ROS2SWARM - A ROS 2 Package for Swarm Robot Behaviors

Tanja Katharina Kaiser,¹ Marian Johannes Begemann,¹ Tavia Plattenteich,¹ Lars Schilling,² Georg Schildbach,³ and Heiko Hamann¹

Abstract-Developing reusable software for mobile robots is still challenging. Even more so for swarm robots, despite the desired simplicity of the robot controllers. Prototyping and experimenting are difficult due to the multi-robot setting and often require robot-robot communication. Also, the diversity of swarm robot hardware platforms increases the need for hardware-independent software concepts. The main advantages of the commonly used robot software architecture ROS 2 are modularity and platform independence. We propose a new ROS 2 package, ROS2SWARM, for applications of swarm robotics that provides a library of ready-to-use swarm behavioral primitives. We show the successful application of our approach on three different platforms, the TurtleBot3 Burger, the TurtleBot3 Waffle Pi, and the Jackal UGV, and with a set of different behavioral primitives, such as aggregation, dispersion, and collective decision-making. The proposed approach is easy to maintain, extendable, and has good potential for simplifying swarm robotics experiments in future applications.

I. INTRODUCTION

Robot swarms [1] are decentralized systems in which relatively simple robots solve tasks collectively. One of their great advantages is the robustness of the system; there is no single point of failure as the defect of some swarm members will not prevent successful task execution.

A common software engineering challenge in mobile robotics is the development of reusable robotic software, which "is difficult primarily due to the variability in robotic platforms" [2]. Software engineering challenges specific to swarm robotics are "the lack of general tools for experimentation", "common 'swarm libraries' do not yet exist, and reusing code is difficult owing to the lack of swarm-centric development platforms" [3]. The arguably most prominent effort in mobile robotics to push for common software architectures and reusable code is the Robot Operating System (ROS) [4]. While ROS [4] is widely used in both single mobile robot applications and multi-robot systems, it is rarely used for swarm robotics to date [5]. This is, among other reasons, due to the central ROS master node which counteracts the paradigm of decentralization in swarm robotics. However, ROS 2 [6] replaces the central ROS master node with a Data Distribution Service (DDS) as a middleware and is thus decentralized. This change makes ROS 2 interesting for swarm robotics research as it brings about several advantages, such as easy reusability of code on different robot platforms. Several works on distributed



(a) Five TurtleBot3 Waffle Pis

(b) Jackal UGV

Fig. 1. Two of the three used robot hardware platforms: (a) a swarm of the ground mobile robot platform TurtleBot3 Waffle Pi from Robotis and (b) a Jackal unmanned ground vehicle (UGV) from Clearpath Robotics.

multi-robot systems and robot swarms using ROS 2 were recently published, ranging from robots that synchronize and swarm [7], [8] over a toolbox for distributed control schemes for heterogeneous multi-robot systems [9] to concepts for large-scale scalable micro-aerial vehicle swarms [10].

In swarm robotics systems, both hardware and control are usually kept simple relying on self-organization and emergence [1]. Despite the simple reactive controllers of individual robots, complex behaviors emerge due to the collaboration between robots. There are several primitive behaviors in swarm robotics, that are frequently used and studied, such as aggregation (i.e., robots group in one spot), dispersion (i.e., robots distribute while staying in contact), and flocking (i.e., bird-inspired coordinated motion).

As ROS 2 allows for the reuse of code on various robot platforms, a ROS 2 package providing such swarm behavior primitives allows for an easy and rapid setup of swarm robot experiments. We introduce ROS2sWARM, a ROS 2 package that provides an easy-to-extend framework for and a library of swarm robot behaviors for mobile robots. Included swarm behaviors can be used out of the box with any ROS or ROS 2 compatible mobile robot platform that provides sensor data via a scan message and sets the linear speed along the x-axis and the rotational speed around the z-axis. Since we are building on the modularity of ROS 2, new behaviors can easily be added and used as modules to build more complex behaviors.

