A Stochastic Approach to Generate Emergent Behaviors in Robotic Swarms

Paulo Rezeck and Luiz Chaimowicz

Laboratório de Visão Computacional e Robótica (VeRLab)
Programa de Pós-Graduação em Ciência da Computacão
Universidade Federal de Minas Gerais (UFMG)
Belo Horizonte, MG – Brazil

{rezeck,chaimo}@dcc.ufmg.br

Abstract. *This research introduces a novel stochastic methodology leveraging Gibbs Random Fields (GRFs) to promote a variety of emergent behaviors in robotic swarms, such as flocking, segregation, cooperative object transport, and pattern formation. By utilizing only local information and decentralized control mechanisms, the approach ensures robust and scalable swarm behaviors. Through both numerical simulations and real-world experiments using HeRo 2.0, a low-cost swarm robotic platform developed in this work, we demonstrate the benefits of the proposed methodology, including adaptability, robustness, and resilience. The stochastic approach proposed here is positioned to significantly advance the field of swarm robotics by enabling diverse behaviors, contributing to the evolution of the field, and shaping the way for new applications.*

***Keywords**: Robotics, Multi-agent system, Swarm robotics, Probabilistic robotics.*

Ph.D. dissertation defended on December 4, 2023 (PPGCC – UFMG). Committee: Prof. Luiz Chaimowicz (DCC – UFMG), Prof. Douglas G. Macharet (DCC – UFMG), Prof. Renato M. Assunção (DCC – UFMG), Prof. Valdir G. Júnior (SEL – USP), and Prof. David J. Saldaña Santacruz (CSE – Lehigh University/USA). Full text, publications, and videos are available at: https://rezeck.github.io/#phd.

1. Introduction

Swarm robotics has become an exciting and promising area within robotics and multi-agent systems. It focuses on the coordination of large groups of robots to collaboratively perform complex tasks. Unlike traditional robotics, which emphasizes the capabilities of individual robots, swarm robotics leverages the interactions between multiple robots to achieve emergent behaviors, drawing inspiration from natural systems like bird flocks and ant colonies. This paradigm shift offers transformative potential across various industries, including manufacturing, agriculture, disaster response, and environmental monitoring.

However, despite rapid advancements, a major challenge persists: developing flexible methodologies capable of generating a broad spectrum of swarm behaviors. Many existing approaches are tailored to specific tasks, limiting their adaptability and scalability. This inflexibility restricts their application in diverse real-world scenarios and hinders innovation. The research presented here seeks to address these limitations by developing an adaptable methodology capable of producing a variety of swarm behaviors. This approach aims to fully harness the potential of swarm robotics, enabling it to tackle complex real-world challenges and be integrated into a broader range of applications.

This study has two primary objectives in advancing swarm robotics. The first objective involves introducing a novel methodology that draws inspiration from statistical mechanics to synthesize control strategies that induce diverse collective behaviors in robotic swarms, even with minimal sensing information. Our approach uniquely extends the application of Gibbs Random Fields (GRFs) to swarm robotics, establishing decentralized control mechanisms that allow robots to coordinate their actions based on local information. This decentralized approach is critical for enhancing system robustness and scalability, as it eliminates reliance on a central controller and reduces vulnerability to single points of failure. The second objective focuses on demonstrating the flexibility and adaptability of this methodology in addressing key challenges in swarm robotics. As depicted in Figure 1, these challenges encompass a spectrum of complex behaviors, including the simultaneous integration of flocking and segregation, the cooperative transportation of objects, and the formation of intricate patterns. These tasks require a sophisticated understanding of inter-robot interactions and environmental dynamics. Through rigorous theoretical analysis, extensive simulations, and practical real-world experiments, we showcase the effectiveness and efficiency of our methodology in tackling these challenges, highlighting the versatility of GRFs as a control paradigm in swarm robotics.

