With whom to communicate

Learning efficient communication for multi-robot collision avoidance

Serra-Gomez, Alvaro; Brito, Bruno; Zhu, Hai; Chung, Jen Jen; Alonso-Mora, Javier

DOI

[10.1109/IROS45743.2020.9341762](https://doi.org/10.1109/IROS45743.2020.9341762)

Publication date

2020

Document Version

Final published version

Published in

Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020

Citation (APA)

Serra-Gomez, A., Brito, B., Zhu, H., Chung, J. J., & Alonso-Mora, J. (2020). With whom to communicate: Learning efficient communication for multi-robot collision avoidance. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020* (pp. 11770-11776). IEEE . <https://doi.org/10.1109/IROS45743.2020.9341762>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

With Whom to Communicate: Learning Efficient Communication for Multi-Robot Collision Avoidance

Álvaro Serra-Gómez¹, Bruno Brito¹, Hai Zhu¹, Jen Jen Chung², Javier Alonso-Mora¹

Abstract—Decentralized multi-robot systems typically perform coordinated motion planning by constantly broadcasting their intentions as a means to cope with the lack of a central system coordinating the efforts of all robots. Especially in complex dynamic environments, the coordination boost allowed by communication is critical to avoid collisions between cooperating robots. However, the risk of collision between a pair of robots fluctuates through their motion and communication is not always needed. Additionally, constant communication makes much of the still valuable information shared in previous time steps redundant. This paper presents an efficient communication method that solves the problem of “when” and with “whom” to communicate in multi-robot collision avoidance scenarios. In this approach, every robot learns to reason about other robots’ states and considers the risk of future collisions before asking for the trajectory plans of other robots. We evaluate and verify the proposed communication strategy in simulation with four quadrotors and compare it with three baseline strategies: non-communicating, broadcasting and a distance-based method broadcasting information with quadrotors within a predefined distance.

I. INTRODUCTION

Being able to account for the planned path of other robots is of utmost importance for safe navigation in Micro Aerial Vehicle (MAV) environments. Centralized systems achieve this objective by having a central robot manage all of the robots’ information and plans. Instead, in decentralized systems robots estimate or communicate their teammates’ future trajectories. Common communication policies are broadcasting or distance-based communication of trajectory plans. However, much of this information becomes redundant or unnecessary when robot motions are clearly not intersecting. This is inefficient and sometimes unfeasible, especially in communication-restrictive environments such as underwater, outer space or for large groups of robots. In this work we focus on the following two issues: a) solving the problem of *when* and with *whom* to communicate and b) how to couple this communication policy with existing motion planning methods.

In this paper we propose an efficient communication policy method combined with an optimal control motion planner for multi-robot collision avoidance. The approach leverages the strengths of learning methods for decision-making and nonlinear receding horizon control, or Non-Linear Model

Predictive Control (NMPC) for multi-robot path planning. In particular, we use Multi-Agent Reinforcement Learning (MARL) to learn the robots’ communication policies. For every robot and time instance, the policy selects a set of other robots and requests their trajectory plans. Non-selected robots are assumed to follow a constant velocity trajectory or their previously communicated one. Then, we formulate a nonlinear optimization problem to generate a safe trajectory. The planned trajectory takes into account the requested and estimated trajectories represented as constraints in the receding horizon framework.

The main contributions of this work are:

- A combined communication policy and trajectory planning method for micro-aerial vehicles (MAVs), which utilizes the strengths of non-linear model predictive control (NMPC) and multi-agent reinforcement learning (MARL) to plan safe trajectories with minimal communications in three-dimensional scenarios.
- An on-line efficient communication policy that uses Multi-Agent Reinforcement Learning (MARL) to learn (off-line) *when* and with *whom* it is useful to communicate, performing collision avoidance while minimizing communication.

We evaluate our method with a team of quadrotors in simulated scenarios requiring different levels of communication for safe navigation and compare it to three other heuristic baseline methods. We show that our learning method enables the emergence of intuitive communication behaviours while maintaining the performance of broadcasting policies.

