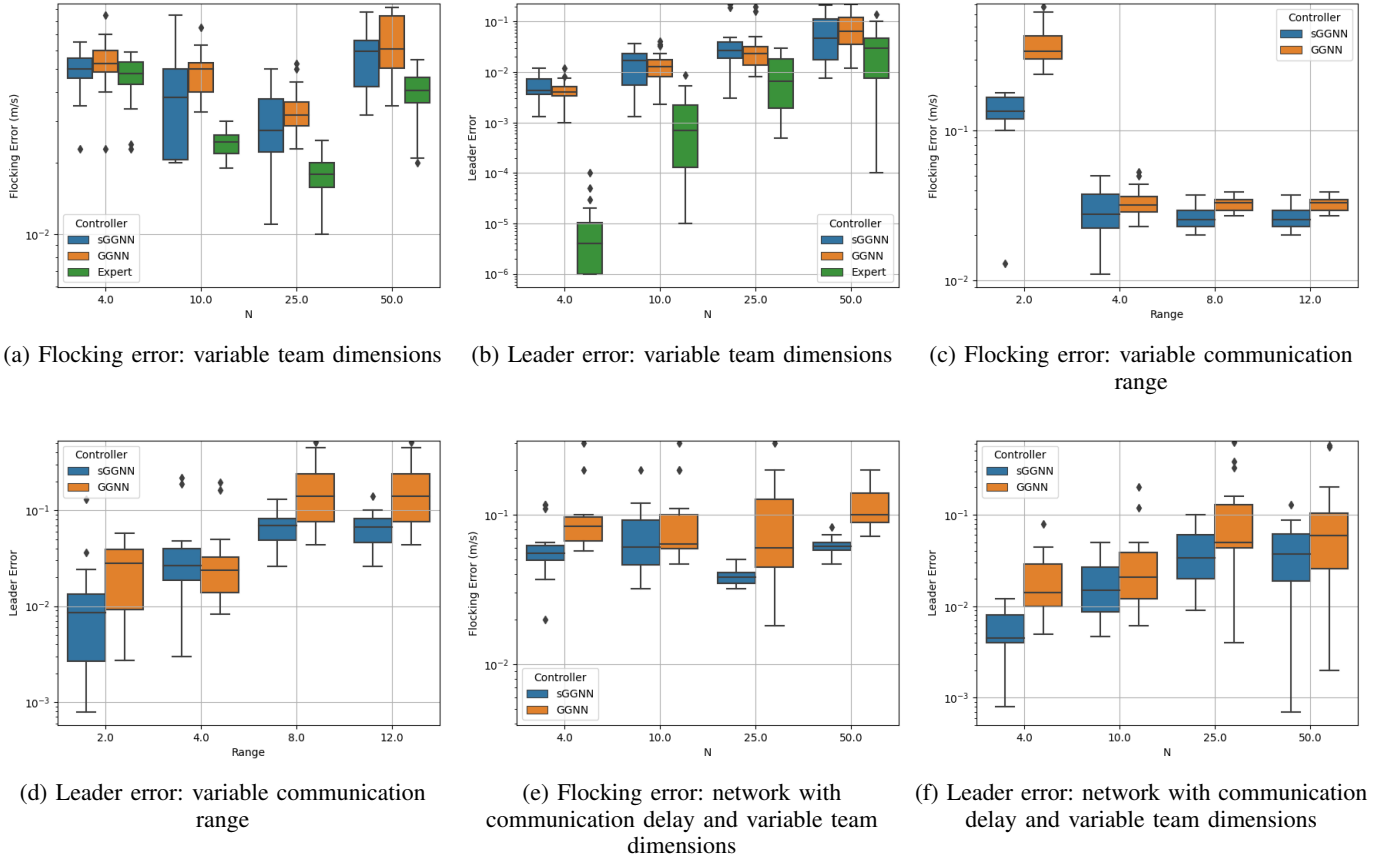


Fig. 2: Flocking and Leader Error for stable (sGGNN) and non-stable GGNN controllers, varying the team size N with fixed communication range of 4m (with and without instantaneous communication) and the communication range with $N = 25$, reported using box plots that display median, minimum, maximum, 25th/75th percentiles, and outliers.

and GGNN trained and tested with the Laplacian matrix to confirm what we stated in the Remark 3. As we can see, with a higher communication range than the one used at training time, the differences between the two controllers disappear except for differences caused by training errors. Even if the condition in Theorem 2 is only sufficient and possibly conservative, from the results, it is clear that imposing stability over the distributed controller improved the agent performances compared to scenarios where stability was not imposed.

B. Multi Robot Motion Control Example

In this section, we address multi-robot motion in a cluttered space (Fig. 4), a problem already studied in [7], [45], [46]. Our approach is based on [47], which aims at solving the multi-robot path planning using GNN. In contrast to this work, we consider smooth trajectories and continuous space instead of discrete space and decisions. This choice highlights the robustness and stability properties of GGNN when controlling robots with smooth inputs. The objective is to guide a group of robots, initially located at random positions, towards their individual targets in the cluttered space.