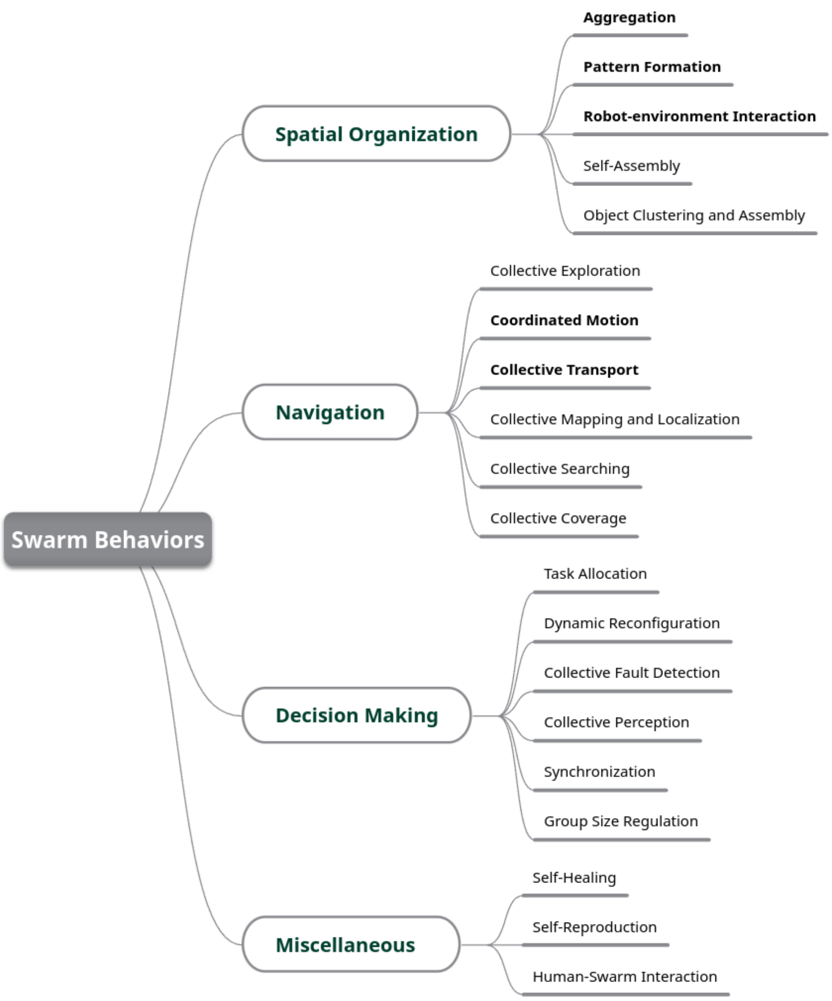

Figura 1. Swarm robotics behaviors adapted from [Schranz et al. 2020]. The highlighted behaviors are tackled in this work.

This work makes substantial contributions to the field of swarm robotics, culminating in four publications in prestigious academic journals and conferences. One of the key innovations is our novel stochastic and decentralized approach, utilizing dynamic GRFs to enable heterogeneous robot swarms to exhibit a wide range of behaviors through local interactions. Notably, we pioneered methods for simultaneous segregation and flocking behaviors [Rezeck et al. 2021b], uniquely handling these behaviors from a random initial state with exclusively local information. This innovation is crucial for tasks requiring coordinated yet differentiated actions within a swarm. Additionally, our research on cooperative object transportation [Rezeck et al. 2021a] demonstrates a method that operates without explicit communication or centralized control, accommodating various object properties and showcasing remarkable adaptability and robustness in dynamic conditions. Our work on pattern formation [Rezeck and Chaimowicz 2022] further illustrates the system's capability to produce various and adaptable spatial configurations, interesting for applications like environmental monitoring and temporary structure construction. Enhancements to the HeRo swarm robotic platform [Rezeck et al. 2023] bolster its modularity and practical deployment, validating the effectiveness of our methodologies and bridging the gap between theoretical concepts and real-world implementation. These contributions collectively position our methodology as a unique solution in swarm robotics, demonstrating significant potential for real-world applications and future research directions.

2. Related Work

This section reviews the existing literature on the use of Markov and Gibbs Random Fields in swarm robotics, exploring their applications and underlying methodologies. Markov Random Fields (MRFs) and Gibbs Random Fields (GRFs) are mathematical frameworks commonly utilized in areas like statistical mechanics and computer vision to model interactions among individual system components. Although these models have found extensive use in various domains, their application in swarm robotics is still relatively underexplored, indicating a need for more focused research. An early contribution to this area was made by [Baras and Tan 2004], who demonstrated the effectiveness of using MRFs to control swarm behavior. They modeled a swarm of robots as an MRF, with robots and their sensing links forming the vertices and edges of a graph, respectively. By using Gibbs potentials to define the interactions within the swarm, they facilitated global objectives such as aggregation or dispersion through simulated annealing guided by the Gibbs sampler. This methodology proved to be efficient and scalable for controlling the movement of large robot swarms. Further developments by [Xi et al. 2006] and [Tan et al. 2010] refined these ideas, introducing hybrid algorithms that combined gradient descent with Gibbs sampling to optimize the behavior of multi-robot systems. These methods showed robustness and efficiency, even under conditions of uncertain sensing, by converging to minimal potential energy configurations.

In recent years, the application of MRFs in swarm robotics has been extended to include flocking behaviors. For example, [Fernando 2021] developed a flocking algorithm that utilized differential flat dynamics and MRFs to model interactions between robots, thus facilitating real-time coordination and collision avoidance within a swarm. Building on the foundational work of [Baras and Tan 2004], our research introduces innovative methods to achieve behaviors such as segregation, flocking, cooperative object transportation, and pattern formation using a GRF-based approach. Our methodology offers several significant improvements. First, it integrates robot kinematics, allowing for continuous movement modeling within confined spaces, a notable improvement over previous models that relied on discrete lattice structures. Additionally, we introduce a new potential function – the Coulomb-Buckingham Potential combined with kinetic energy– to encode low-level swarm behaviors. This approach supports seamless aggregation, interaction with objects, and autonomous navigation, all without the need for centralized control. These advancements highlight the versatility and broad applicability of our methodology, positioning it as a robust, general-purpose solution for swarm control.