II. RELATED WORK

A. Communication in Collision Avoidance

There has been a large amount of work in multi-robot collision avoidance. One of the approaches is the reciprocal velocity obstacles (RVO) method [1]. From the basic RVO framework there are now several extensions: the optimal reciprocal collision-avoidance (ORCA) method [2] casts the problem into a linear programming formulation; the generalized RVO method [3] applies to heterogeneous teams of robots; and the ϵ -cooperative collision avoidance (ϵ CCA) method [4] accounts for the cooperation of nonholonomic robots. While these RVO-based methods are computationally efficient, the robot dynamics are not fully modeled and the robot motion is typically limited by only planning one time step ahead. These issues can be overcome by using a model predictive control (MPC) framework for collision-free trajectory planning. This includes the decentralized MPC [5],

¹Department of Cognitive Robotics, Delft University of Technology
{a.serragomez; bruno.debrito; h.zhu;
j.alonsomora}@tudelft.nl

²Autonomous Systems Lab, ETH Zurich
{jenjen.chung}@mavt.ethz.ch

This work is supported by the U.S. Office of Naval Research Global(ONRG) NICOP-grant N62909-19-1-2027

[6] assuming other robots are moving at a constant velocity, the distributed MPC [7] with on-demand collision avoidance, and the chance-constrained MPC [8] that accounts for robot localization and sensing uncertainties. In this paper, we study the multi-robot collision avoidance problem in the MPC-based framework.

Typically in multi-robot collision avoidance, robots are assumed to be able to observe other robots' positions and estimate their velocity using a filter. However, each robot's intentions and planned trajectories are not known by other robots. One approach to tackle this issue is to let each robot communicate its planned trajectory with every other robot in the team. Robot's can then update their own trajectories to be collision free with other robots' trajectory plans, e.g. as in these distributed MPC works [7], [8], [9]. While these methods can achieve efficient and safe collision avoidance, the communication burden across the team is huge, particularly when the number of robots is large and much of the communication between robots may be redundant and unnecessary. Without communicating trajectory plans, a robot can achieve collision avoidance by constraining its motion to be within a safe neighborhood computed using only other robots' current positions, e.g. the BVC method [10] and the B-UAVC method [11]. While these methods can guarantee collision avoidance without inter-robot communication, the planned robot motions are very conservative and inefficient. Alternatively, the decentralized MPC in [5], [6] employs a constant velocity model when predicting other robots' future trajectories. While communication among robots are not required, the planned robot motions are not safe, in particular when the robots are moving at a high speed [8].

B. Communication Scheduling

A lot of works tend to formulate the problem of efficient communication in a receding horizon fashion. Some methods formulate the problem as a decentralized version of a Markov Decision Process (Dec-MDP) [12] or Partially Observable MDP (Dec-POMDP) [13] and try to optimize a value function in which communications are penalized. Others, such as [14], choose to formulate a constrained optimization problem where communications must be directly minimized while still guaranteeing data flow throughout the network. These approaches need us to be able to directly quantify a priori the value of communication, which is what we are trying to avoid. Recent work [15] manages to avoid this by triggering communication whenever uncertainty over another agent's actions exceeds a threshold. Ultimately, however, receding horizon methods are limited by their prediction horizon and the need for evaluation heuristics, which can unintentionally bias the resulting communication processes. On the other hand, reinforcement learning methods may discover more general policies without the need for delicate hand-tuning.

C. Learning Methods for Coordination

One of the main issues of Multi-Agent Reinforcement Learning (MARL) is instability of the learning process

caused by the non-stationarity that results from having different interacting policies learning at the same time. In order to deal with this problem successfully, recent work on MARL [16] performs centralized training and decentralized execution. Such a method has been successfully applied in the field of non-communicating collision avoidance tasks [17], [18]. Regarding tasks that require communication, several works have been published recently. Many of them focus on learning what content should be shared among agents, be it in the form of a composition of binary signals [19] and predefined symbols [20], policy hidden layers [21], or by directly sharing parameters among agents [22]. The most relevant to our work additionally focus on learning policies that are able to appropriately choose *when* and with *whom* to communicate. Jiang and Lu [23] assign roles to every agent, making some of them in charge of organizing a common communication channel with their neighbours. However, regions where there is no agent with such a role are left without coordination capabilities. Instead, Das et al. [24] present an end-to-end MARL algorithm that creates an attention module which chooses whom to establish bilateral communications with. Similarly, the method we present in this paper can also be considered as an attention module targeting other agents. However, we set our communications to be unilateral to promote asymmetrical behaviour. Additionally, we decouple the problem of communication and motion planning, allowing the combination of our method with existing and well-tested solutions for motion planning in collision avoidance tasks.

III. PRELIMINARIES

In this paper, we address the problem of deciding when and with whom to communicate during a multi-robot collision avoidance task. Though the proposed formulation is intended to be general, we are inspired by the results obtained in [25], which show how in a collision-avoidance scenario, methods that incorporate communication have a clear advantage over those that do not. We approach the information-sharing process as a MARL problem where the robots must learn to request information effectively. In this section, we set the context for our targeted communication process by providing an overview of the Non-Linear Model Predictive Control method used for motion control, as well as our MARL framework, introducing relevant notation for this work.