Dynamics and Expert Controller

We consider N agents described by the position $\mathbf{r}(t) \in \mathbb{R}^{N \times 2}$ in a single integrator dynamics with the velocity

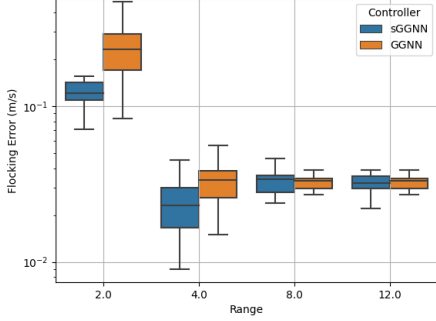
$\mathbf{u}(t) \in \mathbb{R}^{N \times 2}$ taken as the system input. As before, the learning algorithm is agnostic to the agent dynamics. For the expert controller, we used a combination of RRT [48] to find an obstacle-free path and MPC to control the agents in a continuous time. The MPC minimizes the divergence of the agents from the RRT-generated paths while constraining the agent motion to avoid inter-agent and obstacle collisions:

$$\begin{aligned} \min_{\mathbf{u}(t), \mathbf{r}(t)} \quad & \sum_{i=0}^N \|\mathbf{r}_{id}(t)\|_2^2 \\ \text{s.t.} \quad & \dot{\mathbf{r}}_i = \mathbf{u}_i \quad i = 1, \dots, N. \\ & g_o(\mathbf{r}_{io}(t)) \geq \mathbf{b}_o, \quad i = 1, \dots, N. \\ & g_{ij}(\mathbf{r}_{ij}(t)) \geq \mathbf{b}_{ij}, \quad i, j = 1, \dots, N. \\ & \|\mathbf{u}_i\|_2 \leq 1 \quad i = 1, \dots, N. \end{aligned}$$

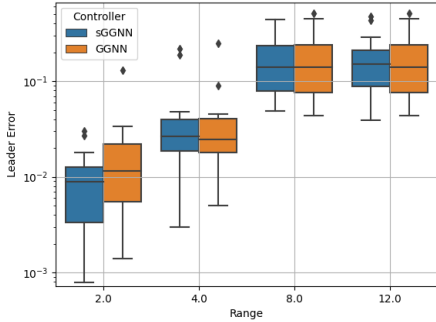
where \mathbf{r}_{id} is the distance from the RRT-path for the agent i , g_o, g_{ij} are the quadratic distances between the agent positions and the obstacles in the space (\mathbf{r}_{io}) and between the agents (\mathbf{r}_{ij}). As we can see, the MPC is naturally a centralized solution, since it uses all the agent dynamics to generate the control inputs.

Neural Network Architecture

As in the previous example, we assume the agents to form a communication graph if they are within a communication



(a) Flocking error



(b) Leader error

Fig. 3: Flocking Error and Leader Error for stable and non-stable GGNN controllers with variable communication range using Laplacian

radius of R . The input features vector $\mathbf{w}_i \in \mathbb{R}^{10}$ of the robot i for the designed neural network is

$$\mathbf{w}_i = \left[\begin{array}{c} \mathbf{r}_{id}, \sum_{o \in \mathcal{NS}_{io}} \frac{\mathbf{r}_{io}}{\|\mathbf{r}_{io}\|_2^4}, \sum_{o \in \mathcal{NS}_{io}} \frac{\mathbf{r}_{io}}{\|\mathbf{r}_{io}\|_2^2}, \\ \sum_{j \in \mathcal{NS}_i} \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^4}, \sum_{j \in \mathcal{NS}_i} \frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|_2^2} \end{array} \right] \quad (21)$$

where $\mathcal{NS}_i, \mathcal{NS}_{io}$ are respectively the set of the agents and obstacles in the sensing range R_C of the robot i . At time T , all this information is available to the agent i . Similarly to the previous case, we use a 2-layer MLP with 128 nodes for input and 2 layers with 128 nodes for readout to control the agent’s velocity. However, for testing deep-GGNN stability, we employ a 2-layer GGNN with 30 features in the hidden state and a filter of length $K=2$. The support matrix used here is the normalized Laplacian.

Training

We recorded 40 trajectories to build the dataset, further separated in three subsets of training, validation and test set using the proportion 70% – 10% – 20%, respectively. The trajectories are the results of RRT+MPC controller running with $N = 10$ agents randomly located in a square space of 20m \times 20m with 15% obstacles density. The agent targets are equally randomly located in the free space. We fixed the communication range to $R = 4$ m and the sensing one to $R_C = 1$ m. The training runs for 200 epochs with the

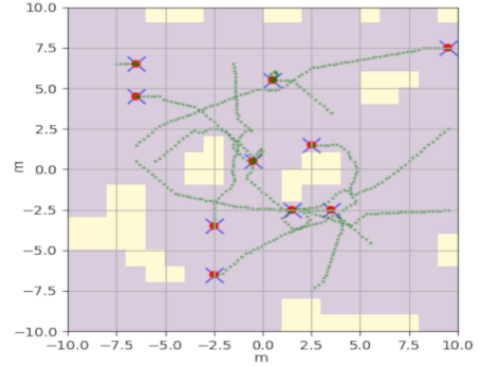


Fig. 4: Multi Robot Motion Control: a group of agents (red dots) move to reach their targets (blue cross) avoiding agent-agent and agent-obstacle (in yellow) collisions.