3. Designing Swarm Behaviors

This section introduces the application of Gibbs Random Fields (GRFs) to the design of swarm behaviors in robotics, detailing our methodology for configuring a swarm's behavior over time. GRFs are utilized to model the spatial distribution of swarm behaviors, where specific potential functions are strategically combined to create a potential energy landscape. This landscape encodes the configuration of the swarm into a single energy value. By sampling velocities through a conditional probability function, we iteratively minimize this energy, steering the swarm towards desired behaviors. Establishing a direct link between a potential energy function and the desired swarm behavior requires intuition and creativity. Thus, we provide a comprehensive overview of our methodology, as illustrated in Figure 2, alongside a sequence of steps guiding its practical application. This approach employs decentralized control, enabling robots to make decisions based on local information without direct communication, which enhances both system robustness and scalability. The methodology's inherent flexibility also supports adaptation to changing environmental conditions or new objectives, highlighting its potential for fostering diverse swarm behaviors.

1. **Potential Functions as Fundamental Primitives:** Potential functions are critical components in our methodology, acting as the fundamental building blocks. For example, the Coulomb-Buckingham potential is used to mathematically describe electrostatic interactions among charged particles in the swarm. Depending on its configuration, this potential can generate attractive or repulsive forces, thus facilitating the emergence of complex swarm behaviors. When paired with kinetic energy, which considers motion interactions, these potential functions dictate the dynamics of the swarm, influencing navigation patterns and group cohesion.
2. **Modeling Basic Swarm Behaviors:** In this phase, primitive behaviors are combined in various configurations to model diverse basic swarm behaviors, such as group formation, cohesive movement, and interactions with the environment. The creation of an effective potential energy function requires a mix of intuition, creative experimentation, and rigorous testing to accurately embody the desired behaviors. These basic behaviors serve as the foundation for constructing more complex, higher-level swarm behaviors.
3. **Designing Complex Swarm Behaviors:** Building on the basic behaviors identified in the previous step, we design higher-level potential functions that encapsulate more complex swarm activities, such as collective transport, exploration, and searching. For particularly intricate swarm behaviors, it may be necessary to dynamically adjust potential functions based on the swarm's evolving configuration to maintain optimal performance and adaptability.

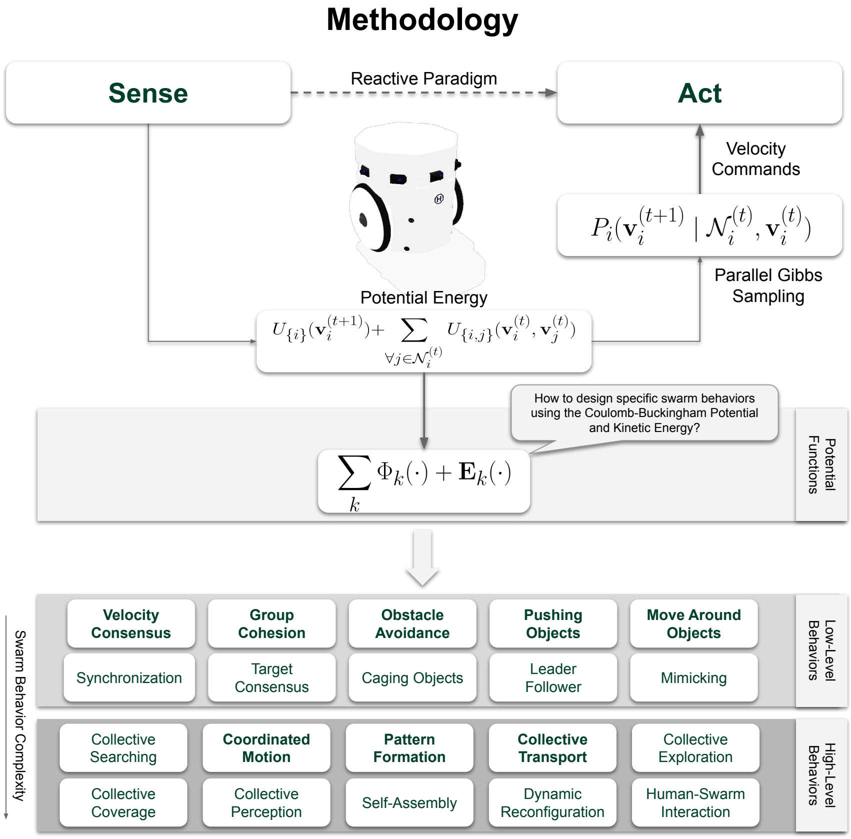

Figura 2. Methodology for inducing diverse swarm behaviors using potential function combinations. Highlighted behaviors are tackled in this work.

4. **Implementing Parallel Gibbs Sampling:** We employ parallel Gibbs sampling to model a probability distribution based on the defined potential energy function. This distribution reflects the likelihood of different swarm configurations. By carefully crafting the potential energy function, we influence the probability distribution, thereby controlling the swarm's behavior. The sampling process serves as a mechanism to generate command velocities that are most likely to reduce the global potential energy, guiding the swarm toward specific behavioral configurations.
5. **Iterative Refinement and Experimentation:** An iterative process is applied to refine and optimize the swarm behaviors. This involves continuous experimentation and adjustment of the potential energy functions to fine-tune the behaviors for specific applications. The iterative nature of this process allows for the continual enhancement of the swarm's performance and adaptability.