A. Multi-Robot Collision Avoidance

Consider a team of n robots moving in a shared workspace $\mathcal{W} \subseteq \mathbb{R}^3$, where each robot $i \in \mathcal{I} = \{1, 2, \dots, n\} \subset \mathbb{N}$ is modeled as an enclosing sphere with radius r . The dynamics of each robot $i \in \mathcal{I}$ are described by a discrete-time equation as follows,

$$\mathbf{x}_i^{k+1} = \mathbf{f}(\mathbf{x}_i^k, \mathbf{u}_i^k), \quad \mathbf{x}_i^0 = \mathbf{x}_i(0), \quad (1)$$

where $\mathbf{x}_i^k \in \mathcal{X} \subset \mathbb{R}^{n_x}$ denotes the state of the robot, typically including its position $\mathbf{p}_i^k$ and velocity $\mathbf{v}_i^k$, and $\mathbf{u}_i^k \in \mathcal{U} \subset \mathbb{R}^{n_u}$ the control inputs at time k . $\mathcal{X}$ and $\mathcal{U}$ are the admissible

state space and control space respectively. $\mathbf{x}_i(0)$ is the initial state of robot i . Any pair of robots i and j from the group are mutually collision-free if $\|\mathbf{p}_i^k - \mathbf{p}_j^k\| \geq 2r, \forall i \neq j \in \mathcal{I}, \forall k = 0, 1, \dots$. Each robot has a given goal location $\mathbf{g}_i$, which generally comes from some high-level path planner or is specified by some user.

The objective of multi-robot collision avoidance is to compute a local motion $\mathbf{u}_i^k$ for each robot in the group, that respects its dynamics constraints, makes progress towards its goal location $\mathbf{g}_i$ and is collision-free with other robots in the team for a short time horizon.

B. Distributed Model Predictive Control

The key idea of using distributed model predictive control to solve the multi-robot collision avoidance problem is to formulate it as a receding horizon constrained optimization problem. For each robot $i \in \mathcal{I}$, a discrete-time constrained optimization formulation with N time steps and planning horizon $\tau = N\Delta t$, where Δt is the sampling time, is derived as follows,

$$\begin{aligned} \min_{\mathbf{x}_i^{1:N}, \mathbf{u}_i^{0:N-1}} \quad & \sum_{k=0}^{N-1} J_i^k(\mathbf{x}_i^k, \mathbf{u}_i^k) + J_i^N(\mathbf{x}_i^N, \mathbf{g}_i) \\ \text{s.t.} \quad & \mathbf{x}_i^0 = \mathbf{x}_i(0), \\ & \mathbf{x}_i^k = \mathbf{f}(\mathbf{x}_i^{k-1}, \mathbf{u}_i^{k-1}), \\ & \|\mathbf{p}_i^k - \mathbf{p}_j^k\| \geq 2r, \\ & \mathbf{u}_i^{k-1} \in \mathcal{U}, \quad \mathbf{x}_i^k \in \mathcal{X}, \\ & \forall j \neq i \in \mathcal{I}; \forall k \in \{1, \dots, N\}. \end{aligned} \quad (2)$$

At each time step, each robot in the team solves online the constrained optimization problem (2) and then executes the first step control inputs, in a receding-horizon fashion.

C. With Whom to Communicate

Note that for each robot to solve problem (2), it has to know the future trajectory of other robots in the team. At time t , denote by $\hat{\mathcal{T}}_{i,j}^t = \{\mathbf{p}_j^{t+1:t+N}\}_t$ the trajectory of robot $j \in \mathcal{I}, j \neq i$ that robot i assumes and uses in solving the problem (2), where the hat indicates that it is what robot i knows about the other agent's trajectory. Further denote by $\mathcal{T}_i^t = \{\mathbf{p}_i^{t+1:t+N}\}_t$ the trajectory for robot i planned at time t . Typically, there are two methods for robot i to obtain the future trajectory information of other robots j :

- **Without communication:** robot i predicts another robot's future trajectory based on their current states, that is

$$\hat{\mathcal{T}}_{i,j}^t = \text{prediction}(\mathbf{x}_j^t), \quad \forall j \neq i \in \mathcal{I}. \quad (3)$$

In this paper, each robot employs a constant velocity model for the prediction as described in [5].

- **Full communication:** All robots in the team communicate their planned trajectories to each other at each time step, that is

$$\hat{\mathcal{T}}_{i,j}^t = \mathcal{T}_j^{t-\Delta t}, \quad \forall j \neq i \in \mathcal{I}. \quad (4)$$

Both of the two methods have their advantages and disadvantages. While the full communication method can achieve safe collision avoidance, it requires a large amount of communication among robots. However, if there is no communication, the robot may plan an unsafe trajectory if its prediction on other robots' trajectories deviates from their real ones.