In the next section, we discuss design patterns for collective behaviors as well as existing ROS and ROS 2 packages for swarm robotics. Section III presents the structure and implementation of ROS2SWARM and introduces the robot platforms that are already supported. In Sec. IV, we show the functionality and benefits of our ROS2SWARM package in simulation and real robot experiments using the TurtleBot3

¹TKK, MJB, TP, HH are with the Institute of Computer Engineering, University of Lübeck, Germany kaiser@iti.uni-luebeck.de

²LS is with the Institute of Robotics, University of Lübeck, Germany

 $^{^3{\}rm GS}$ is with the Institute for Electrical Engineering in Medicine, University of Lübeck, Germany

Burger, the TurtleBot3 Waffle Pi (see Fig. 1a), and the Jackal UGV (see Fig. 1b) robots. Finally, we summarize our work and conclude with an outlook on on-going and future work that will be included in the package in Sec. V.

II. RELATED WORK

To our knowledge, we are the first to propose a ROS 2 package that implements modular and reusable design patterns for swarm behaviors using a decentralized setup in simulations and on real robot hardware. We were inspired by several works on design patterns to facilitate the engineering of distributed multi-agent systems when creating the modular framework. De Wolf and Holvoet [11], for example, consolidate best practices for engineering selforganizing emergent systems in a catalog of design patterns for decentralized coordination mechanisms, such as digital pheromones. Fernandez-Marquez et al. [12] go one step further by proposing modular and reusable design patterns for bio-inspired mechanisms for self-organizing systems. They organize patterns into three layers, including basic mechanisms and two levels of composed patterns. Pitonakova et al. [13] focus specifically on robot swarms by presenting information exchange design patterns for the example of foraging while other swarm tasks are left subject to future work. Loosely inspired by these works, we create a structure of basic and combined patterns that allows for the implementation of modular and reusable swarm behaviors in our ROS2SWARM package.

Already existing ROS and ROS 2 packages for swarm robotics applications do not create such a modular and reusable framework, or they even rely on central components. Most of the existing packages rely on ROS and its central ROS master node. ROSBuzz [14], for example, integrates the swarm-oriented programming language Buzz [3] with ROS to ease the deployment of heterogeneous robot swarms in the field. However, a global positioning system is required. Similarly, the ROS programming framework micros_swarm_framework [15] is inspired by the Buzz programming language and provides abstractions and swarm functions to facilitate the implementation of interaction and communication between swarm members in swarm applications. CMUSWARM [16] is a full stack swarm architecture in ROS that allows for the design, deployment, and evaluation of swarm algorithms. But the framework relies on a centralized core layer for the deployment of the algorithms to the swarm. Hrabia et al. [17] present a modular and reusable ROS framework for self-organization in multirobot systems. Self-organization mechanisms are integrated based on an adapted version of the bio-inspired design patterns by Fernandez-Marquez et al. [12]. Other swarm related ROS packages focus on providing swarm simulations. swarm_stage_ros [18], for example, configures a swarm robot scenario for a cleaning task in the Stage simulator automatically, and swarm_robot_ros_sim [19] provides a generic swarm robot simulation platform for decentralized formation control. Unlike all the previously described ROS packages, CHOIRBOT [9] is based on ROS 2. It provides a toolbox



Fig. 2. UML diagram of ROS2SWARM's package architecture. The *AbstractPattern* class inherits from the standard ROS 2 *node* class. *MovementPattern* and *VotingPattern* are subclasses of the *AbstractPattern* class and differentiate implemented behaviors into patterns for robot movement and patterns for collective decision-making. Left out of the pattern hierarchy is the *HardwareProtectionLayer* that prevents collisions with obstacles. Not shown are several utility classes for processing laser scan data, better readability of state machines, and handling of common voting procedures and standard ROS 2 processes.

for cooperative robots to facilitate the implementation of optimization-based distributed control schemes, but requires global position information of all robots. In contrast, our approach of ROS2SWARM is decentralized as it relies on ROS 2, does not require any global information, such as robot positions, and allows for the easy implementation and execution of swarm behaviors in simulations and on real robots.

III. ROS2SWARM PACKAGE

ROS2SWARM¹ provides a library of ready-to-use behavioral primitives for swarm robotics applications that can easily be extended. Following the key concept of distributed control in robot swarms, the package and consequently the swarm behaviors are executed on each robot autonomously and independently. For some behaviors, data sharing between robots is needed and implemented using global namespace ROS 2 topics requiring all swarm members to be in the same network. In the following, we introduce the general package structure of ROS2SWARM, the pattern hierarchy for the implementation of swarm behaviors, and how the package can easily be used with different ground mobile robots in the Gazebo simulator [20] and on real robots.