Our methodology effectively combines mathematical models with computational algorithms to design swarm behaviors. The following sections will delve into its practical application in producing complex behaviors such as flocking, segregation, cooperative object transportation, and pattern formation. Additionally, we will discuss the design and deployment of a robotic platform used for proof-of-concept experiments, demonstrating the versatility and effectiveness of these methodologies in real-world scenarios.

4. HeRo 2.0: a Low-Cost Robot for Swarm Robotics Research

A key contribution of this dissertation is the development of HeRo 2.0, an open-source, cost-effective robotic platform designed to validate our methodology in real-world experiments. Building on the original HeRo platform [Rezeck et al. 2017], HeRo 2.0 [Rezeck et al. 2023] integrates commercially available components with additive manufacturing techniques, offering affordability without sacrificing functionality. As shown in Figure 3, HeRo 2.0 includes an Espressif ESP8266 microcontroller for robust wireless communication and motion control, an Inertial Measurement Unit (IMU) for sensing orientation and motion, and differential-driven wheels for locomotion. The robot supports long-term autonomy (up to nine hours) and features Firmware Over-The-Air (FOTA) updates. Furthermore, it integrates seamlessly with the Robot Operating System (ROS) and offers comprehensive simulation capabilities in Gazebo, making it a versatile platform for both development and experimentation. These features collectively foster collaboration and innovation within the robotics community, positioning HeRo 2.0 as a valuable tool for swarm robotics research.

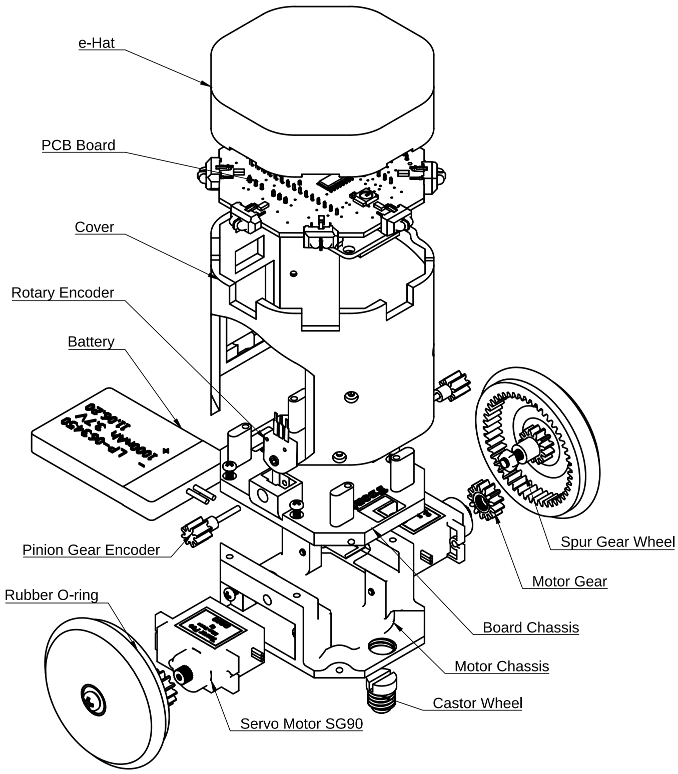

Figura 3. An expanded view of the robot's components and body parts. Video link: https://verlab.github.io/hero_common/.

Beyond its role in swarm robotics research, HeRo 2.0 is also a valuable resource in computer science and engineering education. Its open-source design and versatility make it an ideal platform for exploring swarm robotics and experimenting with decentralized control systems. The platform offers hands-on opportunities for students and researchers to engage with programming, mechanical design, electronics, and experimental robotics, fostering a well-rounded skill set. By bridging these interdisciplinary areas, HeRo 2.0 enhances educational experiences and equips learners with the tools and knowledge needed to tackle complex challenges in both computer science and engineering fields.

5. Flocking-Segregative Behavior

The flocking-segregative behavior in robotic swarms involves coordinated movement (flocking) and the formation of distinct groups (segregation) based on shared characteristics. Prior to this research, the simultaneous achievement of these behaviors, particularly from a random initial state using only local sensing, had not been thoroughly addressed. This challenge is especially complex in scenarios with multiple distinct groups. Figure 4 illustrates an example of this combined behavior.

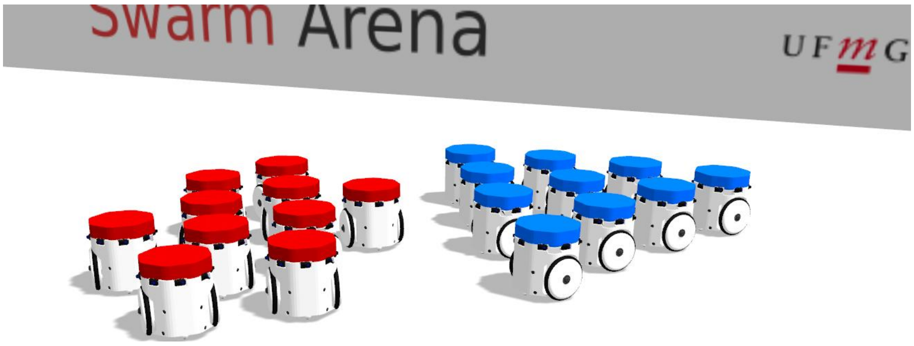

Figura 4. Swarm demonstrates simultaneous segregation and flocking, forming distinct groups while maintaining cohesive navigation.