Motivated by these facts, this paper aims to solve the problem of "with whom to communicate" for each robot in the team for collision avoidance. More precisely, at each time step, each robot i decides whether or not to request a communication message from every other robot j . If robot i decides to request robot j , robot j should communicate its planned trajectory to robot i . If robot i decides not to request robot j , it predicts robot j 's future trajectory based on its observed current state of robot j .

Denote by $\pi_i^t = \{c_{i,j}^t | j = 1, \dots, n\}$ the communication vector of robot i at time t , in which $c_{i,j}^t = 1$ indicates that robot i requires a communicated trajectory from robot j . Otherwise $c_{i,j}^t = 0$. Note that $c_{i,i}^t = 0$ since the robot does not need to communicate with itself. Let $\pi^t = \{\pi_1^t; \dots; \pi_n^t\}$ be the communication matrix of the multi-robot system at time t . We define the communication cost of the system to be

$$C(\pi^t) = \sum_i^n \sum_j^n c_{i,j}^t. \quad (5)$$

The objective of this paper is to find a policy for each robot i ,

$$\pi_i^t = \pi_i(\mathbf{x}_1^t, \mathbf{x}_2^t, \dots, \mathbf{x}_n^t), \quad (6)$$

that minimizes $C(\pi^t)$ while ensuring that the robots are collision-free with each other in the system.

IV. METHOD

A. Overview

An overview of the proposed method is given in Fig.1. It consists of two components: a communication policy and a constrained MPC planner.

Every time step, based on its partial observation of the current joint state $\mathbf{z}_i^t$, every robot targets a set of other robots π_i^t and requests their trajectory plans $\hat{\mathcal{T}}_{i,j}^t = \mathcal{T}_j^{t-\Delta t}$ according to a learnt parametric policy $\pi_{i,\theta_i}(z_i^t)$. Those robots not targeted are estimated to follow a previously communicated trajectory or, in case it is no longer useful, a constant velocity model $\hat{\mathcal{T}}_{i,j}^t = \text{prediction}(\mathbf{x}_j^t)$ as described in Section III.C.

A receding horizon optimization is then employed to plan the future trajectory $\mathcal{T}_i^t$ for robot i . To guarantee the safety of such a trajectory, the resulting trajectory is constrained to not intersect with the requested and estimated trajectories. The first action input from the computed plan is applied.

B. Multi-robot Reinforcement Learning

We formulate a multi-robot reinforcement learning problem to compute an efficient communication policy. By considering the optimization process as part of the transition model, this problem can be transformed into a decentralized POMDP [26]. The decentralized POMDP is composed of six

Fig. 1: Schema of the proposed method for efficient communication. $\pi_{i,\theta}^t(z_i)$ is the communication policy dependent on the observation z_i . $\mathcal{T}_j^{t-\Delta t}$ is the trajectory plan of robot j at the previous time step. And $\hat{\mathcal{T}}_{i,j}^t$ are the combination of obtained and estimated trajectories of the other robots.

components, including state space, action space, observation space, transition model, observation model.

1) *State Space* $\mathcal{X}$: For every robot i , $x^i \in \mathcal{X}$ must account for the current physical state, its sequence of intended actions from the previous time step, as well as any information it has of other robots on their positions, velocities and action sequence intentions. Therefore, the state at time t can be defined as:

$$x_i^t := [p_i^t, v_i^t, \{u_i^{k'}\}_{k'=t-1, \dots, t+K-1}, \pi_i^{t-1}], \quad (7)$$

$$X^t := \{x_1^t, x_2^t, \dots, x_n^t\}, \quad (8)$$

where $p_i^t \in \mathbb{R}^3$ and $v_i^t \in \mathbb{R}^3$ are the position and velocities of robot i at time t , and $\{u_i^{k'}\}_{k'=t-1, \dots, t+K-1}$ is the action sequence planned for a K -time-window at the previous time step $t-1$ by robot i . $\pi_i^{t-1} \in \{0, 1\}^{n-1}$ is the binary vector indicating whether robot i has requested any other robot $j \neq i$ trajectory intentions at time step $t-1$. Then, X^t is the joint state of the whole multi-robot system.

2) *Observation Space* $\mathcal{Z}$: We assume each robot can always observe the positions and velocities of all other robots and knows the position of its goal through its sensors. For robot i , partial observations on the joint state at time t are defined as follows:

$$z_i^t = [v_i^t, p_{i,g}^t, \{p_{i,j}^t\}_{j \in I \setminus i}, \{v_{i,j}^t\}_{j \in I \setminus i}], \quad (9)$$

where $p_{i,j}^t$ and $v_{i,j}^t$ are the relative positions and velocities of the other robots with respect to the i^{th} robot, and $p_{i,g}^t$ is the relative position of robot i 's goal. The joint observation from all robots is denoted by $z = \{z^1, \dots, z^n\} \in \mathcal{Z}$.