DAGGER algorithm executed every 20 epochs. We used the ADAM algorithm with learning rate $1e - 3$ and forgetting factors 0.9 and 0.999. The loss function used for imitation learning is the mean squared error between the output of the model and the expert control action.

Results

We evaluate the stability condition for the multi-robot motion control and report the results in Fig. 5. The comparison is carried out on three sets of experiments evaluating transferable at scale, robustness on communication range and obstacle density. We recorded 40 trajectories for each case on a 20m \times 20m map. For this application, we are primarily interested in the control success rate showed in Figs. 5b, 5d and 5f, where we consider successful the trajectories free of collisions and deadlocks. On successful trajectories, we also computed the travel time increase with respect to the expert controller. We obtained all the trajectories in the non-instantaneous communication setting and, when not stated otherwise, with team size, the communication range and the obstacle density respectively of $N = 10, R = 4.0$ m and 20%.

As in the flocking example, we first test the scalability of the two controllers to the team size (N) variations among [10, 30, 50]. sGGNN has a success rate that attests between 65% and 45%, showing better performances compared to the GGNN one, that is always below 30% (see Fig. 5b). The travel time increase is comparable for the two control solutions as we can see in the Figs. 5a, 5c and 5e, even if the stable one presents more situations of negative travel time increasing. This latter may have values below zeros due to the presence of the RRT in the expert controller that finds an obstacle-free but not the shortest path. Hence GGNNs control can lead to trajectories that are faster since they work with the target location directly. This phenomenon is less evident when the number of robots (Fig. 5a) or the obstacle density increases (Fig. 5e). The sGGNN does not show particular robustness to variations in the obstacle density since, as reported in Fig. 5f, the successful rate of 80% for 10% of obstacles drops to 45% for 30% of obstacles, the same success rate of the $N = 50$ and 20% of obstacles, even if it generalizes better than the non-stable learned controller as confirmed by the higher success rate and the lower average flow time with respect to the

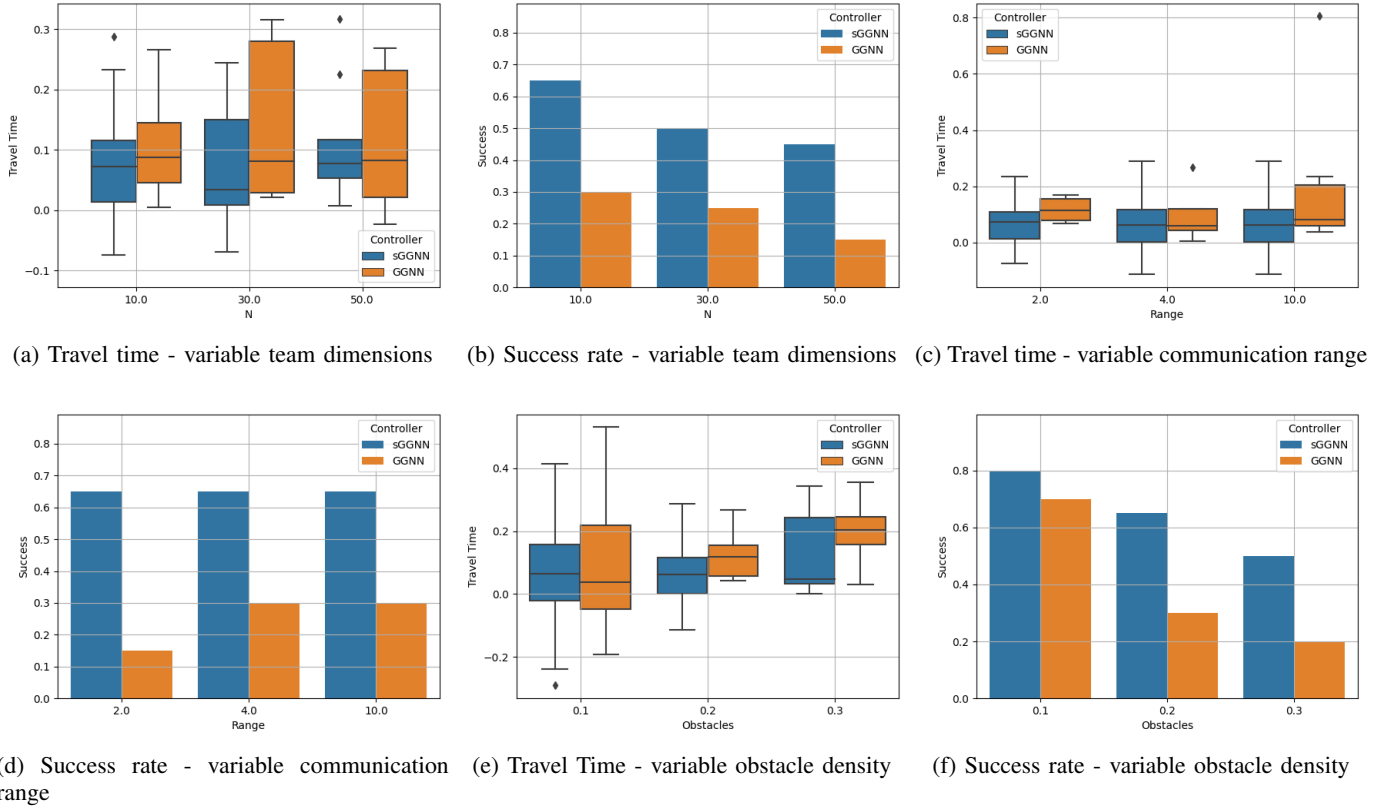


Fig. 5: Success rate and flow time for stable and non-stable GGNN controllers evaluation varying the team size, the communication range and the obstacles density for a $20m \times 20m$ map; the flow time increasing is computed as $(Tf - Tf^*)/Tf^*$ with Tf^* expert controller time of arrival.