Our method's effectiveness in achieving flocking-segregative behavior was rigorously tested through a series of experiments. Simulated experiments assessed the method's performance in various configurations, focusing on metrics such as velocity consensus and cohesion despite sensor noise. These simulations demonstrated the method's robustness in maintaining desired behaviors even when faced with robot failure (Figure 5).

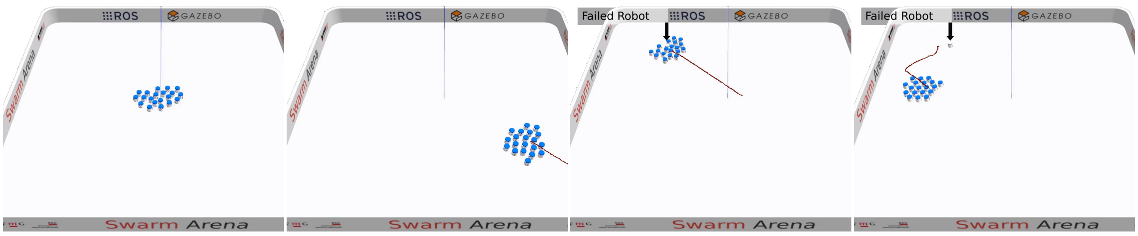

(a) $t = 0$ **s.** **(b)** $t = 50$ **s.** **(c)** $t = 70$ **s.** **(d)** $t = 90$ **s.**

Figura 5. Snapshots show method robustness amid robot failure at 70 seconds. Video link: `https://youtu.be/Rq9ld4gHfo8`.

We also assessed the impact of environmental complexity on the method's performance. Experiments were conducted in simulated environments with varying obstacle configurations, as shown in Figure 6. The results indicated that more complex environments posed greater challenges, resulting in longer segregation times. However, the method consistently demonstrated adaptability, enabling the robots to achieve segregation and navigate through challenging scenarios. These findings highlight the robustness and effectiveness of our approach, even under diverse and complex conditions.

Real-world experiments with HeRo robots provided a proof-of-concept, showcasing the swarm's ability to achieve both segregation and flocking behaviors in constrained physical spaces (Figure 7). These comprehensive experiments underscore the practical

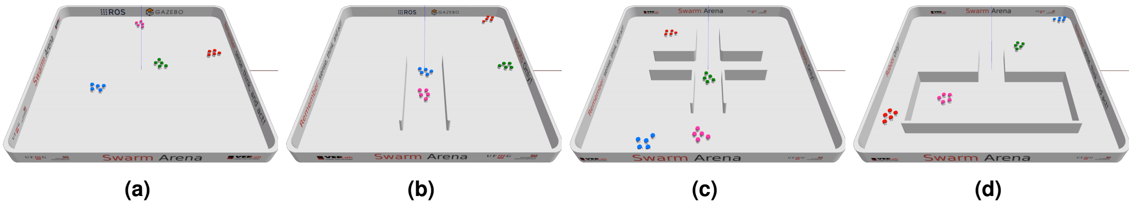

Figura 6. Flocking-segregative behavior demonstrated in simulated environments of varying complexity. Video link: https://youtu.be/playlist?list=PLe1xjXfLIEAfuvBR4KyJMVVhHf5XXgVOy.

significance of simultaneously achieving flocking and segregation in robotic swarms, suggesting potential applications in various areas, such as coordinated transport.

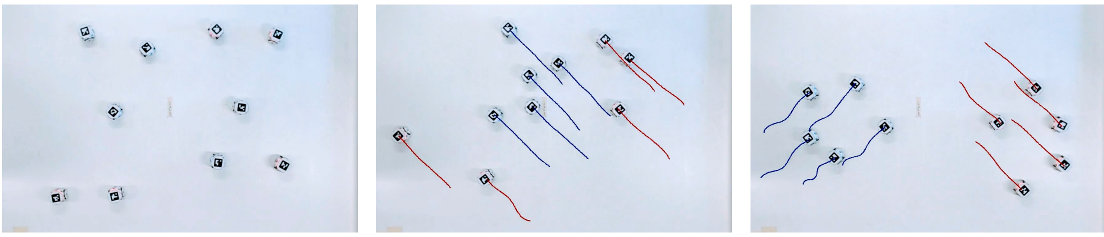

(a) $t = 0$ s. (b) $t = 30$ s. (c) $t = 130$ s.

Figura 7. Snapshots of two robot groups exhibiting segregation and flocking. Video link: https://youtu.be/s1eLOmECcwc.