3) *Action Space* $\mathcal{A} = \times_{i \in I} \mathcal{A}_i$: As it has already been introduced in Section III.C, we denote by $\pi_i^t = \{c_{i,j}^t | j \neq i\}$ the communication vector of robot i at time t . Note we have dropped the i^{th} element as the robot cannot communicate with itself. Therefore the action space for robot i is:

$$\mathcal{A}_i = \{0, 1\}^{n-1}$$

4) *Reward* $R(x^t, \pi^t)$: The reward function is chosen based on the behaviors we want to learn. It aims for the learned communication policy to communicate as little as possible while allowing each robot on the team to reach its goal and avoid collisions. The reward value $R(x^t, \pi^t)$ is the immediate reward that all robots get at a state $x \in \mathcal{X}$ after

applying the communication matrix π^t . All robots getting the same reward accounts for indirect interactions e.g. robot i colliding with another robot j' (whose trajectory was not requested) because of trying to avoid the trajectory plan of robot j . The reward function is composed of the following weighted combination of terms:

$$R(x^t, \pi^t) = w_g R_g(x^t) + w_{coll} R_{coll}(x^t) + w_c R_c(\pi^t) \quad (10)$$

where

$$R_g(x^t) = \begin{cases} r_g & \forall i \in I, \|p_{i,g}^t - p_i^t\| \leq r_i \\ 0 & \text{otherwise} \end{cases}$$

where $r_g > 0$ is a tuned reward given for every time step that all robots are within its goal, r_i is the radius of the smallest sphere containing the robot. The sooner all robots reach their destination, the more reward they receive during the episode, not only encouraging collision avoidance but also to reach the goal quickly.

$$R_{coll}(x^t) = \begin{cases} -r_{coll} & \forall i, j \in I, i \neq j, \|p_i^t - p_j^t\| \leq r_i + r_j \\ 0 & \text{otherwise} \end{cases}$$

where $r_{coll} > 0$ is a tuned penalty term for the collision between any two robots.

Finally the global penalization term for path plan requests has been introduced before in Section III.C and has the form:

$$R_c(\pi^t) = -C(\pi^t) = -\sum_i^n \sum_j^n c_{i,j}^t.$$

5) *Observation Model* $\mathcal{O}(z^{t+1}, x^{t+1}, \pi^t)$: We assume that every robot i can directly observe the positions and velocities of other robots. The main uncertainty lies in their trajectory plans and their communication matrix π_j^t .

6) *Transition model* $T(x^{t+1}, \pi^t, x^t)$: The transition model can be decomposed into a communication step and a physical action step:

$$p(x^{t+1} | \pi^t, x^t) = p(x^{t+1} | u^t, x^t) p(u^t | \pi^t, x^t) \quad (11)$$

where u^t are the control actions applied at time step t , which are obtained from the motion planner. $p(u^t | \pi^t, x^t)$ models the effects of communication π^t on the constrained optimization problem used to compute actions u^t . Then, $p(x^{t+1} | u^t, x^t)$ is the state transition for every robot. The robots employed in this paper are quadrotors, thus the state transition can be interpreted as the quadrotor model explained in Sec.III.A.

C. Multi-Agent Actor-Critic

In order to find a policy π_θ maximizing a cost function $J(\theta) = \mathbb{E}_{x \sim p_{\pi_\theta}, a \sim \pi_\theta(x)} [R(x, a)]$ where $a \in \mathcal{A}$, Policy Gradient methods (PG) [27] directly adjust the parameters θ of the policy by taking steps in the direction of the gradient of J with respect to the policy parameters:

$$\nabla_\theta J(\theta) = \mathbb{E}_{x \sim p_{\pi_\theta}, a \sim \pi_\theta(x)} [\nabla_\theta \log \pi_\theta(a | x) Q_\pi(x, a)] \quad (12)$$

where $Q_\pi(x', a') = \mathbb{E}_{x \sim p_{\pi_\theta}, a \sim \pi_\theta(x)} [R(x, a) | x', a']$ is the expected value from the total expected return conditioned

on taking action a' at state x' and follow policy π from then onwards. Actor-Critic methods are a family of algorithms that learn an approximation of the Q-function using deep neural networks. In particular, Deep Deterministic Policy Gradient (DDPG) algorithms [28] are a variant of off-policy Actor-Critic methods that learn deterministic policies $\pi_\theta(a|x) = \mu_\theta(x)$ instead of stochastic ones and use another policy β to explore the state-action space. Due to this modification, the expression of the gradient changes to the following:

$$\nabla_{\theta^\mu} J = \mathbb{E}_{\mathbf{x}^t \sim \rho^\beta} [\nabla_a Q(\mathbf{x}, a | \theta^Q) |_{\mathbf{x}=\mathbf{x}^t, a=\mu(\mathbf{x}^t)} \nabla_{\theta^\mu} \mu(\mathbf{x} | \theta^\mu) |_{\mathbf{x}=\mathbf{x}^t}]. \quad (13)$$