A. General Package Structure

ROS2SWARM consists of a main ROS 2 Python package providing the swarm behaviors and an accompanying C++ package with custom ROS 2 message interfaces. The main package is structured by a class hierarchy as shown in Fig. 2. The *AbstractPattern* class serves as the basis for the implementation of all swarm behaviors. We distinguish between two different types of patterns: *movement patterns*, which control the motion of the robot swarm, and *voting patterns*, which implement collective decision-making. For the voting patterns, additional ROS 2 message interfaces that allow exchanging opinions between robots are provided by

¹https://gitlab.iti.uni-luebeck.de/ROS2/ros2swarm. Currently available for ROS 2 Dashing Diademata.



Fig. 3. Communication between the ROS 2 nodes of ROS2SWARM using the ROS 2 publisher/subscriber model when executing movement patterns. Laser scan data from a LiDAR is used by the movement pattern to calculate a drive command and by the hardware protection layer to check for potential collisions. The hardware protection layer sends either the drive command from the movement pattern or an adjusted command to avoid obstacles to the robot's motors. ROS 2 nodes are represented by circles and ROS 2 topics by ellipses. Publishers (*Pub*) and subscribers (*Sub*) are represented by rectangles. Arrows indicate the data flow.

the accompanying C++ package. Both movement patterns and voting patterns can be implemented as standalone *basic patterns* or by combining multiple patterns into a *combined pattern*. Termination conditions for each pattern can easily be included, for example, stopping the execution of a pattern until timeout. Hence, we obtain a modular package structure that allows for reusing patterns.

The parameters of each pattern are set via YAML files so that the behaviors can easily be adjusted for individual applications. We include a set of these parameter files for each robot platform (i.e., TurtleBot3 Waffle Pi, TurtleBot3 Burger, and Jackal so far) to enable easy switching between different platforms. Patterns are executed independently for each robot and can easily be started via launch scripts that handle the start of all required ROS 2 nodes with the specified parameterization. In addition, the launch scripts remap ROS 2 nodes to a robot's individual namespace, where it is required to ensure correct data handling in the swarm. Topics in the global namespace are only used for essential neighbor-to-neighbor communication within the swarm, for example, as required by voting patterns. In the current setup of robot experiments with ROS2SWARM, a robot has no information about the distance to a swarm member that shared a message via the global namespace topic. Thus, all other swarm members are considered as a robot's neighbors to date. Our architecture allows for easily adding new patterns by creating a new subclass of either the MovementPattern class or the VotingPattern class that implements the pattern logic as well as a launch script and a parameter file per robot platform.

We include a *Hardware Protection Layer* next to the pattern hierarchy to prevent collisions with obstacles, including other robots, and possible hardware damage to robots, that runs permanently and independently from the executed swarm behavior. As it is especially relevant when executing



Fig. 4. Arbitration architecture of the hardware protection layer. The light blue area indicates the sensor range which is used by the movement pattern to calculate drive commands. Hardware protection is active and suppresses (using 'suppression' S as defined in the subsumption architecture) drive commands from the movement patterns when there are obstacles in the dark blue area as shown for the upper robot.

movement patterns, we explain the hardware protection layer in more detail in Sec. III-B. The hardware protection layer is left out of the pattern hierarchy to ensure its stability independent from the implemented swarm behaviors. Furthermore, several utility classes provide common methods to process laser scan messages and to handle voting procedures, and an extendable Enum with a basic set of states for better readability of implemented state machines. We verified our code by running several tests. Next to the ROS 2 packages constituting ROS2SWARM, we provide several ROS and ROS 2 packages and launch scripts for running the implemented behavioral primitives on simulated swarms in Gazebo and on real robots.

B. Movement Patterns

Movement patterns implement behavioral primitives that initiate and guide the robot motion. Six movement patterns are included in the current version of ROS2SWARM: *attraction, dispersion, discussed dispersion, drive, random walk,* and *flocking* (cf. Table I).