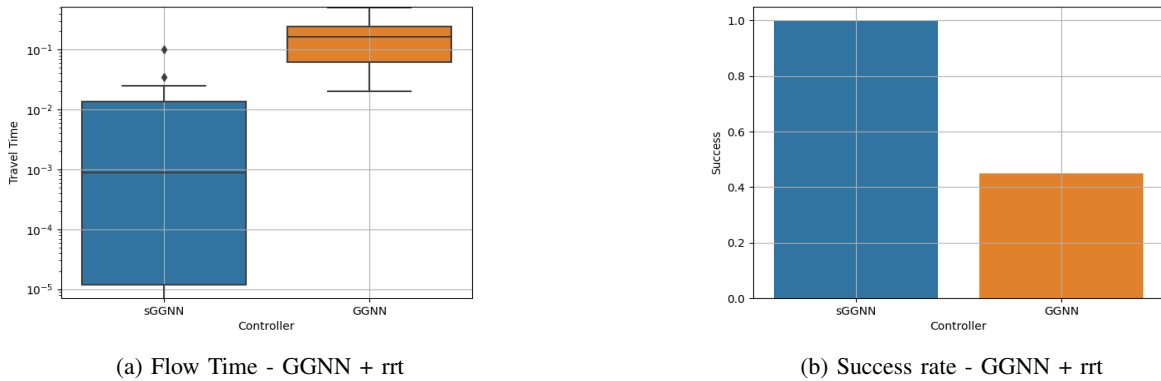


Fig. 6: Success rate and flow time for stable and non-stable GGNN controllers following a precomputed rrt path to avoid the obstacles. The controllers run on the map of $20m \times 20m$ with 30% of randomly generated obstacles and 50 agents. The flow time increasing is computed as $(Tf - Tf^*)/Tf^*$ with Tf^* expert controller time of arrival.

non-stable controller. This suggests a poor obstacle avoidance behaviour, further demonstrated by the results in Fig. 6b where we isolate failures caused by agent-agent collision by testing a combination of gated GNN and RRT. For an obstacle density of 30% and 50 agents, we feed the RRT path to neural models projecting it on the sensing range by replacing the target location in the input layer (21). In this case, sGGNN+RRT reaches 100% success and an average travel time of $10e-3$, confirming

that sGGNN realizes a well-distributed approximation of the MPC. On the contrary, GGNN+RRT still results in a high failure rate due to agent-agent collisions. Varying the communication range, as shown in Fig. 5d, produces comparable results, wherein sGGNN upholds a steady success rate. In contrast, GGNN encounters heightened instances of collisions among agents when communication ranges are reduced. This outcome underscores the resilience of sGGNN to fluctuations

in graph connectivity levels. When communication ranges are lower, robots tend to establish disjointed graphs more easily. Additionally, the performance remains unchanged since agents in close proximity form connected sub-graphs, allowing the neural network's internal state to quickly adapt and control the agent effectively under the new conditions.

VII. CONCLUSIONS

In this work, we devise sufficient conditions for the ISS and Incremental ISS of gated graph neural networks. When GGNN are used to learn distributed policies, the proposed stability conditions allow to guarantee that the trained networks enjoy the ISS/ δ ISS property, which is particularly useful during the synthesis of distributed controllers. The proposed condition has been tested on the flocking control and multi robot motion control, showing good modelling performances. Results suggest that enforcing stability properties on the learned controller makes it closer to the expert centralized one and more robust to parametric changes in a deployment scenario, such as communication radius and team size.

APPENDIX I

A. proof to Theorem 1

In the following, we report the proof of theorem (1).

Proof. Without loss of generality, we assume that the initial state belongs to the invariant set $\mathcal{X} = [-1, 1]^{N \times F}$. This is not a restrictive assumption since even if $\mathbf{x}(0) \notin \mathcal{X}$, at the next iteration it will be in \mathcal{X} due to the activation function σ_c . In light of the definitions given in (8), it holds that

$$\begin{aligned} \|\mathbf{x}^+\|_\infty &\leq \|\sigma_c(\hat{q} \circ A_S(\mathbf{x}) + \tilde{q} \circ B_S(\mathbf{u}) + \mathbf{b})\|_\infty \\ &\leq \|\hat{q} \circ A_S(\mathbf{x}) + \tilde{q} \circ B_S(\mathbf{u}) + \mathbf{b}\|_\infty \\ &\leq \sigma_{\hat{q}} \|S_K\|_\infty \|A\|_\infty \|\mathbf{x}\|_\infty + \\ &\quad \sigma_{\tilde{q}} \|S_K\|_\infty \|B\|_\infty \|\mathbf{u}\|_\infty + \|\mathbf{b}\|_\infty \\ &\leq \mathcal{A} \|\mathbf{x}\|_\infty + \mathcal{B} \|\mathbf{u}\|_\infty + \|\mathbf{b}\|_\infty \end{aligned} \quad (22)$$