6. Cooperative Transport Behavior

Robotic swarms are highly valued for their emergent collective behaviors, making them versatile and scalable for tasks such as cooperative transportation. However, decentralized control methods present challenges, particularly in aligning forces and adapting to objects of varying shapes and sizes. This work introduces a novel method for autonomous navigation and cooperative object transportation, illustrated in Figure 8, where robots collaboratively move an object. Unlike most previous approaches, this method is entirely based on local interactions without requiring explicit communication between robots. Each robot estimates the positions and velocities of its neighbors, distinguishes obstacles, and dynamically adjusts navigation and transportation parameters without relying on predefined behaviors or learning algorithms.

We conducted several experiments using ROS middleware and the Gazebo simulator to validate the effectiveness of this cooperative transportation method. One set of experiments tested the transportation of objects with different shapes and masses, including a triangular prism, a rectangular prism, and an octagonal prism. The results indicated that robots were less efficient in pushing the triangular prism than other objects. When pushing at the acute corners of the triangle, fewer robots could group effectively, leading to increased task completion time. As expected, larger objects allowed more robots to participate in the pushing task, increasing transport time overall. Notably, the impact of increasing the size of the triangular prism on transport time was more significant than that of the rectangular or octagonal prisms. This was likely due to difficulty detecting and pushing sharp corners with radial sensors. Despite these variations, the robots

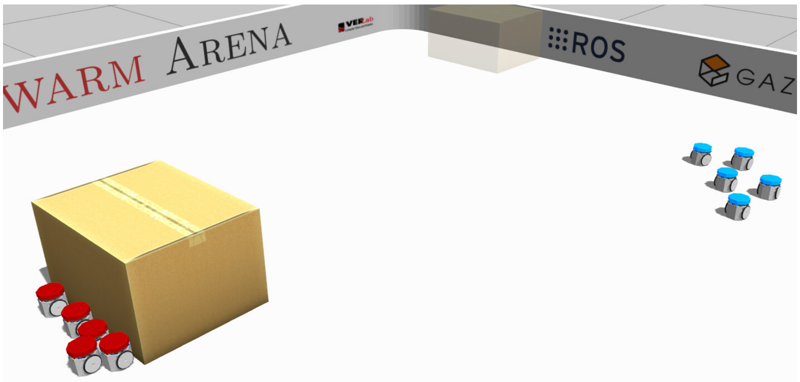

Figura 8. Robots collaboratively transport a solid cardboard object toward a transparent goal. Red robots carry the object, while blue ones search.

successfully transported all objects to their target locations, demonstrating the method's robustness (Figure 9).

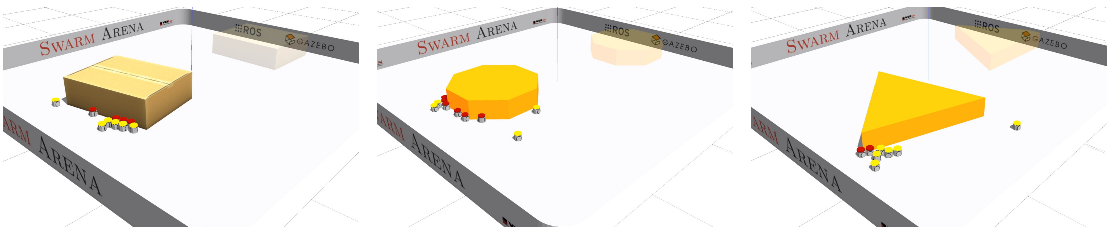

(a) (b) (c)

Figura 9. Swarm's efforts in the transportation of objects with different shapes towards their target location.

Further experiments were conducted in simulated environments with varying levels of complexity. These simulations assessed the method's scalability, adaptability, and robustness in navigating complex environments and scenarios with potential mechanical failures. Additionally, real-world experiments with physical robots provided a practical proof-of-concept. In these trials, the robots demonstrated their ability to navigate complex environments and execute cooperative transportation tasks effectively. Figure 10 shows the robots maneuvering through a complex environment, while Figure 11 captures real-world experiments, underscoring the method's versatility. A video on Youtube[1] provides additional insight.

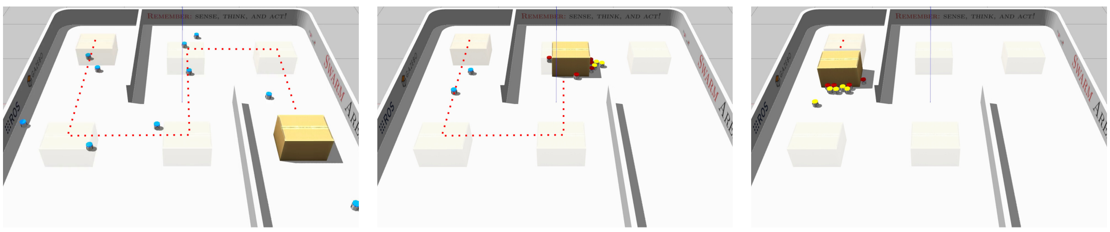

(a) $t = 0$ **s.** **(b)** $t = 485$ **s.** **(c)** $t = 919$ **s.**

Figura 10. Simulated robots transporting an object toward target locations.