Fig. 3 visualizes the interplay between the different ROS 2 nodes that are used for movement patterns via the ROS 2 publisher/subscriber model. Each robot executes the movement pattern independently from the rest of the swarm. The movement pattern receives the robot's laser scan data (e.g., from a LiDAR), processes it, and calculates a drive command (i.e., linear speed and rotation speed). The hardware protection layer receives both the laser scan data and the drive command from the movement pattern. It uses an arbitration architecture to decide based on the distance to the robot's nearest obstacle whether to send the drive command determined by the movement pattern or a drive command for collision avoidance to the robot's motors, see Fig. 4. The distance threshold to trigger hardware protection can be set via a parameter and can thus be adjusted for different mobile robot platforms. If there is at least one obstacle, including other robots, within the specified range, hardware protection sends a drive command to the robot calculated based on a repulsive potential field approach. This way robots should keep a distance to obstacles, but we cannot guarantee to prevent all collisions. Otherwise, the drive command determined by the movement pattern is

 TABLE I

 Behaviors that are currently included as patterns in the ROS2swarm package.

	pattern	description
movement	attraction	Aggregation of robots based on an attractive potential field.
	dispersion	Distribution of robots based on a repulsive potential field.
	discussed dispersion	Distribution of robots while maintaining a distance decided on by the swarm.
	drive	Driving straight ahead.
	flocking	Minimalist flocking algorithm based on Moeslinger et al. [21].
	random walk	Random walk based on a simple state machine switching between driving straight ahead and turning randomly.
voting	majority rule voter model	Opinion update based on the majority opinion. Opinion update based on adopting the opinion of a random neighbor.

executed. Hardware protection determines whether obstacle avoidance is necessary every time new laser scan data is received, regardless of whether also a new drive command has been received from the movement pattern. Therefore, movement patterns have to publish drive commands regularly to ensure their correct execution.

C. Voting Patterns

Voting patterns implement collective decision-making behaviors. Collective decision-making is an important capability of a robot swarm to implement consistent actions on swarm level [22], [23]. Currently, ROS2SWARM includes two voting patterns: *majority rule* and *voter model* (cf. Table I).

Each voting pattern uses a common ROS 2 topic in the global namespace to enable swarm members to exchange their opinions via a custom ROS 2 message type that includes the robot's ID and its opinion represented by an Integer. All swarm members have to be in the same network in order to have access to a shared global namespace topic. Based on the other robots' opinions, robots then update their own opinion using the decision rule implemented by the pattern. The currently included voting patterns make use of a tumbling time window approach, that is, the data stream is split into fixed-size, contiguous time intervals that do not overlap. Other options, such as a sliding time window approach (i.e., time intervals can contain overlapping data) or including synchronization, are possible within the framework and can be implemented if required. The data received during one time frame is then processed by the decision rule to find a new opinion. Robots share their opinions after each update and for the most part, the swarm eventually reaches consensus. By contrast to classical swarm settings that rely purely on local communication, each swarm member can currently access the opinions of the full swarm via the global namespace topic.

D. Robot Platforms

ROS2SWARM is designed to work with any ground mobile robot that (i) supports ROS or ROS 2, (ii) has sensors that cover the robot's surroundings and publish laser scan data, and (iii) can be controlled by specifying a linear and a rotational speed. Although ROS2SWARM is implemented in ROS 2, it sill can be used with robot platforms running on ROS by using a network bridge [24] that enables the exchange of messages between ROS and ROS 2. To include support for a new robot platform, a configuration file setting the robot's parameters, launch scripts for the robot platform, and parameter files for the patterns have to be added to the existing package. We provide out-of-the-box functionality in the Gazebo simulator and for real robots for three ground mobile robot platforms: Robotis TurtleBot3 Waffle Pi, Robotis TurtleBot3 Burger, and Clearpath Robotics Jackal UGV.

The TurtleBot3 mobile robots are small and modular ROS and ROS 2 research and education platforms for indoor environments that are controlled via a Raspberry Pi. The TurtleBot3 Burger and the TurtleBot3 Waffle Pi (cf. Fig. 1a) differ mainly in size, that is, the TurtleBot3 Burger has a smaller footprint and is taller than the TurtleBot3 Waffle Pi. Both have differential drive, an IMU, and a LiDAR with 360° field of view and a range of 0.12 m to 3.5 m. Additionally, the TurtleBot3 Waffle Pi has a Raspberry Pi camera.