From theorem (1), by iterating the (22) for t steps we get

$$\begin{aligned} \|\mathbf{x}(t)\|_\infty &\leq \mathcal{A}^t \|\mathbf{x}(0)\|_\infty + \sum_{k=1}^t \mathcal{A}^k (\mathcal{B} \|\mathbf{u}\|_\infty + \|\mathbf{b}\|_\infty) \\ \|\mathbf{x}(t)\|_\infty &\leq \mathcal{A}^t \|\mathbf{x}(0)\|_\infty + (1 - \mathcal{A})^{-1} \mathcal{B} \|\mathbf{u}\|_\infty + \\ &\quad (1 - \mathcal{A})^{-1} \|\mathbf{b}\|_\infty \end{aligned} \quad (23)$$

which proves the ISS property according to the definition II.1 \square

B. proof to Theorem 2

In the following, we report the proof of theorem 2

Proof. Given two states $\mathbf{x}_1, \mathbf{x}_2$, it holds that

$$\begin{aligned} \mathbf{x}_1^+ - \mathbf{x}_2^+ &= \sigma_c(\hat{q}_1 \circ A_{S1}(\mathbf{x}_1) + \tilde{q}_1 \circ B_{S1}(\mathbf{u}_1) + \mathbf{b}) \\ &\quad - \sigma_c(\hat{q}_2 \circ A_{S2}(\mathbf{x}_2) + \tilde{q}_2 \circ B_{S2}(\mathbf{u}_2) + \mathbf{b}). \end{aligned}$$

Owing to the Lipschitz assumption of σ_c and σ for Lipschitz constants respectively of 1 and $\frac{1}{4}$, the distance between the

two state trajectories is bounded by

$$\begin{aligned} \|\mathbf{x}_1^+ - \mathbf{x}_2^+\|_\infty &\leq \\ &\|(\hat{q}_1 \circ A_{S1}(\mathbf{x}_1) + \tilde{q}_1 \circ B_{S1}(\mathbf{u}_1)) - \\ &\quad (\hat{q}_2 \circ A_{S2}(\mathbf{x}_2) + \tilde{q}_2 \circ B_{S2}(\mathbf{u}_2))\|_\infty \leq \\ &\|(\hat{q}_1 \circ A_{S1}(\mathbf{x}_1) - \hat{q}_2 \circ A_{S2}(\mathbf{x}_2)) + \\ &\quad (\tilde{q}_1 \circ B_{S1}(\mathbf{u}_1) - \tilde{q}_2 \circ B_{S2}(\mathbf{u}_2))\|_\infty \leq \\ &\|\hat{q}_1 \circ (A_{S1}(\mathbf{x}_1) - A_{S2}(\mathbf{x}_2))\|_\infty + \\ &\|(\hat{q}_1 - \hat{q}_2) \circ A_{S2}(\mathbf{x}_2)\|_\infty + \\ &\|\tilde{q}_1 \circ (B_{S1}(\mathbf{u}_1) - B_{S2}(\mathbf{u}_2))\|_\infty + \\ &\|(\tilde{q}_1 - \tilde{q}_2) \circ B_{S2}(\mathbf{u}_2)\|_\infty. \end{aligned} \quad (24)$$

Different input features will correspond to different graph topologies and different support matrices S_1, S_2 ; thus different graph filters with the same parameters. For this reason, the differences of graph filters in the previous equation require further development. Let us focus on the state dependent graphs of the previous inequality. It holds that

$$\begin{aligned} \|\hat{q}_1 \circ (A_{S1}(\mathbf{x}_1) - A_{S2}(\mathbf{x}_2))\|_\infty &\leq \\ \sigma_{\hat{q}} \|S_{K1}(I_K \otimes \mathbf{x}_1)A - S_{K2}(I_K \otimes \mathbf{x}_2)A\|_\infty &\leq \\ \sigma_{\hat{q}} \|S_{K1}(I_K \otimes \mathbf{x}_1 - I_K \otimes \mathbf{x}_2) + \\ (S_{K1} - S_{K2})(I_K \otimes \mathbf{x}_2)\|_\infty \|A\|_\infty. \end{aligned} \quad (25)$$

Under the assumption of theorem 2, we have $\|S_{K1}\|_\infty, \|S_{K2}\|_\infty \leq \|\bar{S}_K\|_\infty$. In light of $\|\mathbf{x}\|_\infty \leq 1$, equation (25) becomes

$$\begin{aligned} \|\hat{q}_1 \circ (A_{S1}(\mathbf{x}_1) - A_{S2}(\mathbf{x}_2))\|_\infty &\leq \\ \sigma_{\hat{q}} (\|\bar{S}_K\|_\infty \|A\|_\infty \|\mathbf{x}_1 - \mathbf{x}_2\|_\infty + \\ (\|S_{K1} - S_{K2}\|_\infty) \|A\|_\infty). \end{aligned}$$