[1] Cooperative transport link: https://youtu.be/hrkJKL3W3pQ.

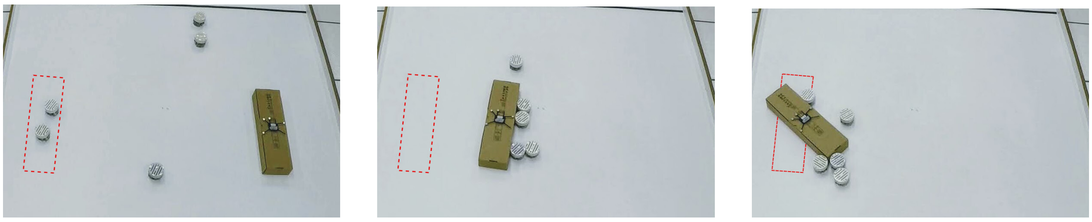

(a) $t = 0$ **s.** **(b)** $t = 240$ **s.** **(c)** $t = 307$ **s.**

Figura 11. Robots transporting an object to a goal marked by a red rectangle.

Overall, this decentralized approach to cooperative transportation emphasizes scalability, robustness, and adaptability. It represents a significant contribution to the field of swarm robotics, demonstrating its potential for real-world applications where flexible and reliable transportation solutions are essential.

7. Pattern Formation Behavior

Pattern formation in swarm robotics involves coordinating robots to form specific shapes, inspired by phenomena observed across various scientific disciplines. Our approach leverages local interactions within a swarm of heterogeneous robots to create intricate patterns, drawing analogies from particle interactions and molecular structures. By modeling the swarm as a dynamic Gibbs Random Field (GRF), our method achieves diverse pattern formations without relying on global information or predefined seeds. Figure 12 showcases an example where the swarm forms a pattern resembling a molecular structure, highlighting the emergent behaviors resulting from local interactions. Unlike existing methods, our minimalist approach demonstrates the feasibility of achieving complex patterns through emergent behaviors, which enhances scalability and adaptability. This capability opens up practical applications in fields such as modular robotics and the construction of temporary structures. Our work introduces a novel perspective by integrating concepts from particle physics into swarm robotics, thereby advancing decentralized and distributed control methodologies for robotic systems.

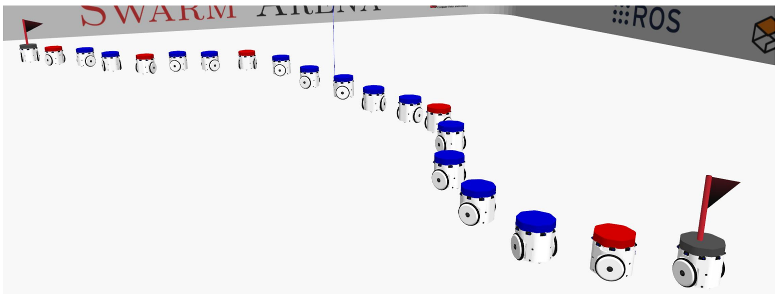

Figura 12. Robots mimic carbon (red) and oxygen (blue) atoms, forming chain patterns for dynamic bridges. Red flags mark the chain's start and end.

To validate our approach, we conducted a series of experiments using both numerical simulations and real-world settings. Initial experiments focused on the diversity of patterns generated by the swarm, including aggregating sub-structures, complex, varied formations, and long-chain dynamic behaviors reminiscent of biological systems (Figure 13). These simulations highlighted the method's ability to produce a wide range of patterns based solely on local interactions.

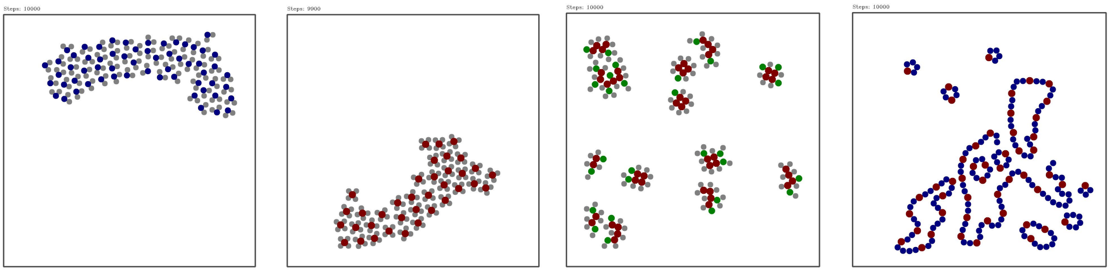

(a) (b) (c) (d)

Figura 13. Four simulated experiments showing that different types of robots create different patterns resembling chemical structures of the (a) water, (b) methane, (c) polyamines and (d) oxocarbons.