The Jackal robots (cf. Fig. 1b) are mobile research platforms for outdoor environments. They are equipped with an onboard computer, GPS, and IMU by default. Our version is additionally equipped with a Jetson TX2 board, a ZED stereo camera, and an Ouster OS1-16 LiDAR with 16 vertical layers, 360° field of view, and a range of 0.8 m to 5 m. In contrast to the TurtleBot3 mobile robots, the Jackal supports only ROS to date. Nevertheless, ROS2SWARM can be used with the robot platform by using the aforementioned network bridge [24] that enables the exchange of messages between ROS and ROS 2.

IV. EXPERIMENTAL EVALUATION

We show the versatility of our ROS2SWARM package by running several experiments using TurtleBot3 Burger, TurtleBot3 Waffle Pi, and Jackal robots (cf. Sec. III-D). In Experiment 1, we execute the basic attraction pattern on the three different robot platforms in simulation and on the real TurtleBot3 Waffle Pi to illustrate how swarm behaviors can be used out of the box using ROS2SWARM. In Experiment 2, we combine basic patterns into a more complex pattern to showcase the modularity of our package. We limit ourselves to these two experiments here due to space constraints, but additional swarm behaviors (cf. Table I) are shown in the supplementary video.

A. Experiment 1: Attraction

The patterns included in the swarm behavior library of ROS2SWARM can be executed out of the box on multiple



(a) Initial robot positions exemplified (b) Final configuration for the by the TurtleBot3 Waffle Pi grouping of seven TurtleBot3 experiment Burgers



(c) Final configuration for the (d) Final configuration for the grouping of seven TurtleBot3 grouping of five Jackals Waffle Pis

Fig. 5. Experiment 1 in the Gazebo simulator: swarms of (a, c) seven TurtleBot3 Waffle Pis, (b) seven TurtleBot3 Burgers, and (d) five Jackal robots executing the attraction pattern. Robots are initially positioned at one meter intervals whereby the rightmost robot is positioned in the arena's center. The arena has a size of $18 \text{ m} \times 18 \text{ m}$, but we only show a subarea for better visualization. Grid cells have a size of $1 \text{ m} \times 1 \text{ m}$ each.

platforms. Here, we show the execution of the basic attraction movement pattern on TurtleBot3 Burger, TurtleBot3 Waffle Pi, and Jackal robots in the Gazebo simulator and in hardware on real TurtleBot3 Waffle Pi robots. The pattern uses an attractive potential field approach to determine drive commands that guide the robots towards all detected obstacles, including other robots, within a sensor range (hereafter referred to as attraction range) that is specified via the pattern's parameter file (see Table I). The LiDAR does not allow to differentiate robots and other obstacles and thus, robots can be attracted to walls or other obstacles instead of to other swarm members (sensors to discriminate walls from robots could be added). The hardware protection layer prevents collisions during pattern execution by ensuring that robots maintain a minimum distance to all obstacles, even when they attempt to group. In all experiments, we set the distance threshold for hardware protection to 0.5 m for TurtleBot3 robots and 1.2 m for Jackal robots taking the minimum detection ranges of the LiDARs into account.

First, we run the attraction pattern on three different swarms in the Gazebo simulator, namely a swarm of seven TurtleBot3 Burgers, a swarm of seven TurtleBot3 Waffle Pis, and a swarm of five Jackals. A launch script provided by ROS2SWARM handles the startup of the simulation environment. The script creates a swarm of the respective robot platform and executes the pattern as specified by its parameters. We use an empty $18 \text{ m} \times 18 \text{ m}$ arena that is bounded by walls and position all robots initially close to the arena center, such that the walls are out of the LiDAR's detection range. As the different robot platforms differ in dimensions and components, we have to parameterize the



Fig. 6. Experiment 1 in hardware: a swarm of seven TurtleBot3 Waffle Pis executing the attraction pattern. (a) Top view of initial robot positions in a $4.8 \text{ m} \times 6.6 \text{ m}$ arena and (b) aggregated robots. Six robots aggregate during pattern execution while the rightmost robot lost the group and is attracted to the wall. The floor tiles have a size of 0.6 m \times 0.6 m.