Applying the same reasoning to the other terms in the inequality (24), we obtain

$$\begin{aligned} \|\mathbf{x}_1^+ - \mathbf{x}_2^+\|_\infty &\leq \\ &(\sigma_{\hat{q}} \|\bar{S}\|_\infty \|A\|_\infty + \frac{1}{4} \|\bar{S}\|_\infty^2 \|\hat{A}\|_\infty \|A\|_\infty + \\ &\frac{1}{4} \|\bar{S}\|_\infty^2 \|\tilde{A}\|_\infty \|B\|_\infty) \|\mathbf{x}_1 - \mathbf{x}_2\|_\infty + \\ &(\sigma_{\tilde{q}} \|\bar{S}\|_\infty \|B\|_\infty + \frac{1}{4} \|\bar{S}\|_\infty^2 \|\hat{B}\|_\infty \|A\|_\infty + \\ &\frac{1}{4} \|\bar{S}\|_\infty^2 \|\tilde{B}\|_\infty \|B\|_\infty) \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + \\ &\mathcal{W}(\|S_{K1} - S_{K2}\|_\infty) \leq \\ &\mathcal{A}_\delta \|\mathbf{x}_1 - \mathbf{x}_2\|_\infty + \mathcal{B}_\delta \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + \\ &\mathcal{W}(\|S_{K1} - S_{K2}\|_\infty) \end{aligned}$$

where \mathcal{W} gathers all the coefficient multiplying the difference $\|S_{K1} - S_{K2}\|_\infty$. We can consider this latter as an additional bounded input, $\|S_{K1} - S_{K2}\|_\infty \leq \|\bar{S}_K\|_\infty - 1$ which, analogously to the input features, is defined by the team state. Hence, as stated in the theorem (2), it holds

$$\begin{aligned} \|\mathbf{x}_1(t) - \mathbf{x}_2(t)\|_\infty &\leq \mathcal{A}_\delta^t \|\mathbf{x}_1(0) - \mathbf{x}_2(0)\|_\infty + \\ (1 - \mathcal{A}_\delta)^{-1} \delta \mathcal{B} \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + \\ (1 - \mathcal{A}_\delta)^{-1} \mathcal{W} \|S_{K1} - S_{K2}\|_\infty. \end{aligned} \quad (26)$$

The state trajectories have then a maximum distance that is asymptotically bounded by a function monotonically increasing with the maximum distance between the input sequences for

$$\gamma_\delta = \left[(1 - \mathcal{A}_\delta)^{-1} \delta \mathcal{B} \quad (1 - \mathcal{A}_\delta)^{-1} \mathcal{W} \right] \left\| \begin{matrix} \mathbf{u}_1 - \mathbf{u}_2 \\ S_{K1} - S_{K2} \end{matrix} \right\|_\infty. \quad (27)$$

Therefore the system is incrementally ISS under the definition II.2. \square

C. proof of Theorem 4

Proof. To analyse the incremental stability it is useful to separate each layer. From the proof of theorem I-B, for the first layer we know that:

$$\begin{aligned} \|\mathbf{x}_1^{1+} - \mathbf{x}_2^{1+}\|_\infty &\leq \mathcal{A}_\delta^1 \|\mathbf{x}_1^1 - \mathbf{x}_2^1\|_\infty + \mathcal{B}_\delta^1 \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + \\ &\quad \mathcal{W}^1 \|S_{K1} - S_{K2}\|_\infty \end{aligned} \quad (28)$$

As a result, for the second it holds

$$\begin{aligned} \|\mathbf{x}_1^{2+} - \mathbf{x}_2^{2+}\|_\infty &\leq \\ &\mathcal{A}_\delta^2 \|\mathbf{x}_1^2 - \mathbf{x}_2^2\|_\infty + \mathcal{B}_\delta^2 \|\mathbf{x}_1^{1+} - \mathbf{x}_2^{1+}\|_\infty + \\ &\mathcal{W}^2 \|S_{K1} - S_{K2}\|_\infty \leq \\ &\mathcal{A}_\delta^2 \|\mathbf{x}_1^2 - \mathbf{x}_2^2\|_\infty + \mathcal{B}_\delta^2 \mathcal{A}_\delta^1 \|\mathbf{x}_1^1 - \mathbf{x}_2^1\|_\infty + \\ &\mathcal{B}_\delta^2 \mathcal{B}_\delta^1 \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + \\ &\mathcal{B}_\delta^2 \mathcal{W}^1 \|S_{K1} - S_{K2}\|_\infty + \mathcal{W}^2 (\|S_{K1} - S_{K2}\|_\infty \end{aligned} \quad (29)$$