The algorithm we are using in our approach is the natural extension of DDPG to multi-agent environments, that is, the Multi-Agent Deep Deterministic Policy Gradient algorithm (MADDPG) [16]. While very similar to DDPG, MADDPG proposes for every agent to learn a decentralized policy with partial observations, while using its own centralized action-value function for learning using global state information, once again changing the expression for the gradient:

$$\nabla_{\theta_i} J(\mu_i) = \mathbb{E}_{\mathbf{z}^t \sim \rho^\mu} [\nabla_{a_i} Q_i^\mu(\mathbf{z}, a_1, a_2, \dots, a_N) |_{a_i=\mu_i(\mathbf{z}_i)} \nabla_{\theta_i} \mu(\mathbf{z}_i | \theta_i)], \quad (14)$$

where $\mathbf{z}$ consists of the observations of all agents $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_N)$. This technique allows us to cope with the non-stationarity resulting from having agents simultaneously learning interacting policies.

For every robot i , we want to learn a decentralized communication policy that targets other robots whose path plan is useful based on current observations. Thus, the policy of robot i at time step t will follow the expression:

$$\pi_{i,\theta_i}(\mathbf{z}_i) = \{1[f_{\theta_i}^{i,j}(\mathbf{z}_i) > \delta]\}_{j \in I \setminus i} = \pi_i \in \{0, 1\}^{N-1}, \quad (15)$$

where $\pi_{i,\theta_i}(\mathbf{z}_i)$ is the parameterized communication policy, $f_{\theta_i}^{i,j}$ is the parameterized function (e.g. neural network) mapping partial observations from robot i to a communication score between $[0, 1]$, δ is the threshold discriminating whether robot i should request robot j 's path intentions and 1 is the indicator function. The value of δ can be chosen as a hyperparameter. To learn the communication policy using this method, careful consideration must be given to the exploration of the state-action space. This policy does not directly influence the state transitions, which makes the task of finding a good combination of communication requests through the episode complex. While it might pose a risk to use them in tasks where safety is a requirement, stochastic policies allow good exploration at training time making the task of finding efficient communication behaviors easier.

The method used to learn this policy, MADDPG [16], is off-policy which means we can learn a deterministic target policy while using a stochastic policy at training time to encourage exploration. To do so, we substitute the preset threshold by a $\delta' \sim U(0, 1)$, which we sample at every time step. This enables us to apply exploration coherently as all

scores attributed to other robots by robot i during the same time step will be subject to the same sampled threshold. Requests from robot i to robot j following the exploration policy β are denoted by:

$$\beta_{i,\theta_i}(c_{i,j}^t | z_i^t) = \mathbb{P}_{\theta_i}(c_{i,j}^t = 1 | z_i^t) = \mathbb{P}(f_{\theta_i}^{i,j}(z_i^t) > \delta'). \quad (16)$$

V. RESULTS

In this section we describe our implementation of the proposed method and evaluate it in simulation.

A. Simulation Setup

The simulation environment and NMPC controller were implemented in Matlab. We rely on the solver Forces Pro [29] to generate optimized NMPC code. The learning algorithm was implemented in Python and ROS as middle-ware to connect both simulator and learning method. The Critic and Actor models are parameterized by two fully connected layers with 64 units and ReLU activation and were trained for 10000 episodes in an Intel i7 CPU@2.6GHz computer. We use the same hyperparameters reported in [16] for training except for $\gamma = 0.98$. Values for the reward weights were $w_g = 1$, $w_{coll} = 1$, $w_c = 0.1$. Tuned reward and penalty terms were $r_g = 1.3$, $r_{coll} = 150$. Episodes finished after reaching a collision or 100 time steps.