attraction pattern individually for each platform. We set the attraction range from the minimum detection range of the LiDAR, that is, 0.12 m for both TurtleBot3 robot platforms and 0.8 m for Jackal robots, to a maximum of 2 m for the TurtleBot3s and of 3 m for Jackals. In Fig. 5 we show screenshots of Gazebo simulation runs of the attraction pattern on robot swarms of the three different robot platforms. In all simulations, the swarms successfully gather in one cluster. But the varying distance threshold for hardware protection, which specifies the minimum distance robots try to keep to prevent collisions, leads to differences in the grouping behavior. The TurtleBot3 Burger and the TurtleBot3 Waffle Pi have a low distance threshold for hardware protection. Consequently, the robots can form close groups as shown in Figs. 5b and 5c. By contrast, the Jackal has a rather high distance threshold for hardware protection due to a high minimum detection range of its LiDAR. Hence, the robots do not group as closely as the TurtleBot3 robot platforms.

Next, we run the attraction pattern in hardware on a swarm of seven real TurtleBot3 Waffle Pis in a $6.6 \text{ m} \times 4.8 \text{ m}$ empty arena that is bounded by walls. We set the attraction range from 0.12 m to 0.8 m to avoid that robots are mostly attracted to the arena walls due to the smaller arena. Fig. 6 shows the initial robot positions and the aggregated swarm. We find that six robots form one group that repeatedly splits and merges over time (see supplementary video). Thereby, two groups with three robots each form already after approx. 5 s and merge the first time after around 60 s of execution time. One robot moved away from the group and is guided towards the arena's walls by the attractive potential field as robots and other obstacles cannot be differentiated based on the LiDAR's sensor data only. In summary, our ROS2SWARM package proved to be effective in testing and running swarm behaviors on different robot platforms in simulations and on real robot hardware.

B. Experiment 2: Discussed Dispersion

In the next hardware experiment, we showcase how basic patterns can be used to form combined patterns to create more complex behaviors. This shows the modularity and reusability of the available patterns included in ROS2SWARM's swarm behavior library. As an example, we combine the basic majority rule voting pattern and the basic



Fig. 7. Architecture of the *discussed dispersion* pattern. For 20 s, swarm members execute only the majority rule pattern to determine their majority opinion. This opinion is mapped to a distance by the discussed dispersion pattern using a mapping function f which is set as the minimum distance during dispersion. The dispersion pattern determines the robot's drive commands based on a repulsive potential field approach. Rectangular boxes give patterns and rounded boxes give states.



(a) Initial/final opinions (b) Dispersed swarm at agreed on distance based on the majority rule

Fig. 8. Experiment 2 in hardware: a swarm of seven TurtleBot3 Waffle Pis executing the *discussed dispersion* pattern that combines two basic patterns. The robots keep the start configuration as shown in Fig. 6a for the first 20 s executing only the majority rule pattern. In this case here, robots decide for opinion 1 that is mapped to a minimal distance between robots of 1.0 m. Afterwards, robots disperse in the arena keeping this collectively decided distance (or more) to each other and the arena walls.

dispersion movement pattern and form a new pattern that we call discussed dispersion. The robots are supposed to disperse while maintaining a robot-to-robot distance that they collectively choose themselves. The minimum specifiable robot-to-robot distance equals the distance threshold triggering hardware protection as both dispersion and hardware protection layer use the same repulsive potential field approach. As visualized in Fig. 7, robots execute only the majority rule pattern in the first 20 s staying on their initial positions (as previously shown in Fig. 6a). After this initial decision phase, both the majority rule pattern and the dispersion pattern are executed in parallel. The swarm may still change its opinion, for example, when several new robots are added. The opinion determined by the majority rule pattern is mapped to a distance value using a mapping function f. This is set as the minimum distance that robots maintain from all obstacles including other robots when executing the dispersion pattern.

We execute the combined pattern on a swarm of seven real TurtleBot3 Waffle Pis, see Fig. 8. The opinion of each robot is randomly initialized to $r \in \{0, 1, 2\}$ and mapped to a distance of f(0) = 0.6 m, f(1) = 1.0 m, or f(2) =1.4 m, respectively. In the visualized run, robots collectively choose opinion 1 (distance of 1.0 m). As shown in Fig. 8, the swarm distributes in space until each robot is at least 1.0 m away from all obstacles, that is, other robots and walls. In summary, we have shown how the modularity of our swarm behavior library can easily be leveraged to derive new and more complex swarm behaviors.