Denoting $\Delta X = \begin{bmatrix} \mathbf{x}_1^1 - \mathbf{x}_2^1 & \dots & \mathbf{x}_1^M - \mathbf{x}_2^M \end{bmatrix}^T$, $\Delta U = \mathbf{u}_1 - \mathbf{u}_2$, $\Delta S_K = S_{K1} - S_{K2}$ iterating the same reasoning for M layers we get

$$\begin{aligned} \|\Delta X^+\|_\infty &\leq \begin{bmatrix} \mathcal{A}_\delta^1 & 0 & \dots & 0 \\ \mathcal{B}_\delta^2 \mathcal{A}_\delta^1 & \mathcal{A}_\delta^2 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ \mathcal{A}_\delta^1 \prod_{h=2}^M \mathcal{B}_\delta^h & \dots & \dots & \mathcal{A}_\delta^M \end{bmatrix} \|\Delta X\|_\infty + \\ &\begin{bmatrix} \mathcal{B}_\delta^1 \\ \mathcal{B}_\delta^2 \\ \vdots \\ \prod_{h=1}^M \mathcal{B}_\delta^h \end{bmatrix} \|\Delta U\|_\infty + \begin{bmatrix} \mathcal{W}^1 \\ \mathcal{W}^1 \mathcal{B}_\delta^2 + \mathcal{W}^2 \\ \vdots \\ \prod_{h=1}^{M-1} \mathcal{W}^h \mathcal{B}_\delta^{h+1} + \mathcal{W}^M \end{bmatrix} \|\Delta S\|_\infty \\ &\leq M_\delta \|\Delta X\|_\infty + M_{B\delta} \|\mathbf{u}_1 - \mathbf{u}_2\|_\infty + M_{W\delta} \|S_{K1} - S_{K2}\|_\infty \end{aligned} \quad (30)$$

Since the matrix M_δ is lower triangular, its eigenvalues are the elements of its diagonal \mathcal{A}_δ^i and are equal to the eigenvalues of the matrix M_δ . \square

REFERENCES

- [1] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.
- [2] A. Dorri, S. S. Kanhere, and R. Jurdak, "Multi-agent systems: A survey," *Ieee Access*, vol. 6, pp. 28 573–28 593, 2018.
- [3] J. Cortés and M. Egerstedt, "Coordinated control of multi-robot systems: A survey," *SICE Journal of Control, Measurement, and System Integration*, vol. 10, no. 6, pp. 495–503, 2017.
- [4] G.-B. Huang, L. Chen, C. K. Siew, *et al.*, "Universal approximation using incremental constructive feedforward networks with random hidden nodes," *IEEE Trans. Neural Networks*, vol. 17, no. 4, pp. 879–892, 2006.
- [5] S. Batra, Z. Huang, A. Petrenko, T. Kumar, A. Molchanov, and G. S. Sukhatme, "Decentralized control of quadrotor swarms with end-to-end deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 576–586.
- [6] C. He, Y. Wan, Y. Gu, and F. L. Lewis, "Integral reinforcement learning-based multi-robot minimum time-energy path planning subject to collision avoidance and unknown environmental disturbances," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 983–988, 2020.
- [7] J. K. Gupta, M. Egorov, and M. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *Autonomous Agents and Multiagent Systems: AAMAS 2017 Workshops, Best Papers, São Paulo, Brazil, May 8-12, 2017, Revised Selected Papers 16*. Springer, 2017, pp. 66–83.
- [8] N. Majcherczyk, N. Srishankar, and C. Pinciroli, "Flow-flt: Data-driven federated learning for spatio-temporal predictions in multi-robot systems," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8836–8842.
- [9] L. Wu, P. Cui, J. Pei, L. Zhao, and L. Song, "Graph neural networks," in *Graph Neural Networks: Foundations, Frontiers, and Applications*. Springer, 2022, pp. 27–37.
- [10] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" *arXiv preprint arXiv:1810.00826*, 2018.
- [11] F. Gama, E. Tolstaya, and A. Ribeiro, "Graph neural networks for decentralized controllers," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5260–5264.
- [12] Q. Li, W. Lin, Z. Liu, and A. Prorok, "Message-aware graph attention networks for large-scale multi-robot path planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5533–5540, 2021.
- [13] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8944–8950.
- [14] A. Khan, A. Ribeiro, V. Kumar, and A. G. Francis, "Graph neural networks for motion planning," *arXiv preprint arXiv:2006.06248*, 2020.
- [15] X. Ji, H. Li, Z. Pan, X. Gao, and C. Tu, "Decentralized, unlabeled multi-agent navigation in obstacle-rich environments using graph neural networks," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8936–8943.
- [16] Y. Zhou, J. Xiao, Y. Zhou, and G. Loianno, "Multi-robot collaborative perception with graph neural networks," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2289–2296, 2022.
- [17] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," *arXiv preprint arXiv:2209.12091*, 2022.
- [18] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020.
- [19] A. Davydov, A. V. Proskurnikov, and F. Bullo, "Non-euclidean contractivity of recurrent neural networks," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 1527–1534.
- [20] B. Song, J.-J. Slotine, and Q.-C. Pham, "Stability guarantees for continuous rl control," *arXiv preprint arXiv:2209.07324*, 2022.
- [21] H. Tsukamoto, S.-J. Chung, and J.-J. Slotine, "Learning-based adaptive control using contraction theory," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 2533–2538.
- [22] F. Bonassi, M. Farina, and R. Scattolini, "On the stability properties of gated recurrent units neural networks," *Systems & Control Letters*, vol. 157, p. 105049, 2021.
- [23] E. Terzi, F. Bonassi, M. Farina, and R. Scattolini, "Learning model predictive control with long short-term memory networks," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 18, pp. 8877–8896, 2021.
- [24] F. Bayer, M. Bürger, and F. Allgöwer, "Discrete-time incremental iss: A framework for robust nmpc," in *2013 European Control Conference (ECC)*. IEEE, 2013, pp. 2068–2073.
- [25] W. D'Amico, A. La Bella, and M. Farina, "An incremental input-to-state stability condition for a generic class of recurrent neural networks," *arXiv preprint arXiv:2210.09721*, 2022.