B. Training Environment

We have created a simulation environment where a group of four drones navigate from an initial position to a goal position and must communicate their trajectory plans to perform collision avoidance. We have designed four different scenarios to evaluate our communication policy, as depicted in Fig. 2. Each scenario has a different level of difficulty for the drone swarm to perform collision avoidance, ranging from a simple scenario where no communication is needed (e.g., Fig. 2a) to highly complex scenarios where the drones must communicate (e.g., Fig. 2b) to successfully avoid each other. The employed scenarios are the following:

- **Random:** Each drone must move to a random goal position. To ensure collision avoidance the drones must communicate their trajectory plans when crossing the path of another drone.
- **Random swapping:** each drone is requested to move to another drone's initial position.
- **Asymmetric swapping:** In this scenario we split the $\mathbb{R}^2$ x-y plane into four quadrants and initialize each drone in a different quadrant with random initial position. Then, each drone swaps positions with a drone from the diametrically opposed quadrant. If the drones do not communicate, collision is highly likely to occur.

Figure 3 shows how the number of collisions varies per scenario considering a *full-communication* and *no-communication* policy. Depending on which scenario the agents are trained in a different communication policy may be learned. For instance, if an agent is trained only on the first scenario it will learn a *no-communication* policy. In contrast, if only trained in the last it may learn to always communicate.

Fig. 2: Simulation results for each scenario using our communication policy. Solid lines represent the trajectories executed by the drone-swarm. Yellow represents the positions where the drones communicate their trajectory plans. Blue depicts the positions where the drones do not communicate. Green and Red represent the initial and goal position of each drone, respectively. Increasing opacity represent the episode progression.

Fig. 3: Collision avoidance performance for each scenario of the baseline methods vs. our learned-policy.

Fig. 4: Number of communication requests for each scenario of the baseline methods vs. our learned-policy.

Hence, we employ *curriculum learning* [30], training the agents first in a simple scenario, where no communication is needed, and continuously introducing more difficult and complex scenarios where the agents must learn *when to communicate*.

C. Performance Evaluation

We compare our learned communication policy with two baseline approaches:

- *Full communication* (FC): At each time-step each drone

broadcasts its trajectory plans.

- *No communication* (NC): The drones never exchange their trajectory plans and a Constant Velocity model is used by each drone to infer the others trajectories.
- *A distance-based communication policy* (ϵ -DBCP): If the distance between two agents distance is smaller than a threshold ϵ (in meters) then the agents broadcast their trajectory information.

Fig. 4 shows the number of collisions per scenario for each communication policy. In the first two scenarios the number of collisions is zero for any baseline. This demonstrates that for these two scenarios the simplified constant velocity model is enough and no-communication is required. In contrast, for the Random and Asymmetric scenarios the number of collisions raises significantly if the drones do not communicate. Yet, our learned policy achieved zero collisions in all scenarios. Moreover, Fig. 4 demonstrates that our policy reduced the number of communications requests approximately 77% while ensuring collision avoidance. In comparison with a 4-DBCP policy our method was able to reduce approximately 40% the number of communications requests and the number of 11% collisions to zero. Finally, Fig. 2 depicts the drone-swarm trajectories for each scenario. We can observe that our learned policy triggers communication either in the beginning of the motion or when the drones are in collision course.

VI. CONCLUSIONS

In this paper, we have introduced an effective communication policy integrating the strengths of MARL and NMPC in collision avoidance tasks. Simulation results show that our policy learns “when” to request other agents to perform collision avoidance. Furthermore, our method reduces the amount of communication requests significantly while ensuring collision-free motions. Future work will seek to scale our approach to a higher and variable number of agents and perform experimental results.

REFERENCES

- [1] J. van den Berg, Ming Lin, and D. Manocha, “Reciprocal velocity obstacles for real-time multi-agent navigation,” in *2008 IEEE Interna-*