V. CONCLUSION

Our ROS2SWARM ROS 2 package provides a library of ready-to-use swarm behaviors for collective motion and collective decision-making applications. We showed that the package can be extended easily and that all patterns are reusable and modular. This allows to run combined patterns for more complex applications as illustrated in our Experiment 2 (cf. Sec. IV-B). To date, three ground mobile robots are supported out of the box by ROS2SWARM, namely TurtleBot3 Burger, TurtleBot3 Waffle Pi, and Jackal. Support for further robot platforms running ROS or ROS 2 can be integrated with minimal effort. ROS2SWARM unites the advantages of swarm robotics, namely adaptability, robustness, and scalability, and of ROS 2, that is, platform independence and modularity. In total, our package is versatile and enables a quick and easy setup of swarm robotics applications.

While the package in its current form already offers the most important concepts for swarm applications, we want to add several enhancements in the future. First, all robots must currently be connected to a common network for the execution of voting patterns. So far, we use a central WiFi network with a single access point for this purpose and thus have a potential single point of failure. In future, we will integrate IR communication [25] or switch to an ad-hoc WiFi network that allows to determine a robot's neighbors based on the decreasing WiFi signal strength over distance. Also, our hardware protection layer does not guarantee the prevention of collisions using the current repulsive potential field approach. We plan to investigate more reliable approaches, such as decentralized safety barrier certificates [26]. Furthermore, we continue to extend the library with more collective behaviors and, in particular, with more complex behaviors, such as pattern formation. We will integrate tests into the package for quantitatively evaluating the effectiveness of implemented swarm behaviors, for example, by computing the mean-distance-to-centroid in aggregation. A longer-term plan is to integrate support for more robot platforms that may need different drive commands (e.g., quadcopter swarms [8], [27]) than the currently included platforms and to make the package independent of a LiDAR by supporting different sensor setups for obstacle detection, such as multiple IR sensors and cameras. We plan to release the package to the public ROS 2 buildfarm. We hope to form a developer community that keeps extending and maintaining this package for the profit of a growing swarm robotics research and teaching community.

ACKNOWLEDGMENTS

The authors thank Steffen Fleischmann, Vincent Jansen, and Daniel Tidde for their contributions to the implementation of ROS2sWARM, Yuanchen Yuan for proofreading the paper, and Lucas Kaiser for support in image editing.