- [26] L. Ruiz, F. Gama, and A. Ribeiro, "Gated graph recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 68, pp. 6303–6318, 2020.
- [27] F. Gama, J. Bruna, and A. Ribeiro, "Stability properties of graph neural networks," *IEEE Transactions on Signal Processing*, vol. 68, pp. 5680–5695, 2020.
- [28] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, 2013.
- [29] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized path planning," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2020, pp. 1901–1903.
- [30] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 8944–8950.
- [31] A. Sandryhaila and J. M. Moura, "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1644–1656, 2013.
- [32] F. Gu, H. Chang, W. Zhu, S. Sojoudi, and L. El Ghaoui, "Implicit graph neural networks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 11 984–11 995, 2020.
- [33] D. Lukovnikov, J. Lehmann, and A. Fischer, "Improving the long-range performance of gated graph neural networks," *arXiv preprint arXiv:2007.09668*, 2020.
- [34] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2011, pp. 627–635.
- [35] Z.-P. Jiang and Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," *Automatica*, vol. 37, no. 6, pp. 857–869, 2001.
- [36] D. Angeli, "A Lyapunov approach to incremental stability properties," *IEEE Transactions on Automatic Control*, vol. 47, no. 3, pp. 410–421, 2002.
- [37] J. Jouffroy, "A simple extension of contraction theory to study incremental stability properties," in *2003 European Control Conference (ECC)*. IEEE, 2003, pp. 1315–1321.
- [38] F. Gama, Q. Li, E. Tolstaya, A. Prorok, and A. Ribeiro, "Synthesizing decentralized controllers with graph neural networks and imitation learning," *IEEE Transactions on Signal Processing*, vol. 70, pp. 1932–1946, 2022.
- [39] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, "Learning decentralized controllers for robot swarms with graph neural networks," in *Conference on robot learning*. PMLR, 2020, pp. 671–682.
- [40] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized path planning," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2020, pp. 1901–1903.
- [41] W. Yu, G. Chen, and M. Cao, "Distributed leader–follower flocking control for multi-agent dynamical systems with time-varying velocities," *Systems & Control Letters*, vol. 59, no. 9, pp. 543–552, 2010.
- [42] L. E. Beaver and A. A. Malikopoulos, "An overview on optimal flocking," *Annual Reviews in Control*, vol. 51, pp. 88–99, 2021.
- [43] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents part i: dynamic topology," in *42nd IEEE International Conference on Decision and Control (IEEE Cat. No. 03CH37475)*, vol. 2. IEEE, 2003, pp. 2016–2021.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [45] D. Le and E. Plaku, "Multi-robot motion planning with dynamics via coordinated sampling-based expansion guided by multi-agent search," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1868–1875, 2019.
- [46] J. Tordesillas and J. P. How, "Mader: Trajectory planner in multiagent and dynamic environments," *IEEE Transactions on Robotics*, vol. 38, no. 1, pp. 463–476, 2021.
- [47] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 785–11 792.
- [48] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed rrt: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 2997–3004.



Antonio Marino obtained his M.Sc. degree in robotics engineering at the University of Genova in 2020. A. Marino is currently a PhD candidate in the Rainbow team at IRISA/Inria Rennes working on learning-based multi-robot formation control.



Claudio Pacchierotti (SM'20) is a tenured researcher at CNRS-IRISA in Rennes, France, since 2016. He was previously a postdoctoral researcher at the Italian Institute of Technology, Genova, Italy. Pacchierotti earned his PhD at the University of Siena in 2014. He was Visiting Researcher in the Penn Haptics Group at Univ. Pennsylvania in 2014, the Dept. of Innovation in Mechanics and Management at Univ. Padua in 2013, the Institute for Biomedical Technology and Technical Medicine (MIRA) at University of

Twente in 2014, and the Dept. Computer, Control and Management Engineering of the Sapienza Univ. Rome in 2022. Pacchierotti received the 2014 EuroHaptics Best PhD Thesis Award and the 2022 CNRS Bronze Medal. He is Senior Chair of the IEEE Technical Committee on Haptics, Co-Chair of the IEEE Technical Committee on Telerobotics, and Secretary of the Eurohaptics Society.



Paolo Robuffo Giordano (M'08-SM'16) received his M.Sc. degree in Computer Science Engineering in 2001, and his Ph.D. degree in Systems Engineering in 2008, both from the University of Rome "La Sapienza". In 2007 and 2008 he spent one year as a PostDoc at the Institute of Robotics and Mechatronics of the German Aerospace Center (DLR), and from 2008 to 2012 he was Senior Research Scientist at the Max Planck Institute for Biological Cybernetics in Tübingen, Germany. He is currently a senior

CNRS researcher head of the Rainbow group at Irisa and Inria in Rennes, France.