- tional Conference on Robotics and Automation, vol. 48, no. 1. IEEE, 2008, pp. 1928–1935.
- [2] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, “Reciprocal n-body collision avoidance,” in *Springer Tracts in Advanced Robotics*, 2011, vol. 70, pp. 3–19.
 - [3] D. Bareiss and J. van den Berg, “Generalized reciprocal collision avoidance,” *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1501–1514, 2015.
 - [4] J. Alonso-Mora, P. Beardsley, and R. Siegwart, “Cooperative collision avoidance for nonholonomic robots,” *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 404–420, 2018.
 - [5] M. Kamel, J. Alonso-Mora, R. Siegwart, and J. Nieto, “Robust collision avoidance for multiple micro aerial vehicles using nonlinear model predictive control,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 236–243.
 - [6] S. H. Arul and D. Manocha, “Dcad: Decentralized collision avoidance with dynamics constraints for agile quadrotor swarms,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1191–1198, 2020.
 - [7] C. E. Luis, M. Vukosavljev, and A. P. Schoellig, “Online trajectory generation with distributed model predictive control for multi-robot motion planning,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 604–611, 2020.
 - [8] H. Zhu and J. Alonso-Mora, “Chance-constrained collision avoidance for mavs in dynamic environments,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.
 - [9] L. Ferranti, R. R. Negenborn, T. Keviczky, and J. Alonso-Mora, “Coordination of multiple vessels via distributed nonlinear model predictive control,” in *2018 European Control Conference (ECC)*. IEEE, 2018, pp. 2523–2528.
 - [10] D. Zhou, Z. Wang, S. Bandyopadhyay, and M. Schwager, “Fast, on-line collision avoidance for dynamic vehicles using buffered voronoi cells,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1047–1054, 2017.
 - [11] H. Zhu and J. Alonso-Mora, “B-uavc: Buffered uncertainty-aware voronoi cells for probabilistic multi-robot collision avoidance,” in *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2019, pp. 162–168.
 - [12] M. Roth, R. Simmons, and M. Veloso, “Reasoning about joint beliefs for execution-time communication decisions,” Tech. Rep., 2005.
 - [13] R. Becker, A. Carlin, V. Lesser, and S. Zilberstein, “Analyzing myopic approaches for multi-agent communication,” *Computational Intelligence*, vol. 25, no. 1, pp. 31–50, 2009.
 - [14] A. Kassir, R. Fitch, and S. Sukkarieh, “Communication-efficient motion coordination and data fusion in information gathering teams,” in *IEEE International Conference on Intelligent Robots and Systems*, vol. 2016–November. Institute of Electrical and Electronics Engineers Inc., nov 2016, pp. 5258–5265.
 - [15] G. Best, M. Forrai, R. R. Mettu, and R. Fitch, “Planning-Aware Communication for Decentralised Multi-Robot Coordination,” in *Proceedings - IEEE International Conference on Robotics and Automation*. Institute of Electrical and Electronics Engineers Inc., sep 2018, pp. 1050–1057.
 - [16] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, “Multi-agent actor-critic for mixed cooperative-competitive environments,” in *Advances in Neural Information Processing Systems*, vol. 2017–Decem, 2017, pp. 6380–6391.
 - [17] M. Everett, Y. F. Chen, and J. P. How, “Motion Planning among Dynamic, Decision-Making Agents with Deep Reinforcement Learning,” in *IEEE International Conference on Intelligent Robots and Systems*, 2018, pp. 3052–3059. [Online]. Available: [{_}ros](https://github.com/mfe7/cadrl)
 - [18] —, “Collision avoidance in pedestrian-rich environments with deep reinforcement learning,” *ArXiv*, vol. abs/1910.11689, 2019.
 - [19] J. N. Foerster, Y. M. Assael, N. De Freitas, and S. Whiteson, “Learning to communicate with deep multi-agent reinforcement learning,” in *Advances in Neural Information Processing Systems*. NEURAL INFO PROCESS SYS F, may 2016, pp. 2145–2153. [Online]. Available: <https://dl.acm.org/citation.cfm?id=3157336><http://arxiv.org/abs/1605.06676>
 - [20] I. Mordatch and P. Abbeel, “Emergence of grounded compositional language in multi-agent populations,” *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 1495–1502, 2018.
 - [21] S. Sukhbaatar, A. Szlam, and R. Fergus, “Learning multiagent communication with backpropagation,” *Advances in Neural Information Processing Systems*, no. Nips, pp. 2252–2260, 2016.
 - [22] J. K. Gupta, M. Egorov, and M. Kochenderfer, “Cooperative Multi-agent Control Using Deep Reinforcement Learning,” Tech. Rep., 2017.
 - [23] J. Jiang and Z. Lu, “Learning attentional communication for multi-agent cooperation,” in *Advances in Neural Information Processing Systems*, 2018.
 - [24] A. Das, T. Gervet, J. Romoff, D. Batra, D. Parikh, M. Rabbat, and J. Pineau, “TarMAC: Targeted multi-agent communication,” in *36th International Conference on Machine Learning, ICML 2019*, 2019.
 - [25] H. Zhu and J. Alonso-Mora, “Chance-Constrained Collision Avoidance for MAVs in Dynamic Environments,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.
 - [26] D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein, “The complexity of decentralized control of markov decision processes,” *Mathematics of Operations Research*, vol. 27, 12 2002.
 - [27] R. Sutton, D. Mcallester, S. Singh, and Y. Mansour, “Policy gradient methods for reinforcement learning with function approximation,” *Adv. Neural Inf. Process. Syst*, vol. 12, 02 2000.
 - [28] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *CoRR*, 09 2015.
 - [29] A. Domahidi and J. Jerez, “Forces professional. embotech gmbh (<http://embotech.com/forces-pro>),” 2014.
 - [30] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning,” in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.