Subsequently, we explore the use of robot anchors to guide the formation and positioning of structures, paving the way for more practical applications. By using the Gazebo simulator, we demonstrated the swarm's capacity to form structures like bridges dynamically, showcasing the method's adaptability and responsiveness in simulated environments (Figure 14). Finally, real-world experiments confirmed the method's effectiveness in creating chain-like structures with bridge-like topologies using physical robots (Figure 15). These diverse outcomes underscore the method's versatility, though fine-tuning the formation of specific geometric shapes remains challenging without a coordination mechanism. Videos of these experiments are available online[2], providing further insight into the method's versatility across different environments.

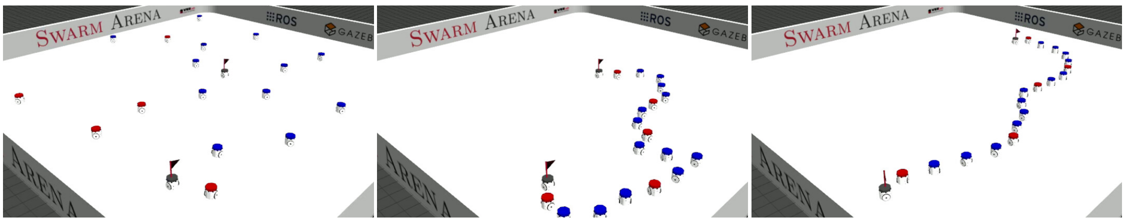

(a) $t = 0$ s. (b) $t = 115$ s. (c) $t = 135$ s.

Figura 14. 20 robots representing carbon (red) and oxygen (blue) atoms form bridge-like structures. Red-flagged robots mark the structure's ends.

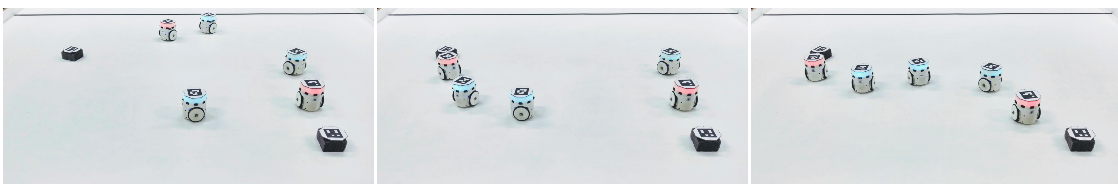

(a) $n = 0$ s. (b) $n = 30$ s. (c) $n = 50$ s.

Figura 15. Snapshots of 5 real robots depict a dynamic bridge pattern with 2 carbon (red) and 3 oxygen (blue) atoms, bounded by black blocks.

[2]Pattern formation link: https://rezeck.github.io/chemistry-inspired-swarm.

8. Conclusion and Future Work

This research introduced a novel methodology using Gibbs Random Fields (GRFs) in swarm robotics to model and control heterogeneous robot groups. Our decentralized approach guides robots toward desired configurations within a GRF framework, ensuring robust and scalable coordination. We addressed complex challenges such as flocking-segregation dynamics, cooperative object transportation, and pattern formation, demonstrating our methods' versatility and effectiveness. Both simulated and real-world experiments validated our approach, showcasing its adaptability and resilience. The development of HeRo 2.0, our open-source low-cost robotic platform, was crucial in demonstrating the practical viability of our methodology. Beyond robotics, our contributions provide valuable insights into decentralized control methods applicable to distributed systems, optimization, and artificial intelligence.

Future work will include more diverse swarm behaviors, such as advanced caging and dynamic spatial coverage. We'll integrate machine learning for better decision-making, explore human-swarm interactions, and adapt to dynamic conditions. Loosening constraints will allow for more complex behaviors and open new research directions.

Referências

Baras, J. S. and Tan, X. (2004). Control of autonomous swarms using gibbs sampling. In *2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601)*, volume 5, pages 4752–4757. IEEE.

Fernando, M. (2021). Online flocking control of uavs with mean-field approximation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8977–8983. IEEE.

Rezeck, P., Assunção, R. M., and Chaimowicz, L. (2021a). Cooperative object transportation using gibbs random fields. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9131–9138. IEEE.

Rezeck, P., Assunção, R. M., and Chaimowicz, L. (2021b). Flocking-segregative swarming behaviors using gibbs random fields. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8757–8763. IEEE.

Rezeck, P., Azpurua, H., and Chaimowicz, L. (2017). Hero: An open platform for robotics research and education. In *Latin American Robotics Symposium (LARS/SBR)*, pages 1–6. IEEE.

Rezeck, P., Azpúrua, H., Corrêa, M. F. S., and Chaimowicz, L. (2023). Hero 2.0: a low-cost robot for swarm robotics research. *Autonomous Robots*.

Rezeck, P. and Chaimowicz, L. (2022). Chemistry-inspired pattern formation with robotic swarms. *IEEE Robotics and Automation Letters*, 7(4):9137–9144.

Schranz, M., Umlauft, M., Sende, M., and Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7:36.

Tan, X., Xi, W., and Baras, J. S. (2010). Decentralized coordination of autonomous swarms using parallel gibbs sampling. *Automatica*, 46(12):2068–2076.

Xi, W., Tan, X., and Baras, J. S. (2006). Gibbs sampler-based coordination of autonomous swarms. *Automatica*, 42(7):1107–1119.