REFERENCES

- [1] H. Hamann, *Swarm Robotics: A Formal Approach*. Cham: Springer International Publishing, 2018.
- [2] I. A. D. Nesnas, R. Simmons, D. Gaines, C. Kunz, A. Diaz-Calderon, T. Estlin, R. Madison, J. Guineau, M. McHenry, I.-H. Shu, and D. Apfelbaum, "CLARAty: Challenges and steps toward reusable robotic software," *International Journal of Advanced Robotic Systems*, vol. 3, no. 1, p. 5, 2006.
- [3] C. Pinciroli and G. Beltrame, "Buzz: a programming language for robot swarms," *IEEE Software*, vol. 33, no. 4, pp. 97–100, 2016.
- [4] M. Quigley, J. Faust, T. Foote, and J. Leibs, "ROS: an open-source Robot Operating System," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.
- [5] M. Dorigo, G. Theraulaz, and V. Trianni, "Swarm robotics: Past, present, and future [point of view]," *Proceedings of the IEEE*, vol. 109, no. 7, pp. 1152–1165, 2021.
- [6] Y. Maruyama, S. Kato, and T. Azumi, "Exploring the performance of ROS2," in 2016 International Conference on Embedded Software (EMSOFT), 2016, pp. 1–10.
- [7] A. Barciś, M. Barciś, and C. Bettstetter, "Robots that Sync and Swarm: A proof of concept in ROS 2," in 2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS), 2019, pp. 98–104.
- [8] A. Barciś and C. Bettstetter, "Sandsbots: Robots that sync and swarm," *IEEE Access*, vol. 8, pp. 218752–218764, 2020.
- [9] A. Testa, A. Camisa, and G. Notarstefano, "ChoiRbot: A ROS 2 toolbox for cooperative robotics," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2714–2720, 2021.
- [10] J. P. Queralta, Y. Xianjia, L. Qingqing, and T. Westerlund, "Towards large-scale scalable MAV swarms with ROS2 and UWB-based situated communication."
- [11] T. De Wolf and T. Holvoet, "Design patterns for decentralised coordination in self-organising emergent systems," in *Proceedings of the* 4th International Conference on Engineering Self-Organising Systems, ser. ESOA'06. Berlin, Heidelberg: Springer-Verlag, 2006, p. 28–49.
- [12] J. L. Fernandez-Marquez, G. Di Marzo Serugendo, S. Montagna, M. Viroli, and J. L. Arcos, "Description and composition of bioinspired design patterns: a complete overview," *Natural Computing*, vol. 12, no. 1, pp. 43–67, Mar. 2013.
- [13] L. Pitonakova, R. Crowder, and S. Bullock, "Information exchange design patterns for robot swarm foraging and their application in robot control algorithms," *Frontiers in Robotics and AI*, vol. 5, p. 47, 2018.
- [14] D. St-Onge, V. S. Varadharajan, I. Švogor, and G. Beltrame, "From design to deployment: Decentralized coordination of heterogeneous robotic teams," *Frontiers in Robotics and AI*, vol. 7, p. 51, 2020.
- [15] C. Xuefeng, C. Zhongxuan, W. Yanzhen, and Y. Xiaodong, "micros_swarm_framework - ROS Wiki," 2019. [Online]. Available: https://wiki.ros.org/micros_swarm_framework
- [16] K. Morris, G. Arpino, S. Nagavalli, and K. Sycara, "Full stack swarm architecture," *RISS Working Papers Journal*, 2018.
- [17] C.-E. Hrabia, T. K. Kaiser, and S. Albayrak, "Combining selforganisation with decision-making and planning," in *Multi-Agent Systems and Agreement Technologies*, F. Belardinelli and E. Argente, Eds. Cham: Springer International Publishing, 2018, pp. 385–399.
- [18] V. C. Kalempa, M. A. S. Teixeira, A. S. de Oliveira, and J. A. Fabro, "Intelligent dynamic formation of the multi-robot systems to cleaning tasks in unstructured environments and with a single perception system," in 2018 Latin American Robotic Symposium, 2018 Brazilian Symposium on Robotics (SBR) and 2018 Workshop on Robotics in Education (WRE), Nov 2018, pp. 71–76.
- [19] Y. Liu, A. Ali, and G. Dare, "swarm_robot_ros_sim," 2020. [Online]. Available: https://github.com/yangliu28/swarm_robot_ros_sim
- [20] Gazebo, "Gazebo," 2020. [Online]. Available: http://gazebosim.org/
- [21] C. Moeslinger, T. Schmickl, and K. Crailsheim, "A minimalist flocking algorithm for swarm robots," in Advances in Artificial Life. Darwin Meets von Neumann, G. Kampis, I. Karsai, and E. Szathmáry, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 375–382.
- [22] G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo, "Collective decision with 100 Kilobots: Speed vs accuracy in binary discrimination problems," *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 3, pp. 553–580, 2016.
- [23] G. Valentini, E. Ferrante, and M. Dorigo, "The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives," *Frontiers in Robotics and AI*, vol. 4, p. 9, 2017.

- [24] D. Thomas, E. Fernandez, and W. Woodall, "State of ROS 2 - demos and the technology behind," in *ROSCon Hamburg* 2015. Open Robotics, September 2015. [Online]. Available: https://doi.org/10.36288/ROSCon2015-900743
- [25] F. Arvin, K. Samsudin, and A. R. Ramli, "A short-range infrared communication for swarm mobile robots," in 2009 International Conference on Signal Processing Systems, 2009, pp. 454–458.
- [26] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.
- [27] C. Steup, S. Parlow, S. Mai, and S. Mostaghim, "Generic componentbased mission-centric energy model for micro-scale unmanned aerial vehicles," *Drones*, vol. 4, no. 4, p. 63, 2020.